



The Ethics of Ignorance

Citation

Miller Larsen, Elís. 2023. The Ethics of Ignorance. Doctoral dissertation, Harvard University Graduate School of Arts and Sciences.

Link

<https://nrs.harvard.edu/URN-3:HUL.INSTREPOS:37375766>

Terms of use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#)

HARVARD UNIVERSITY
Graduate School of Arts and Sciences



DISSERTATION ACCEPTANCE CERTIFICATE

The undersigned, appointed by the
Department of Philosophy
have examined a dissertation entitled
The Ethics of Ignorance

presented by Elís Miller Larsen
candidate for the degree of Doctor of Philosophy and hereby
certify that it is worthy of acceptance.

Signature Edward Hall

Typed name: Prof. Ned Hall

Signature [Handwritten Signature]

Typed name: Prof. Susanna Siegel

Signature [Handwritten Signature]

Typed name: Prof. Selim Berker

Signature Eric Mandelbaum

Typed name: Prof. Eric Mandelbaum

Signature _____

Typed name: Prof.

Date: 5/8/23

The Ethics of Ignorance

A dissertation by

Elís Miller Larsen

May 8, 2023

Submitted to the Harvard University Department of Philosophy in Partial Fulfillment of the
Requirements for the Degree of Doctor of Philosophy in Philosophy

© 2023 Elís Miller Larsen All rights reserved.

Dissertation Advisors: Ned Hall, Susanna Siegel

The Ethics of Ignorance

Abstract

Most of us think that being unwittingly ignorant— as we are when we forget a friend’s birthday or remain blind to glaring social injustices—is a bad thing. But if ignorance is so bad, why aren’t we required not to be ignorant? On the standard view that ignorance is the lack of knowledge, ignorance can be no more epistemically evaluable than absences. I develop a theory of ignorance that illuminates and resolves the major paradoxes that arise for the ethics and epistemology of ignorance. What kinds of epistemic shortcomings attach to ignorance? What obligations, if any, do we have not to be ignorant? This dissertation offers answers to those questions.

Part I addresses the paradox of the evaluability of absences: absences are neither good nor bad, rational nor irrational. I first show that there exists an array of counterexamples to the common view that ignorance is the absence of knowledge, with sources in psychology, scientific theory, and everyday life where someone is ignorant because they fail to consider obvious and relevant possibilities. I then use these examples to argue for a substantive form of ignorance that is an attitude like belief. Many take it as evident that ignorance is just the state of not knowing. But I argue that this is incorrect: not knowing is neither necessary nor sufficient for some cases of ignorance. If epistemic evaluability is going to have any bearing, ignorance must be more than *mere* not knowing.

Part II is the test case for ignorance’s epistemic evaluability. I start by addressing a second paradox about ignorance and responsibility. If we follow the standard line that ignorance is the absence of knowledge, then ignorance is not something that we can be

responsible for. Some have confronted this paradox by suggesting that a theory of responsibility for ignorance (esp. ignorance of racial injustices) rests on our social obligations. On such a view, we are responsible for our ignorance just in case it harms others or further perpetuates social inequities. But such a theory suggests that we are only responsible because we have let down our communities since we owe it to those around us not to be ignorant. I argue for a kind of *substantive* ignorance that takes an attitude form where agents are just as responsible for ignorance as they would be for belief. In the dissertation, I call this attitude *ignoring* to distinguish it from the common, passive form of ignorance. Ignoring, according to my view, structures an inquiry so that we are responsible when we fail to consider obvious and relevant answers to an open inquiry. Importantly, the responsibility is agent-centered; it is not based on others, but on the idea that we owe it to ourselves to be as rational as possible.

Part III applies my theory of ignorance to a problem in perception. Hallucinations, illusions, and attentional blindness are all forms of ignorance where perceivers ‘miss out’ on information in their environment. In this chapter I identify another kind of misperception, which I call *perceptual ignorance*. Perceptual Ignorance is the perceptual analogue to the cognitive state of ignoring. When someone is perceptually ignorant, they have literal blind spots that result in errors in early perception, which lead to judgments that are unwarranted. I use Bayesian perceptual processing to model the problem as one of the missing hypotheses within a set of priors. Perceptual ignorance not only explains how ignorance manifests as a type of prejudice, but also presents a surprising critique for Bayesian theories of perception in general. Omissions are evidence of malfunctions in updating, which undermines Bayesian models as paragons of proper functioning.

So, it is up to us to resolve our ignorance, but this task is one that we can all meet. That is part of the larger narrative of an agent-centered ethics of ignorance. Taken together, the chapters are distinct, yet imbricated arguments for a theory of ignorance. This dissertation makes the case that ignorance is the kind of thing that is epistemically evaluable—it can be

epistemically better or worse; and that as better or worse it places demands on us to be our best rational selves.

Por mi Familia.

“Not ignorance, but ignorance of ignorance is the death of knowledge.”

— *attributed to Alfred North Whitehead*

“Imagine an ignorance that resists,”

“Imagine an ignorance that fights back.”

— *Charles Mills, “White Ignorance”*

Table of Contents

Title Page	<i>i</i>
Copyright	<i>ii</i>
Abstract	<i>iii</i>
Table of Contents	<i>viii</i>
Acknowledgements	<i>ix</i>
Preface: The Ethics of Ignorance	1
Part I: Ignorance and Attitudes	3
Part II: Responsibility for Ignorance	34
Part III: Ignorance in Perception	56

Acknowledgements

From an early age I was a philosopher. I just did not know it yet. My initial questioning was always about the etiology of things. “If God created us, then who created God?” In my early years I felt this question must have an answer, so set out to ask my parents, the pastor at my church, the Sunday school teacher, really anyone who would be able to answer me. I was disappointed when I found out that there was no answer to my question. I was even more disappointed that some people were less annoyed with the fact that the question seemed unanswerable (as I was) and more annoyed with me for thinking this is a question we should be asking in the first place. For a while I gave up on philosophy, until early college when these questions were reawakened in me and I found there existed a community of others with thoughts, and questions, like my own.

From that point on, I have been lucky to encounter many great philosophical minds, both on the page and in person. During graduate school I had unwavering support from my first advisors at CUNY Graduate Center. First, I thank Eric Mandelbaum, who never shut down my questions in class and always helped me develop them into a philosophically fruitful project. Jesse Prinz is also noteworthy as an early supporter of my career. I took his advice (too much to heart) that graduate school is a time to take the *most* philosophical risks. This advice fits me all too well as someone who refuses to fit neatly into any one category. From the beginning of graduate school, it seemed as if one always had to choose: “Are you an evidentialist or a pragmatist?” “Are you an internalist or an externalist?” Unlike many of my peers who found one side immediately intuitive and therefore more defensible, I continue to look for third options.

Thankfully, my Harvard dissertation committee has been supportive of this approach. I thank Selim Berker, who is always willing to decipher and dissect any philosophical idea finding its joints and determining deficiencies. I thank Susanna Siegel, who meaningfully referred to me as the ‘queen of questions,’ but who (in my opinion) is the original since her Susanna-questions always leave you pondering for days, weeks, or even months later, leading to a deeper understanding of a concept you initially approached flat-footedly. And finally, Ned Hall, who is always up for developing inchoate ideas; happy to go along for the twists and turns of nascent thoughts while somehow able to clearly visualize the landscape so that the result produces a precise and useful argument.

Next, I can only list the extensive number of people who have supported me along the way either by reading drafts of early work, helping to format the prospectus for this dissertation, encouraging me when I felt that I could not write another sentence, and finally carrying me through the valleys of the job market: Jenna Donahue, Aaliyah Triumph, Kate Nolfi, Rima Basu, Brianna Toole, Justin Clary, Fiery Cushman, Caitlin Fitchett, Elizabeth Miller, Subrena Smith, Ronni Gura Sadovsky, Wendy Salkin, Lauren Layden-Hardy, Olivia Bailey, Ariana Fabio, David Thorstad, Jada Wiggleson-Little, Matthew Lyskawa, Tyler Doggett, and Zoë Johnson King. I also want to especially note the administrators of the Harvard philosophy department. Emily Ware and Nyasha Borde Bovell often went above and beyond the duties of finding a file or answering a query. If you have ever had to herd cats, I am the cat. I am grateful to those who are willing to do the herding.

Finally, I thank my family. I would not be the thinker I am without our rigorous debates (sometimes executed with complete disregard for politeness or conversational etiquette), including our most recent discussion on the *Familia* group chat about the nature of bias. To my husband, Cory, there are barely words to describe our surprising bond—surprising because neither of us thought we would do well on a team (both independent to a fault) and have somehow created a peaceful and loving home for our pomsy pups (Koda and Nala) and our son (Soren).

Soren. Named after the great philosopher who taught me that no 'sickness is unto death.'
Academic philosophy is no easy path, neither (as it turned out) was bringing you into the world.
But I have both now, and I am grateful that no sickness—no obstacle material or physical—
obscured this beautiful end. Whenever you ask, "What's that?" I can see you aiming for the true
nature of things. Thank you for showing me how to categorize the world anew.

Preface

The Ethics of Ignorance

Like the ethics of belief¹, the ethics of ignorance is a set of questions concerned with the norms or standards that govern the synchronic and habitual manifestations of ignorance. The ethics of belief is concerned with the norms or standards that govern synchronic and habitual manifestations of belief (both implicit and explicit), which generate questions at the intersection of epistemology, psychology, philosophy of mind, and ethics²: *Is it wrong to hold a belief on insufficient evidence? Is it ever irrational to believe based on statistical generalizations? Just as we are morally responsible for actions, are we morally responsible for beliefs? Are we also ever epistemically responsible for our beliefs?*

The ethics of ignorance generates a similar yet distinct set of questions: *Is it always wrong (epistemically and/or morally) to be ignorant? Is it ever rational to be ignorant? Are we epistemically responsible for our ignorance?* Answering these questions also requires research at the intersection of epistemology, psychology, the philosophy of mind, and ethics.

¹ See W.K. Clifford's *The Ethics of Belief*, originally published in 1877.

² For a concise survey, see the Stanford Encyclopedia of Philosophy entry "The Ethics of Belief" by A. Chignell, revised March 5, 2018.

Part I

Ignorance and Attitudes

“It is very difficult to find a black cat in a dark room,” says an old proverb.

“Especially when there is no cat.”

1. Introduction

Ignorance is a hot topic, yet often misunderstood. To be ignorant, according to the prominent views in contemporary epistemology, is to lack a target cognitive epistemic state. The *absence-schema* for the nature of ignorance is expressed in several varieties. There is the standard view, according to which a person is ignorant when they lack knowledge.

Standard View: an epistemic agent is ignorant of the fact that p at t just in case the agent *does not know* that p at t .³

There is also the new view where a person is ignorant when they lack true belief.⁴

The New View: an epistemic agent is ignorant of the fact that p at t just in case the agent *does not have a true belief* that p at t .⁵

And finally, there is the awareness view where a person is ignorant when they lack awareness of a fact.

The Awareness View: an epistemic agent is ignorant of the fact that p at t just in case the agent *is unaware* of the fact that p at t .⁶

Lately many philosophers have also emphasized the *epistemic significance* of ignorance, arguing that ignorance matters.⁷ By epistemically significant, these authors mean that ignorance has a normative component where a person *ought not be ignorant*, which is governed by standards of

³ See Fine (2018), Haack (2001), Le Morvan (2011), Zimmerman (2008).

⁴ The debate is about the status of lucky true belief. For more on the debate see Le Morvan & Peels (2016).

⁵ See Goldman (1986), Goldman and Olson (2009), Guerrero (2007), Mills (2007), Peels (2010), van Woudenberg (2009).

⁶ See Pritchard (2021) (2022), and Silva & Siscoe (2023).

⁷ See Alcoff (2007), Fricker (2016), Mills (2007), Medina (2013) (2016), Pritchard (2021), Silva & Siscoe (2023), Wang & Wang (2022).

rationality in contemporary epistemology.⁸ But by remaining true to the idea that ignorance is the absence of an epistemic state—either the *absence* of knowledge, the *absence* of true belief, or the *absence* of awareness— none of these views can provide what we need for expressing its epistemic significance. On the standard view that ignorance is the lack of a cognitive epistemic state, ignorance can be no more epistemically evaluable than absences. The *absence-schema*, therefore, presents a puzzle for the nature of ignorance and its epistemic significance: either ignorance is an absence so that it cannot be epistemically evaluated, or ignorance is not a mere absence, and we must give up the standard conceptions. In this paper I take the second approach.

So, what if ignorance is not an absence? ‘Dark’ is the absence of ‘light.’ ‘Silence’ is the absence of ‘sound.’ ‘Blindness’ is the absence of ‘sight.’ These negative concepts pairs (‘dark’/ ‘light,’ ‘silence’/ ‘sound,’ ‘blindness’/ ‘sight’) are oppositional while also expressing *privative* notions— concepts that lack some quality or attribute. Not all oppositional pairs will fit this category, however. ‘Drought’ is not the absence of ‘rainfall’—it may also signify low rainfall or a shortage of rain for a region. ‘Hunger’ is not the absence of ‘satiation’—to be hungry is to experience a lack of food rather than to lack satiation. And surprising, ‘bad’ is not the absence of ‘good.’ To be bad is not to be *non*-good but consists in its own unique properties that describe ‘badness’ apart from what it means to be good (e.g., to be unwelcome or unpleasant). The reverse is also true. If you are not a ‘bad’ driver, it does not mean that you are a good one. Similarly, if you are not a ‘good’ driver, it does not mean that you are a bad one—you could be mediocre, or moderately competent. The relation between ‘ignorance’ and ‘knowledge’ is better suited within this second category of oppositional pairs that include ‘drought’/ ‘rainfall’, ‘hunger’/ ‘satiation,’ and ‘bad’/ ‘good.’

⁸ See, for example, Pritchard (2021) for an argument that the normative component should be constitutive of the nature of ignorance given its importance in epistemology.

The purpose of this paper is to argue that previous conceptions of ignorance are mistaken. Ignorance is not (always) the *absence* of a cognitive epistemic state. Rather, it is (sometimes) an *attitude* that constructs a person's inquiry space through the possibilities that occupy them. Although there is a longstanding tradition of viewing ignorance as an absence, novel conceptions should be considered given the aim of providing a theory of the nature of ignorance that is epistemically significant.

2. The core cases

Unlike *mere* absences, ignorance constructs the space of possibilities. For example, when a judge strikes a statement from the record, it can no longer be 'known' by the jury, and the removal of the statement as evidence for the case bears on the jurors' potential deductions. A second example, provided by Szegedy et al. (2015), involves the application of 'negative' space in the form of a static image that alters the outputs of deep neural networks. The static should have no impact on the networks ability to classify images, yet when applied, it alters the outputs of the system from a 57.7% confidence rate that the image is a panda (the correct output) to a 99.3% confidence rate that the same image is a gibbon (the incorrect output).⁹

For a third illustration, Bar Hillel et al. (2018) describe multiple cases of failing to consider relevant possibilities across content. When participants are given *insight problems*—problems that require the examinee to shift their perspective—, they fail to register certain possibilities that would improve their chances of arriving at the solution. Let us suppose a participant, we will call him Joaquin, is given the following insight problem: an accountant says, 'That attorney is my brother,' and that is true, they really have the same parents. Yet the attorney denies having any brothers, and that is also true! How is this possible?¹⁰ Joaquin cannot answer the riddle. It is likely that he, along with the other participants, are blinded by gender stereotypes

⁹ To clearly see the connection, the reader should think of 'negative' space as the equivalent of ignorance for a neural network.

¹⁰ See Bar Hillel et al (2018).

that prevent him from considering the possibility that the accountant is the sister. His bias is said to be so strong that it overrides the plausibility of the correct answer in favor of less likely ones.¹¹ It should be clear that Joaquin is not ignorant because he simply lacks the target epistemic state. His ignorance is due to a deeper cognitive mistake that causes unstable reasoning patterns. These cases suggest the following idea. When certain possibilities are omitted in the space of a person's reasoning, the person manifests a type of ignorance distinct from the mere absence of a cognitive epistemic state.

How should we make sense of such cases? Is it that the people (or neural networks) in these examples lack the relevant epistemic state—either knowledge, true belief, or awareness? Or is there another kind of ignorance in addition to the familiar, *privative* form?

Consider another armchair case.

Who's the Professor? On the first day of classes, Finn—a usually excellent student—is late to his intro to philosophy course. Thinking he has already missed valuable information; Finn frantically searches the room for the course professor. He considers several people in the room, including a male student in the back wearing Oxbridge-like attire. However, Finn fails to consider the woman at the board, despite her performing duties that would normally indicate she is the professor.

Finn is, undoubtedly, in a state of ignorance. But the state is peculiar given the fact that he has access to evidence that would resolve his inquiry. To be specific, he experiences the woman writing on the board and performing professor duties (**P**) but fails to infer that she is the

¹¹ Some participants preferred the possibility that the *accountant had died*, over the thought that the accountant is the sister. Or they rationalized that the salutation brother has multiple meanings to account for why the riddle seemed unsolvable. For more participant responses see Bar Hillel et al. (2018).

professor (Q), omitting the possibility (Q) altogether. In so doing, Finn fails to make the proper inference ($P \rightarrow Q$), which would be a perfectly good deduction given his available evidence.

Is Finn's mistake an issue of legibility? While legibility cases abound¹², Finn's mistake is not about accessibility to evidence under the deception of evil demons¹³, brain machines¹⁴, or bad ideology¹⁵. We can imagine an ordinary person with fully optimal faculties who unwittingly makes a genuine error. For those who remain unconvinced of the plausibility of the phenomena, consider a version of the case that casts Finn's attitudes in terms of belief. Suppose that Finn has a belief that the woman simply cannot be the professor. He thinks, everyone knows that *women are not university professors*.

There are a few things worth saying about such a case. First, the reason to view Finn as ignorant in this version may have nothing to do with lacking a cognitive epistemic state, and everything to do with how it offends our egalitarian beliefs. For these reasons, Finn is ignorant *simpliciter*—a deeply ingrained ignorance illustrative of his defective character.¹⁶ In addition, in this version of the case his ignorance is not the absence of an epistemic state, but rather a faulty epistemic state in the form of a false belief. But now ignorance is epistemically significant when it reflects the character of the ignorant person, and when the *content* of an (ignorant) belief is one that does not match egalitarian standards. Such a view is tempting but should be resisted.¹⁷

¹² See, for example, the externalist/internalist debate. Externalists argue that 'legibility' is a matter of external facts, while internalists argue that 'legibility' is a matter of internal facts about which the person has (reflective) access.

¹³ For a collection of essays on contemporary perspectives on the evil demon problem see Dutant (forthcoming).

¹⁴ See, for instance, Putnam's (1975) brain-in-a-vat and Nozick's (1981) experience machine.

¹⁵ See Srinivasan (2021) for an externalist defense of justified belief in the face of bad ideology and see Johnson King (2022) for an internalist defense of justified belief in the face of bad ideology.

¹⁶ I'll discuss character more in depth in § 8.

¹⁷ Mills (2007) view seems to accord with these conditions: ignorance is epistemically significant given its negative effect on the character of white Americans and its description as 'racial insensitivity'.

One final example will illustrate this point.

Changing Shape: Beverly, a participant in a change blindness study, is tasked with sorting shapes by size. She knows that large shapes should be stacked in one pile and small shapes stacked in another. But when it comes to performing the task, Beverly cannot register the fact that shape sizes change mid-task, despite seeing them do so.¹⁸

Changing Shape is not a classic inquiry case, but it can easily be construed as one. If we think of perception as a kind of ‘inquiry to the world’ we can imagine that every time Beverly looks at the pile of shapes, she is making an inquiry about their size, which informs how she categorizes them. Like Finn, Beverly is in a state of ignorance. Unlike Finn, however, it is not intuitively the case that her ignorance is due to defective character or a false belief about a marginalized group. So, there is no intuition that her ignorance arises as part of the content of her attitudes. Instead, I will suggest that Beverly is ignorant because she fails to consider a relevant possibility—*the possibility that shapes sizes change*. Her mistake is less obvious than Finn’s but is in the same category of epistemic mistakes given her inquiry task.

Is someone ignorant just in case they have a failed inquiry? No. In fact, accounts that equate ignorance with failed inquiry are unable to provide a full normative spectrum for the assessment of ignorance—one that extends in both positive and negative directions.¹⁹ For one, an inquiry may be defective yet result in an epistemic good (either knowledge or true belief).²⁰ Additionally, an inquiry may be good yet not guarantee an epistemic good. To understand this point, consider Jack, a census worker who is inquiring into the number of people living in an abandoned house.²¹ The goal of Jack’s inquiry is to accurately count the number of people who

¹⁸ Based on the Triesch et al. (2013) study.

¹⁹ Both Pritchard (2021) and Meylan (2022) are committed to the idea that ignorance is only *negatively* normative.

²⁰ See Wang and Wang (2022) for a version of this objection for Pritchard’s view of ignorance as failed inquiry.

²¹ See Guerrero (2007) for the original case.

live in a town. Suppose that Jack concludes that there are zero people living in the home and does so on the basis of his reasoning that people in the town do not usually live in abandoned buildings. While conducting his inquiry, Jack does not double-check his conclusion by knocking on the building door, nor does he interview any of the neighbors. It turns out that Jack is correct! —the house has no occupants. But this outcome is not the result of good inquiry. There is a sense in which Jack remains ignorant (since his belief does not enjoy a justified basis), even though his inquiry is a ‘success.’ Conversely, Jack may manifest good inquiry, exemplified by double-checking and interviewing neighbors, and still arrive at a false belief. His good inquiry may lead him to believe that there are no occupants in the house when there, in fact, are. So, good inquiry can leave him unaware of a truth about which he ought to be aware. And bad inquiry can result in a true belief without a justified basis. Inquiry at its best, then, does not safeguard against ignorance. There is much more to be said about what good inquiry amounts to²² and its relation to ignorance, but it would be impossible to take up all these issues in this paper.

For now, it is more important to say something about the nature of ignorance as distinct from good inquiry. Finn and Beverly’s mistake is a failure which is *part* of the inquiry process, but not constitutive of that process. While conducting the inquiry they exhibit an omission that yields unstable reasoning patterns (e.g., a failed inference). Their omissions are epistemic in the sense that they are part of the decision-making process. So, the epistemic significance of ignorance relies on the impact it has on reasoning that is efficacious in ways that *absences* are not. Our core cases lead to the main proposal of this paper. Ignorance is an attitude that constructs a person’s inquiry space through the possibilities that occupy them. The upshot of the attitudinal account of ignorance is that it shows how ignorance is not only *criticizeable* (for we can criticize absences as problematic in some way or another, often according to moral

²² See Friedman (2022) for a start.

standards), but, more importantly, how ignorance can do *epistemic work* by constructing the space of possible thought.

3. The attitudinal account of ignorance

Ignorance, I argue, is a *sui generis* attitude²³ :

The Attitude View: an epistemic agent is ignorant of the fact that p at t just in case the agent ignores that p at t by failing to consider obvious and relevant possibilities in inquiry.

But what is an attitude? Explanations from psychology treat attitudes as *likings* or *preferences*, whereas explanations from philosophy see attitudes as mental states in relation to propositional content. I will prefer the philosophical explanation that attitudes are relations to propositional content, although I am open to the idea that attitudes can have different content structures (propositional, associative, or iconic).²⁴ Specifically for my purposes, propositional content formulations will be useful models for expressing the ignorance attitude in relation to the omitted target information in sentential form: Finn has an ignorance attitude in relation to *the woman is the professor*. This is not to say that the psychological explanation of attitudes has no explanatory significance for our core cases. A fitting psychological explanation of Finn's ignorance may be that he has an implicit bias that impacts his decision-making. The bias would consist of the belief (unbeknownst to him) that women cannot be professors. And the bias,

²³ By *sui generis* I do not mean that ignorance is so unique that it is irreducible to any other descriptions, such as not considering or omitting. I simply mean that ignorance is an attitude distinct from other familiar states, such as knowing, believing, desiring, loving, and the like. Thanks to Angela Mendolovic for drawing my attention to this important distinction.

²⁴ See Quilty-Dunn (2020) for a pluralist account of content for perception. There may be reason to think that some attitudes also have various representational content, including image-like or iconic content. It is worth noting that not all content types will be easily compatible with reason-responsiveness. It is not clear that iconic content is reason-responsive at all. And associative content may be more insulated from reasons than propositional content (see Mandelbaum 2016).

albeit non-conscious, may then impact his ability to accurately ‘read’ the world, causing blindness to available evidence witnessed in the case.²⁵

Such an explanation, while compelling, cannot explain all our core cases. *Does Beverly also have an implicit bias?* To date, biases mostly have explanatory force for socially driven ills against marginalized groups based on race, age, credence, sex, religious affiliation, and so on. To count as (implicitly) biased Beverly would need to have an implicit belief that shapes cannot change size for which she is completely unaware of. And her belief would not only impact her present decision, but also any future decisions with similar content. Unlike other implicit biases that are stubborn and often immutable despite attempts to overturn them, it is reasonable to think that we could train Beverly to recognize shapes changing mid-task. So, while all implicit biases may include some form of ignorance (absence or attitude), not all cases of implicit bias will count as ignorance. Moreover, it is not evident that all cases of attitudinal ignorance will count as an implicit bias. This difference is especially important for the normative dimension of ignorance. Implicit bias, as discussed across philosophy and psychology, is a pejorative phenomenon. Ignorance catches our attention because of its pejorative nature, but it would be too quick to conclude that ignorance is pejorative across all cases.

There is more to be said about the plausible psychological explanations for ignorance (some that come to mind include cognitive dissonance, associative reasoning, fragmented belief, belief perseverance, and base-rate assumptions), but instead of attempting to eliminate every possible psychological explanation that one could use to account for ignorance, I will turn instead to the philosophical explanation for attitudes. There might be many psychological reasons why ignorance crops up (e.g., cognitive dissonance might beget ignorance). The analytical perspective provides a unifying account of ignorance across these reasons for why ignorance

²⁵ See Johnson (2020), Siegel (2016), and Banaji & Greenwood (1994) for support for this kind of explanation.

arises. Thinking that ignorance is reducible to a single psychological explanation would be mistaken.

3.1 Attitudes

Ignoring is the state of having an ignorant attitude in relation to omitted target information. Having introduced the idea that ignorance can be an attitude, I will now provide some characteristics. For simplicity, I have called the relevant attitude ‘ignoring’ and have said that a person has an ignoring attitude when they omit target information that is available and relevant to their inquiry. Thus far, we have already identified one characteristic of ignoring: ignoring is partially constituted by an *omission*. Notice that there is an inactive and active sense of ‘omission.’ It can be understood as items that have been excluded. Or, it can be understood as an act of exclusion, signally an agent-centered property. The inactive sense of omission may be too weak for our purposes, while the active one may be too strong. Bear et al. (2020), for example, argue for an account where mental omissions are simple exclusions that occur for any number of categories (restaurants, pet preferences, or favorite novels).²⁶ When someone excludes a restaurant from choices that ‘come to mind,’ there is no normative expectation that they ought to have considered other choices. This is where the explanation is weak. The active sense of ‘omission’ while promising is too strong. The *act* of excluding builds into the explanation the moral obligation that the person *should not* have excluded the item. The exclusion then is illustrative of an intellectual vice, such as negligence or forgetfulness. For the attitude view, we should consider an understanding of ‘omission’ that is in between the inactive and active sense, so that the omission is agent-centered in the sense that it can be attributed to a subject without the normative component being settled by the mere fact of omission.

²⁶ See Bear et al. (2020) See also Phillips and Cushman (2017).

Further characteristics can be defined as follows. The attitude is (a) representational—the omission is represented in the person's space of mental reasoning; it is also often (b) systemic (tendencies to be ignorant with respect to content is repeated across instances and is habitual); and is (c) action-guiding by impacting immediate behavioral responses, as well as future reasoning.

So, what is it to say that ignorance is no longer the mere absence of a cognitive state? To date, we simply do not have a unified account of the properties of attitudes.²⁷ Instead, one gets their information from certain background theories that might magnify the psychological consequences of attitudes. For example, functionalism about mental states is a wonderful way to make sense of the normative dimension of attitudes in general, but particularly, ignorance. 'Function' is a normative conception, at least in contrast to malfunction, so that an attitude's functioning beautifully illustrates a success condition of that attitude (e.g., belief aiming at the truth), while its malfunction is illustrative of error (e.g., a false belief that *p*). For functionalists, something is an attitude given its typical, or causal relations to sensory stimuli, behavior, and/or other mental states.²⁸ Such relations can be *backward-looking*, articulating *the cause* of the attitude (e.g., directing perceptual attention can cause a belief that *p*), or the relations can be *forward-looking*, articulating effects *caused by* the attitude (e.g., believing that *p* causes her to make an assertion <*p*>). Functionalist explanations that are backward-looking will not work as explanations for attitude ignorance. It does not seem true that directing perceptual attention *causes* ignorance, rather the opposite—failing to direct attention to perceptible properties leads a person to be ignorant.²⁹ Forward-looking relations are more promising. Ignoring that *p* will

²⁷ One exception may be Loar's (1981) outline of attitude properties as functionalist relationships.

²⁸ Credited sources include Armstrong (1968), Fodor (1968), Lewis (1972) (1980), Putnam (1975), and Block (1978). Other version of functionalism will not be directly discussed but may have something to contribute to defining an account of attitude ignorance. For example, psychofunctionalism is the idea that what makes something an attitude or process is the role it plays in the cognitive system. There is certainly a connection between attitude ignorance, as I've defined it, and the kind of psychological functions that may be caused by it, which require further investigation.

²⁹ This point may just signify a different functional relationship from the one given for belief. But it is at least clear that not all functional relationships, which already exist for belief will map on nicely to attitude ignorance.

typically lead to an incomplete inquiry into Q , where $\langle p \rangle$ is a possible or relevant answer to Q .³⁰ The causal relationship between ignorance and inquiry expresses just one example of the kinds of functional relationships that may exist for attitude ignorance.

An easier approach—one that does not require defining causal relations—may be found by applying dispositionalist or interpretationist accounts of mental states. According to dispositionalism, to have an attitude (e.g., a belief that p) is to possess one or more behavioral dispositions pertaining to p .³¹ One can see how a dispositionalist approach is evident in Mill's (2007) discussion of ignorance.

“White ignorance is best thought of as a cognitive tendency—an inclination, a doxastic disposition” (Mills 2007: 23).

While Mills account provides a compelling case for the existence of a particular radicalized ignorance that affords resistance to knowledge of structural injustices, the distinctive cognitive profile of ignorance is unclear. A dispositionalist reading of Mills would explain white ignorance as the behavioral disposition of white Americans to ignore certain facts and willfully refuse to learn American history that would reveal its oppressive past.³² One version of dispositionalism—interpretationalism—would be particularly apt for this task. For interpretationists, attitudes are patterns of action and reaction that manifest as observable behavior. So, seeing (e.g.) a white person get up from their seat on the subway and move away from a black patron who comes to sit next to them would illustrate patterns of action and reaction that reveal a mental life riddled with white ignorance.

³⁰ I use an italicized ‘ Q ’ to denote a question in inquiry and a bold ‘ Q ’ to denote the conclusion of an inference.

³¹ Some contemporary advocates of the view include Schwitzgebel (2002) (2013), Price (1969), Baker (1995).

³² Note that a dispositionalist account of mental states is different from dispositional belief. It could also be the case that Mills (2007) is describing a kind of *dispositional* ignorance that exists as part of a person's background mental life, as contrasted with their *occurrent* beliefs that would express immediate (and perhaps available) mental reflections.

One of the main setbacks for dispositionalism is that it cannot tell us about attitudes that do not attach to behavioral profiles (e.g., implicit attitudes).³³ For our purposes, observable behavior should not play such a prominent role in identifying the attitude. *Would it matter if someone had an ignorance attitude, yet had no outward manifestations of ignorance?* Given our core cases, I think that a definable attitude would still be present if, for example, Finn demonstrated no outward actions or reactions to his ignorance of the woman as the professor. Unwitting ignorance, without being attached to behavior, can nonetheless result in unstable processes for current and downstream thought.

3.2 Is ignorance a state or a process?

Without determining the perfect background theory of mind to support attitude ignorance, identifying properties nevertheless exist that can illustrate the differences between attitudes and other mental activities, such as processes. Attitudes as states are thought of as *static* properties, while processes are thought of as *dynamic* properties. To say that a property is static, such as ‘being red’ or ‘knowing that p’ or ‘owning something,’ is to say that it has a perfective aspect, which is the fact that something is *completed*. Conversely, to say that property is ‘dynamic,’ such as ‘buying’ or ‘starting’ or ‘shaking,’ is to say that a property has yet to be completed—the property is still in the *process* of completion. Notice that attitudes—knowing, believing, loving—are said to be static.³⁴ So, saying that ignorance is an attitude will make it seem as if it is a static cognitive state.

The attitude view of ignorance claims that ignorance is part of the inquiry process, which is where it stops being a *mere* static cognitive state and takes on a new characteristic. Ignorance is

³³ See Quilty-Dunn & Mandelbaum (2018) for a compelling refutation of dispositionalism.

³⁴ For support, see Setiya (2013) and Friedman (2022).

also not the only attitude that is not static. There is a lively debate about beliefs as static or dynamic. Setiya (2013), for example, defends the idea that believing (despite having the active -*ing* suffix) is static, even though it can be done for reasons.³⁵ Contra Setiya, Hieronymi (2011) thinks that believing is also done for reasons, but it is not completed. It is evident that our beliefs evolve—we can acquire new beliefs and we can revise old ones—so beliefs do not appear to be static in the sense that they are completed.³⁶

Another distinction between *telic* and *atelic* properties can be applied to the assessment of attitudes to help make sense of their metaphysical status. Some attitudes, like belief, appear to have an aim or *telos* that is constitutive of the practice. For example, it is argued that the aim of believing is to have true beliefs. Telic properties (like static ones) have an endpoint in the sense that they are completed, but the completion has a teleological structure that includes, for instance, a final epistemic value. Discussion of the properties of states, and what it means to say that a property is static has roots in Vendler's (1957) categories that crosscut the two distinctions I have discussed in consistent ways. Vendler categories include *states* (knowing, believing, loving), *activities* (running, pushing a cart, celebrating), *accomplishments* (running a marathon, painting a picture), and *achievements* (winning, recognizing, reaching the top). It is no surprise that states, accomplishments, and achievements are both *static* and *telic* since they are all categories of completion. While it is not the case that all states, accomplishments, or achievements must be *telic*, although many tend to be (e.g., it is constitutive of the practice of winning to actually win).

³⁵ The debate actually turns on this claim about beliefs being 'done for reasons' since it articulates the deliberative aspect of belief, which appears dynamic in form.

³⁶ There remains a question about the temporal constraints on belief. When are evaluating a belief? At any time, *t* it might be the case that the chose time slot reveals a belief that is 'completed'—at that time there is nothing further occurring as far as the development of the belief. But the intuition that beliefs change continues to weigh on us. The tension between the idea that we evaluate beliefs at a given time and the intuition that beliefs clearly change over time, illustrates the limitation of a synchronic evaluation of beliefs. The tension is more pressing when it comes to evaluating processes that are part of inquiry.

There are, however, types of activities that do not fit our neat categorization. The idea that a person is going to ‘quit smoking’ is one such case. To say that a person ‘quit smoking’ is an activity—it a process or journey that the quitter undergoes. The journey may also have accomplishments and achievements in the sense that they are quitting (which is an accomplishment) and they may have quit for several months (an achievement). So, to quit smoking appears both static and dynamic. I suspect that there are other cases that do not fit the static/dynamic framework as well. And my suspicion is that ignorance is one of these exceptions. To say that ignorance is a state conforms to the idea that it is static (‘static’ being a cognate of the word state). But what it means to *say* that something is a state differs from what it *is* for something to be a state. In his arguments, Setiya claims that our linguistic practices directly connect to the metaphysical joints of the nature of attitudes. But there is no reason for us to immediately accept Setiya’s direct metaphysical connection from language.³⁷

Ignorance is not a *mere* static cognitive state, but it is not quite a process either. As I have claimed it is *part* of the inquiry process—it does not consist of the inquiry process itself.³⁸ Like Hieronymi (2011) the attitude of ignorance is dynamic in the sense that it is not completed (although it may be done for reasons)³⁹. When people have an ignoring attitude, they have ignored possibilities by (a) representing an omission as part of their decision-process within inquiry. Unfortunately, their omission may take on an (b) habitual nature where it is repeated across similar content, suggesting a blindness to certain stimuli. Think of Mills white ignorance, for example, manifestations of ignorance are not likely one-off. Rather, the white ignorance is repeated across many cases of similar content. Finally, I have argued that the attitude is (c) action-guiding since it constructs the space of possibilities in inquiry, which yields unstable reasoning patterns for present and downstream thought.

³⁷ I will hedge here and say that there is a reason based on *precedent*. Setiya (2013) cites the longstanding tradition that grounds the metaphysics of properties being static versus dynamic, but I am not one to base my philosophical commitments on precedent alone, even good precedent.

³⁸ This will be one of the major divergences between my account and Pritchard (2022).

³⁹ see § 5.

4. A linguistic defense of the attitudinal account

Suppose we grant that ignorance is an attitude. Our proposal thus far does not preclude ignorance from also being the absence of an epistemic state. When a subject has an ignoring attitude, they omit target information (as characteristic of the attitude), which could also manifest, in part, as a lack of knowledge. There are many reasons to hold tight to the old conception of ignorance. For one, ignorance comes from the Latin word 'ignorantia' which means 'not knowing.' So pairing ignorance with a description of not knowing is a natural move. There are also descriptive similarities for linguistic comparisons of 'ignorance' and 'not knowing,' as evident by the following semantic descriptions for our core cases:

Joaquin is 'ignorant' of the fact that the accountant is the sister.

Finn is 'ignorant' of the fact that the woman is the professor.

Beverly is 'ignorant' of the fact that shapes change size.

For the standard view we get:

Joaquin does 'not know' the fact that the accountant is the sister.

Finn does 'not know' the fact that the woman is the professor.

Beverly does 'not know' the fact that shapes change size.

For the new view we get:

Joaquin 'lacks a true belief' of the fact that the accountant is the sister.

Finn 'lacks a true belief' of the fact that the woman is the professor.

Beverly 'lacks a true belief' of the fact that shapes change size.

Finally, for the awareness view we get:

Joaquin 'lacks awareness' of/is 'not aware' of the fact that the accountant is the sister.

Finn 'lacks awareness' of/is 'not aware' of the fact that the woman is the professor.

Beverly 'lacks awareness' of/is 'not aware' of the fact that shapes change size.

So, the privative descriptions are acceptable semantic substitutions for the term 'ignorance.' *But are the privative descriptors explanatorily sufficient for our core cases?*

Proponents of the awareness view, for example, argue that 'awareness' is a superior description of ignorance given the semantic similarities between saying that a subject is 'ignorant' of the fact that p and saying that a subject is 'not aware' of the fact that p .⁴⁰ Even with these semantic similarities, substituting the descriptors of each view does not explain how Joaquin, Finn, or Beverly are ignorant. Take Joaquin for example. It is true that he lacks knowledge/true belief/awareness that the accountant is the sister. But the absence of these states does not explain his ignorance. In fact, just considering insight problems, it is the case that subjects can know that p while remaining ignorant of the fact that p : Joaquin can know that woman are accountants and remain ignorant of the fact that the sister is the accountant. There is some dispute given these remarks about what exactly Joaquin is ignorant of. *Is he ignorant of the general fact that woman are accountants?* It seems not. *Or is he ignorant of the fact that the particular woman is the accountant?* At most, the privative accounts can provide a descriptive linguistic analysis for why Joaquin is ignorant of the specific fact. What remains is an analysis

⁴⁰ See Silva & Siscoe (2023).

of the disconnect between his general knowledge and the specific tokening of that knowledge. Importantly, Joaquin's mistake seems quite different from (say) forgetting. Imagine an armchair version of the example where Joaquin just forgets that woman can be accountants even though he knows this fact all along. When we forget we may have a fragmented or divided mind where we know on the one hand but fail to know on the other.⁴¹ This could partially explain Joaquin's ignorance. Such an explanation would still put pressure on the biconditional relation between ignorance and not knowing. It would suggest that to count as ignorant it is not the case that someone needs to lack knowledge, where lacking knowledge is when the target information is absent from their entire mental database. Lacking knowledge is just one *absence* input, but the point stands for the absence-schema in general.

The resulting nature of ignorance is one where the absence of a cognitive state is neither necessary nor sufficient for a subject to count as ignorant. The new conception provides an explanatorily proficient account of ignorance in linguistic terms. And although it introduces a novel account of the nature of ignorance, it remains compatible with the active form of the Latin root, *ignorare*, which can mean 'to not know,' but also means 'to ignore.'

5. Does the nature of ignorance (prima facie) exclude reason-responsiveness?

Absences are not reason responsive. So, ignorance—as an absence— excludes reason-responsiveness: If an epistemic agent is ignorant of the fact that p at t , then the agent cannot respond to the fact that p at t as a reason to believe that p at t . Two cases are taken as supportive of this claim. The first is a version of Nozick's (1981) experience machine as discussed by Clayton Littlejohn.⁴² While hooked to the machine, Agnes is said to have had a

⁴¹ For example, see Eagan (2007) and Bendana & Mandelbaum (2021) on fragmented belief.

⁴² See Littlejohn (2017).

veritic (yet manipulated) experience of her brother walking across the stage and smiling at her during his graduation. Unbeknownst to Agnes this event takes place in the ‘outside world’ at the very same time she experiences it while connected to the machine. Her machine-experience, however, cannot count as a belief because she is ‘cut off from reality.’ Littlejohn takes the phrase ‘cut off from reality’ as indicative of her ignorance, so that she is said to be ignorant of the fact her brother is walking across the stage. Agnes’ ignorance excludes her from being able to respond to the fact that he is walking across the stage as a reason to believe that fact. The armchair example is meant to be illustrative of how neither the standard view nor the new view can address the disconnect between Agnes’s seeming to have the target belief while being simultaneously ignorant given the fact that she is ‘cut off from reality.’

The second case that is meant to support the claim that ignorance is not reason-responsive is a possible case developed by proponents of the awareness view. They suppose that Felice, a wary traveler, is worried about flying across the country.⁴³ To allay her worries, Felice conducts research on safety guidelines and protocols for airlines in the United States and discovers that flying is relatively safe. However, despite reading this evidence, Felice is unable to believe that her flight does not pose a significant risk. The awareness view is meant to have a superior explanation for why Felice is not ignorant, which rivals the standard and new view. Unlike Agnes in the experience machine, Felice is not ‘cut off from reality’ so that she is aware of the evidence that would lead her to believe that her flight is safe. Her lack of belief, proponents of the awareness view argue, is not a signifier of ignorance—she simply is not able to believe at will. She is, however, able to respond to reasons in ways that Agnes is ‘unable’ given her state in the experience machine. Felice can buy tickets or address her fear to others.

The attitude view can also address the second case but takes a different approach to the reason-responsiveness claim. The reason that Felice does not count as ignorant on the attitude

⁴³ See Silva & Siscoe (2023).

view is because she considers the relevant and obvious possibilities when conducting her inquiry into the safety conditions of the flight. The attitude view does not require that she has a resulting belief, true or otherwise, nor does it require that her inquiry results in knowledge.

The armchair case of the experience machine is simply beyond the purview of the attitudinal account. For one, it is not evident that Agnes is conducting an inquiry where she has yet to settle on what to believe. She *has* settled and believes that her brother is walking across the stage. There exist separate analytic tools to further address Nozick's case that do not bear on the nature of ignorance, and I will not unpack those tools here.

For our purposes, there is an important connection between reason-responsiveness and the empirical case of insight problems. When Joaquin does not know that the accountant is the sister, he cannot respond to this fact as a reason to believe that she is the accountant. So, privative ignorance *does* exclude reason-responsiveness. Similarly, one might think that attitude ignorance excludes reason-responsiveness as well. So, either we need an argument for why attitude ignorance is reason-responsive, or we must accept another account of ignorance where it is not (the standard view, new view, or awareness view).

Understanding attitude ignorance as responsive to reasons requires the introduction of a distinction between *being able to respond to a fact for a reason* (e.g., the reason I got on the plane is because of the fact that it was boarding) and *having an attitude that is done for reasons* (e.g., I believed that he was a good father because of his attentiveness to his child). The latter grouping does not require that the attitude be something we must be aware of. I can experience fear because an arachnid crawls across my foot. Or I can feel ashamed because I am a terrible cook. Both fear and shame can be done for reasons without my knowing that I am in the epistemic state of 'being fearful' or 'being ashamed.' Moreover, something can be responsive to

reasons without being an epistemic state at all.⁴⁴ On such proposals reason-responsiveness is about how something operates within rational processes. Ignorance is operative given the epistemic impact of failing to infer Q from P. So, even if one thinks that ignorance itself is not responsive to reasons because the person does not know, or is unaware of, the reasons that they are meant to be responding to, it is still evident that ignorance impacts deductive (and abductive) reasoning processes.

6. A worry for normatively significant views of ignorance

There is a worry for all views of ignorance, including my own, that ignorance can only be *negatively* normative. Some accounts will be comfortable with this consequence. Meylan (2020), for example, argues that ignorance is bad because ‘it is the manifestation of a vice’. Meylan points out that the kind of badness at stake is *epistemic* and *non-instrumental*. We can think of many cases of instrumentally good ignorance where there are *pro tanto* reasons to be ignorant (e.g., avoiding knowledge of bad outcomes to which we have no control). Meylan thinks instrumental considerations do not take away from the fact that there is “something *prima facie* bad about being ignorant” (Meylan 2020, 436). ‘Badness’ here takes on an epistemic sense, where it expresses false, irrational, or incoherent beliefs, the result of a person’s ignorance of a true proposition.

But this kind of intrinsic connection between ignorance and badness does not hold up for cases where a person arrives at a true proposition yet is still ignorant. Consider another insight problem similar in structure to the one about the accountant: A father and son are in a car accident and the father is killed. The son is rushed to the hospital, but the doctor says, “I cannot operate on him—he is my son!” How is this possible? Suppose that when solving the

⁴⁴ See Siegel (2016) for an argument that perception is responsive to reasons. Munton (2019) makes a similar case for perception, and Flores (2021) argues that delusions can be reason-responsive.

insight problem, you arrive at a true proposition that the doctor is the other father. This conclusion illustrates open-mindedness about different existent family structures (at least in the United States). However, for my purposes, failing to consider the possibility that the doctor could be the mother is enough to count as being ignorant. So, simply arriving at a true proposition (or the correct answer to an insight problem) does not eliminate a person's ignorance.

Moreover, it is possible that ignorance can be 'good' in more than an instrumental sense; sometimes ignorance is also rational. Three examples make this point:

- a. you fail to infer that a politician is connected to a child trafficking ring (**Q**) from her choice to eat cheese pizza (**P**)⁴⁵
- b. you fail to infer that a skyscraper is unsafe (**Q**) from the fact that it is swaying in the wind (**P**)
- c. you fail to infer that there will be six more weeks of winter (**Q**) from the fact that Punxsutawney Phil saw his shadow (**P**)

These cases are meant to get at the idea that sometimes failing to infer **Q** from **P** is the epistemically prudent thing to do. In each of these examples there is a consideration (**Q**) that is relevant to some inquiry. So, failing to infer **Q** from the state of affairs (**P**) means that you are ignorant but does not mean that you are irrationally ignorant. This result is precisely what we want for the normative component of attitudinal ignorance — a normative project where ignorance can be epistemically better, as well as worse.⁴⁶

7. Comparisons to other substantive forms

⁴⁵ Thanks to Susanna Siegel for this case.

⁴⁶ For more on the epistemic evaluability of ignorance see Miller Larsen (2023b).

All the accounts of ignorance discussed in this paper argue for the importance of ignorance. But the reasons why we should care about ignorance vary across views. The privative accounts of ignorance care about the value of ignorance given that it is an epistemic deficiency—an absence. Other substantive accounts of ignorance care about the cultural productions of ignorance that promote individual and societal prejudicial structures, as well as manifestations of oppression (epistemic and material). This perspective is often found among philosophers of standpoint epistemology, feminist epistemology, and critical race theory, which are sometimes grouped together as the *epistemologies of ignorance*.

What is the nature of ignorance picked out by the epistemologies of ignorance? The first of these views—the practice view—is concerned with the structural nature of ignorance and provides an externalist picture of ignorance that connects epistemic agents to the societies in which they live.

The Practice View: an epistemic agent is ignorant just in case the agent participates in structural practices of ignorance—willful ignorance, for example, and socially accepted but faulty justificatory practices.⁴⁷

Willful ignorance of structural injustices is one structural practice, according to Linda Alcoff (2007). Another kind of practice involves Mills' white ignorance (discussed above), which is a specific kind of ignorance of racial structural injustices. White ignorance can be seen as a practice that maintain justificatory practices, especially in the United States. And the reason to care about white ignorance is due to the effects of the practices on material goods (e.g., jobs, housing, and other financial opportunities). Recall that for Mills white ignorance is also a disposition to uphold pernicious racial structures that prioritize the histories and plights of white Americans, while erasing the narratives (especially the oppressive ones) of Black

⁴⁷ See Alcoff (2007).

Americans. This disposition, arguably, takes an active role in the mental processes of an agent, as indicated by the active view.

The Active View: an epistemic agent is ignorant just in case the agent actively upholds false outlooks via mechanisms of resistance to belief-forming practices that would increase access to knowledge.⁴⁸

Like the practice view, the active view provides an externalist picture of social structures, for example, in the United States. But, to care about ignorance is to care about how ignorance impacts an individual's mental activists. The problems that Mills and Alcoff raise lie at the interface of individual minds and the social world. Addressing these problems, I argue, requires that we have an equally defensible account of the internalist picture of ignorance alongside the externalist one.

A recent proposal, known as the integrated view, is meant to provide us with an internalist explanation of ignorance as a disposition.

The Integrated View: an epistemic agent is ignorant just in case the agent has an ignorance disposition that manifests in the agent's beliefs— by having a false belief, or lacking belief—, and manifests in the agent's epistemic attitudes (doxastic attitudes, epistemic virtues, and epistemic vices).⁴⁹

The integrated view is closest to the attitude view but faces a fundamental problem with respect to the attitude component. Epistemic attitudes on the integrated view include doxastic attitudes—false belief and suspension of judgment— and epistemic virtues, such as open-mindedness, as well as epistemic vices, such as close-mindedness. Nadia El Kassar, who defends

⁴⁸ See Mills (2007) and Medina (2013) (2016).

⁴⁹ See El Kassar (2018).

the integrated view, provides a comparative example to illustrate the role of the attitude component as part of the nature of ignorance. Two people—Hannah and Kate—are both deeply ignorant about the fact that cruise ships have high emissions of carbon and sulfur dioxide that are bad for the environment. Both women have booked an upcoming trip on a cruise ship and share their plans with a third friend, Monika. Monika informs them of the harm of cruise ships to the environment based on a documentary she recently watched. Hannah stops Monika before she can even provide the negative information, while Kate is interested and asks Monika for more details. El Kassir's conclusion is that Hannah is close-minded, whereas Kate is open-minded.

Close-mindedness and open-mindedness, however, are not by themselves uniquely indicative of an attitudinal component of ignorance. Such epistemic virtues and vices might be *causes* of ignorance but are not themselves *constitutive* of ignorance. If this is right, the integrated view faces a fundamental problem. El Kassir argues that ignorance is constituted by these virtues, but her argument misses a crucial distinction. While it is true that ignorance can be *constituted by* epistemic virtues, such as open-mindedness, it is something else to say that ignorance is *constitutive of* such epistemic virtues. Since the aim of El Kassir's paper, as well as my own, is to discuss the nature of ignorance, a view that incorporates virtues and vices, which are variable features of a person's doxastic life would be committed to a kind of ignorance that changes given the dispositions of the particular person. So, Hannah's ignorance would be the kind constitutive of an epistemic vice and Kate's ignorance would be the kind constitutive of an epistemic virtue, leading to two drastically distinct kinds of ignorance, rather than one unifying and explanatorily robust account. The attitude view also will not be perfectly unifying across all descriptions of ignorance, but the characteristic properties of the attitude will be consistent across subjects so that when an epistemic agent is ignorant, virtue and vice will be separate assessments that bear on, for example, a person's character.

8. Character

—No virtue is constituted by, or based upon, ignorance.⁵⁰

The above quote from Julia Driver summarizes the consequence of an Aristotelian view of virtue. Such virtue ethics is the preferred approach for discussion of the quality of will someone exhibits when they are ignorant, leaving no room for manifestations of ignorance that are virtuous as part of being praiseworthy or rational. One reason to prefer the Aristotelian framework is due to our discussion of people being ignorant *simpliciter*. The preference is most apt when the kind of ignorance at hand is less about factual content and more about a mental blind spot to target information that we think the person should be aware of. To describe such people, we sometimes say that they are just ‘ignorant’ at the core. Or, in African American culture, the person is said to be ‘ignant’—articulating a difference between ignorance of everyday factual knowledge and deeply ingrained (perhaps immutable) ignorance.

The distinction between being ignorant *simpliciter* and being mistakenly ignorant on occasion can be cashed out as a difference of ‘character,’ where character ignorance is the disposition to be factually ignorant. The attitudinal account may seem like a natural ally of character ignorance. The virtue-theoretic framework describes ignorance as a failure to manifest intellectual flourishing (such as arriving at true belief), as well as a failure to eliminate vices (such as dogmatism or close-mindedness). Aristotelean virtue theory also lends well to an account where considering relevant possibilities in inquiry is an epistemic good.

The downside of cashing ignorance in virtue-theoretic terms is the inevitable entanglement between the intellectual and moral components. This entanglement is true of the Aristotelian picture, as well as contemporary views about intellectual culpability.⁵¹ Character ignorance can

⁵⁰ See Driver (1989).

⁵¹ See, for example, Medina (2013) for the latter.

nicely capture the moral failings of ignorance, as well as capture the reasons why we tend to ‘write someone off’ when they are ignorant *simpliciter*.

While it is tempting to incorporate virtue-theoretic terms into the attitude view, such temptations should be resisted. For one, it leads to a view that lacks normative depth. If someone is ignorant *simpliciter* then we might think that they are also factually ignorant (i.e., there is some fact of the matter about which they lack knowledge/true belief/awareness). Yet when someone is factually ignorant, it is not the case that they manifest character ignorance. In addition, the normative significance of their ignorance is precisely what should be up for analysis. On my view, the assessment of ignorance as ‘bad’ is second to the analysis of the nature of ignorance.

An additional reason to avoid the virtue-theoretic approach for the ethics of ignorance comes from the nature of our core cases. It is too easy to describe the subjects as ignorant as the result of negative character flaws, such as being a *sexist*, *bigot*, or *racist*. This is not to say that there is never a need for moral condemnation for people with sexist, bigoted, or racist attitudes. My point is that the epistemic assessment of such attitudes should be disentangled from the moral aspects, thus allowing us to note how both Finn and Beverly perform poorly epistemically, while (perhaps) only Finn performs poorly morally.⁵²

9. Conclusion

Ignorance is epistemically significant. The purpose of this paper has been to argue for a, *heretofore*, overlooked type of ignorance that is an attitude. The attitudinal account introduces an important distinction between ignorance as a mere *absence* and ignorance as an *attitude* capable of doing being epistemically efficacious. The upshot of the attitude view is that can

⁵² Moreover, the disentanglement is possible with a proposed tension between the epistemic and moral (contra Gendler 2011). For more on the difference between the epistemic and the moral with respect to responsibility for ignorance see Miller Larsen *ms*.

address the systemic nature of ignorance as a mental habit across various content that may be amoral and asocial. It can also account for the full spectrum normative ignorance.

The proposal sets out to establish an ethics of ignorance that makes rational evaluation plausible. I have argued that ignorance is epistemically significant because it is more than a mere absence of a cognitive epistemic state. From here, we can epistemically evaluate ignorance since it is not only *criticizeable*, but also does *epistemic work*.

REFERENCES

- Alcoff, L. (2007). Epistemologies of Ignorance: Three Types. In *Race and Epistemologies of Ignorance*, Sullivan, S. & Tuana, N. (eds.) New York: SUNY Press
- Armstrong, D. (1968). *A Materialist Theory of the Mind*. New York: Routledge & Kegan Paul.
- Baker, L. R. (1995). *Explaining Attitudes*. Cambridge: Cambridge University Press.
- Banaji, M. R. & Greenwood, A. G. (1994). Implicit Stereotyping and Prejudice. *The Psychology of Prejudice: The Ontario Symposium* 7: 55-76.
- Bar-Hillel, M. Noah, T., & Frederick, S. (2018). Learning Psychology from Riddles: The Case of Stumpers. *Judgement and Decision Making* 13 (1): 112-122.
- Bendana, J. & Mandelbaum, E. (2021). The Fragmentation of Belief. In C. Borgoni, D. Kindermann & A. Onofri (eds.), *The Fragmented Mind*. Oxford: Oxford University Press, pp. 78-107.
- Block, N. (1978). Troubles with Functionalism. *Midwest Studies in the Philosophy of Science*, 9:261-325.
- Bear, A. Bensinger, S, Jara-Ettinger, J. Knobe, J., & Cushman, F. (2020). What Comes to Mind? *Cognition* 194: 104057.
- Driver, J. (1989). The Virtues of Ignorance. *The Journal of Philosophy*. 86 (7): 373-384.
- Dutant, J. (forthcoming). *The New Evil Demon Problem*. Oxford University Press.
- Eagan, A. (2007). "Seeing and Believing: Perception, Belief Formation, and the Divided Mind. *Philosophical Studies* 140 (1): 47-63.
- El Kassar, N. (2018). What Ignorance Really Is. Examining the Foundations of Epistemology Ignorance. *Social Epistemology* 32 (5): 300-310.
- Fine, K. (2018). Ignorance of Ignorance. *Synthese* 195 (9): 4031-4045.
- Flores, C. (2021). Delusional Evidence-Responsiveness. *Synthese* 199 (3-4): 6299-6330.
- Fodor, J. (1968). *Psychological Explanation*. New York: Random House

- Friedman, J. (forthcoming). The Aim of Inquiry? *Philosophy and Phenomenological Research*.
- Gendler, T. (2011). On The Epistemic Costs of Implicit Bias. *Philosophical Studies* 156 (1): 33-63.
- Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, A. & Olsson, E. J. (2009). Reliabilism and the Value of Knowledge. In *Epistemic Value*.
- Guerrero, A. A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136(1): 59-97.
- Haack, S. (2001). "The Ethics of Belief" Reconsidered. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Steup, M. (ed.) Oxford University Press, pp. 21-33.
- Hieronymy, P. (2011). Believing at Will. In D. Hunter (ed.), *Belief and Agency*. Calgary: University of Calgary Press. Pp. 149-187.
- Le Morvan, P. (2011). On Ignorance: A Reply to Peels. *Philosophy* 39 (2): 335-344.
- Le Morvan, P. & Peels, R. (2016). The Nature of Ignorance: Two Views. In R. Peels & M. Blaauw (eds.), *The Epistemic Dimensions of Ignorance*. Cambridge: Cambridge University Press, pp. 12-32.
- Lewis, D. (1972). Psychophysical and Theoretical Identifications. *Australasian Journal of Philosophy*, 50: 249-258.
- Lewis, D. (1980). Mad Pain and Martian Pain. In N. Block (ed.), *Readings in the Philosophy of Psychology* (Volume 1), Cambridge: Harvard University Press
- Littlejohn, C. (2017). How And Why Knowledge Is First. In J.A. Carter, E.C. Gordon, & B. Jarvis (eds.), *Knowledge First*. Oxford: Oxford University Press, pp. 191-45.
- Loar, B. (1981). *Mind and Meaning*. Cambridge: Cambridge University Press.
- Johnson, G. M. (2020). The Structure of Bias. *Mind* 129 (516): 1193-1236.
- Johnson King, Z. (2022). Radical Internalism. *Philosophical Issues* 32 (1): 46-64.
- Miller Larsen, E. (2023b). Responsibility for Ignorance.
- Mandelbaum, E. (2015). Attitude, Inference, Association: On the Propositional Structure of Implicit Bias. *Noûs* 50 (3): 629-658.
- Medina, J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and Resistant Imaginations*. Oxford: Oxford University Press.
- Medina, J. (2016). Ignorance and Racial Insensitivity. In *The Epistemic Dimensions of Ignorance*. Peels, R. & Blaauw, M. (Eds.). Cambridge: Cambridge University Press.

- Meylan, A. (2020). Ignorance and Its Disvalue. *Grazer Philosophische Studien* 97 (3): 433-447.
- Mills, C. (2007). White Ignorance. In *Race and Epistemologies of Ignorance*, Sullivan, S. & Tuana, N. (eds.) New York: SUNY Press.
- Munton, J. (2019). Perceptual Skill and Social Structure. *Philosophy and Phenomenological Research* 99(1): 131-161.
- Nozick, R. (1981). *Philosophical Explanations*. Harvard University Press.
- Peels, R. (2010). What is Ignorance? *Philosophia* 38(1): 57-67.
- Phillips, J. & Cushman, F. (2017). Morality Constrains the Default Representation of What is Possible. *Proc Natl Acad Sci U S A* 114 (18): 4649-4654.
- Price, H. H. (1969). *Belief*. London: Allen & Unwin.
- Pritchard, D. (2021). Ignorance and Inquiry. *American Philosophical Quarterly* 58 (2): 111-124.
- Pritchard, D. (2022). Ignorance and Normativity. *Philosophical Topics* 49 (2): 225-243.
- Putnam, H. (1975). *Mind, Language, and Reality*. London: Cambridge University Press.
- Quilty-Dunn, J. (2020). Perceptual Pluralism. *Noûs* 54 (4): 807-838.
- Quilty-Dunn, J. & Mandelbaum, E. (2018). Against Dispositionalism: Belief in Cognitive Science. *Philosophical Studies* 175: 2353-2372.
- Setiya, K. (2013). Epistemic Agency: Some Doubts. *Philosophical Issues* 23 (1) 179-198.
- Siegel, S. (2016). *The Rationality of Perception*. Oxford: Oxford University Press.
- Silva, P. & Siscoe, R. W. (2023). Ignorance and Awareness. *Forthcoming in Noûs*.
- Schwitzgebel, E. (2002). A Phenomenal, Dispositional Account of Belief. *Noûs* 36: 249-275.
- Schwitzgebel, E. (2013). A Dispositional Approach to Attitudes: Thinking Outside the Belief Box. In N. Nottelmann (ed.), *New Essays on Belief*. New York: Pgrave Macmillan, 75-99.
- Srinivasan, A. (2020). Radical Externalism. *Philosophical Review* 129(3): 395- 431.
- Triesch, J., Ballard, D. H., Hayhoe, M. M., Sullivan, B. T. (2003) What You See is What You Need. *Journal of Vision*. 3 (1): 86-94.
- Van Woudenberg (2009). Ignorance and Force: Two Excusing Conditions for False Beliefs. *American Philosophical Quarterly* 35 (138): 1-31.
- Wang, C. & Wang, J. (2022). Pritchard on Ignorance and Normativity. *Asian Journal of Philosophy* 2(1):1-18.
- Zimmerman, M. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press

Part II

Responsibility for Ignorance

1. Introduction

This paper is not about moral responsibility for ignorance. Instead, it concerns two assumptions for literature on the ethics of attitudes, particularly with respect to their *rational* status. The first is the standard view of ignorance.

The Old Guard: Ignorance is the lack or absence of knowledge.⁵³

The second is a norm for responsibility.

Responsibility Norm: We are epistemically responsible for a state just in case it is the 'kind of attitude' that is normatively evaluable.⁵⁴

The combination of the old guard and the responsibility norm results in a picture of ignorance where we are never responsible for our ignorance. Since the old guard assumes that ignorance is the lack of knowledge, ignorance can be no more epistemically evaluable than absences. And since the responsibility norm applies only to attitudes—which ignorance is not grouped among—it cannot extend to ignorance.

I argue that this picture fails to accommodate our commonsense understanding of rational responsibility for ignorance. When we are ignorant of a friend's birthday or are ignorant of the time of a staff meeting, or are ignorant of glaring social injustices, we are rationally responsible for such ignorance, in the sense that there are epistemic reasons for why we *ought not* be ignorant.

⁵³ Within the literature on ignorance, THE OLD GUARD view is more commonly referred to as the 'Standard View'. Recently, authors have rejected the Standard View in favor of the 'New View', which is the idea that ignorance is the lack of true belief. The difference between the view that ignorance is the lack of knowledge (Fine 2018; Haack 2001; Le Morvan 2011; Zimmerman 2008) and the view that ignorance is the lack of true belief (Goldman 1986; Goldman & Olsson 2009; Guerrero 2007; Peels 2010; van Woudenberg 2009) will not matter for my purposes.

⁵⁴ I purposefully leave the 'kind of state' vague so that one can fill in the features that they think are necessary for an epistemically evaluable state. For instance, some conditions for responsibility for attitudes include reason-responsiveness (McCormick 2011), answerability (Hieronymi 2008), or indirect control (Peels 2017).

To make my case, I will defend the thesis that—contra the old guard— ignorance is an attitude, and epistemic agents bear responsibility for that attitude. The pathway to my conclusion requires a three-part argument. First, we need an argument for the kind of responsibility we have not to be ignorant (§2). Second, we need an argument for the kind of ignorance to which epistemic agents are responsible (§3). And finally, we need an argument for the epistemic evaluability of ignorance (§4).

2. What is rational responsibility for ignorance?

To begin I must address a preemptive objection to the idea that we can be responsible for ignorance. One might argue that responsibility for ignorance (or other attitudes for that matter) necessarily entails blameworthiness. In what follows, I make the case that this objection comes from a serious misstep in normative categorization. For one, though it is reasonable to think so given some of the discussions on responsibility for belief⁵⁵, epistemic responsibility should not be thought of as not simply moral responsibility for epistemic entities. There are, however, several philosophers who are picking out something distinct from moral responsibility in their appraisal of the responsibility of beliefs.⁵⁶ Following their lead, I aim to make progress on ‘epistemic responsibility’ as a concept for normative evaluation distinct from its moral counterpart.

One problem is that the normative categories for epistemic appraisal are a sock drawer. Responsibility, justification, and blame have been theoretically mixed and matched so that it is impossible to determine which category entails the other. Here I try to give a precise sketch of epistemic responsibility, pulling apart these categories from their theoretical entanglements.

⁵⁵ Hieronymi (2008), for example, argues that viewing belief as something that a person (intentionally) settles on allows us to transfer all of the assessments from moral responsibility, including *blame* and *praise*, to epistemic responsibility.

⁵⁶ Goldberg (2017) and Kauppinen (2018) seem to be in this camp.

First, ignorance (and other attitudes) is subject to assessments of epistemic responsibility as part of the rational standards that we ‘owe to ourselves.’ An epistemic agent is epistemically responsible just in case these rational standards have not been met. Such standards measure when thinking has gone well, as well as when thinking has gone poorly. Measures of evaluation count for little without the agents who work towards them. In this sense, responsibility is an activity an agent engages in, and responsibility is activated when a state is attributable to that agent. We can ask, *who is responsible for dinner?* suggesting that someone should have been responsible but failed to do the cooking. Or we can ask a similar question in a positive tone, *who is responsible for dinner?* indicating an epicurean delight.

For epistemic states, philosophers have primarily been concerned with thinking ‘gone wrong,’ where our assessments seek to understand why someone failed to comply with rational standards. For example, someone who believes that there will be six more weeks of winter because Punxsutawney Phil saw his shadow is responsible for their belief because they have not met our rational standards to have credible evidence that counts in favor of the belief.⁵⁷

When philosophers discuss the ethics of belief, they have more than just a failure of rational standards in mind. For some thinkers, the demand of epistemic responsibility does not only indicate a failure to be rational, but it also expresses a way that we have wronged others. Take Sanford Goldberg’s argument for should-have-known cases.⁵⁸ Goldberg discusses a father who places his daughter in a dangerous situation while watching a football game by leaving a pocketknife within her reach. The argument is that the father *should have known* that the pocketknife would present a danger for his daughter. The reasons he ought not be ignorant are social/occupational reasons. According to Goldberg, the father’s rational responsibility is not only a matter of the justification of his ignorance but is also grounded in his social role as the

⁵⁷ Punxsutawney Phil is the groundhog that is used in Punxsutawney, Pennsylvania, to predict the weather in Spring. If the groundhog does not see his shadow, then Spring is predicted to ‘come early’. But if the groundhog *does* see his shadow, then supposedly there will be six more weeks of winter.

⁵⁸ See Goldberg (2017).

child's father. Antti Kauppinen (2018) also has a view that grounds responsibility in social norms and expectations. For Kauppinen, we are responsible for beliefs because we need to be reliable sources of knowledge within our epistemic communities.⁵⁹ Kauppinen's paradigmatic case is about an amateur political pundit who has no special knowledge about the political climate he in which he is opining but makes claims as if he does. Such amateur pundits, Kauppinen thinks, should be held responsible for their assertions because they make themselves unreliable within the epistemic community. Goldberg and Kauppinen argue that our epistemic norms of responsibility are grounded in social practices and what we epistemically owe to each other.

Views that tie responsibility to what we owe other people, naturally build in the idea that we are blameworthy when we are epistemically irresponsible. The assumed link between blame and responsibility for epistemic assessment can be summarized as the following idea. If an epistemic agent is unjustified in believing that p , then she is also blameworthy for believing that p . So, if you believe that there will be six more weeks of winter because Punxsutawney Phil sees his shadow (an unjustified belief), then you are also blameworthy for this belief. But are you blameworthy in the same way that the father in Goldberg's should have known case is blameworthy if he has the (unjustified) belief that his daughter is safely playing out of harm's way? There seems to be another notion of blameworthy belief needed to differentiate between the two examples. For these reasons Brown (2020) argues that justification and blameworthiness should be pulled apart. Justification, she thinks, is about 'improving norm conformity,' and epistemic blame, whatever it is, should be more than this. With the distinction made, Brown argues we can better accommodate blameworthy belief by extending our notion of blame.

⁵⁹ The idea of an epistemic community can be thought of as any grouping of people that share information. Communities can include, friend groups, universities, clubs, families and so on.

Brown is correct in her challenge but mistaken in her conclusion. Intuitively, Brown argues that beliefs can be justified and blameworthy, or unjustified and blameless. The belief that any particular person in a high crime neighborhood may be justified given regional statistics, but it may, nonetheless, be blameworthy since it stereotypes a group of people. And basing one's weather predictions on the habits of groundhogs is unjustified, as I have argued, but may, nonetheless, be blameless. The upshot of Brown's view is that it puts pressure on the assumptive justification-blame relation. Contra Brown, I do not think the answer is to extend the notion of epistemic blame to accommodate blameworthy belief. We can say that a belief is morally blameworthy just in case it fails pre-established moral standards. If epistemic blame is just moral blame for beliefs, then there is no worry for the account. But since Brown and others insist on a unique species of epistemic blame, more is needed to argue that it directly correlates to the justificatory status of a belief.

2.3 Is epistemic responsibility possible without epistemic blame?

Can we account for the idea that epistemic agents are responsible for their ignorance without also assigning blame? There is a compelling case to be made that epistemic agent can be responsible yet blameless. Take the case of Radha, for example, who lives in a community steeped in toxic patriarchy. Radha believes that she deserves the brutal beatings she receives from her husband.⁶⁰ Strong sympathies for Radha may lead us to think that she is not responsible for her belief even though she is, in a sense, ignorant. But these sympathies, I think, are due primarily to our hesitancy to blame her.

Part of the problem is that we require a distinction between grounding claims and motivational claims for responsibility. Feelings of guilt or shame can be motivational tools for ignorance revision, but they are not, in themselves, the grounding features of responsibility. If guilt were

⁶⁰ This example is from Srinivasan (2020).

the grounds for responsibility, then someone would only be responsible for their ignorance if they felt guilty about it. Instead, I think responsibility should be grounded in our rational standards and practices of ‘improving norm conformity.’ Instead of depending on the blaming practices of our communities, we owe it to ourselves to be as rational as possible.

In Radha’s case she can be responsible yet blameless, owning up to her false belief without also enduring shame from her community. The way I have categorized responsibility has two distinct types of assessment. There is a ‘backward-looking’ assessment where people are responsible for their attitudes because it was a ‘bad’ thing to believe; and there is also a ‘forward-looking’ assessment of responsibility where people are responsible for improving their attitudes. The resulting framework of ‘forward-looking’ responsibility extends the usefulness of the normative category. I will show how the extension works in §4.2.

Lastly, it is worth noting that epistemic responsibility may have a connection to praise and praiseworthiness—the often-forgotten category at the bottom of the normative sock drawer. Praise is commonly viewed as the positive counterpart of blame, which suggests the following idea. If an epistemic agent is justified in believing that p , then she is praiseworthy for believing that p . So, there is a justification-praise relation that should account for praiseworthy beliefs.

In the moral domain, it might be that we are praiseworthy for what we do, or that we are praiseworthy for what we try to do (e.g., our intentions, motivations, desires, or will), or we might be praiseworthy for both (Johnson King 2021). The analogue for belief then would be that we are praiseworthy for what we believe, what we try to believe (e.g., inferences, heuristics, and other epistemic transitions), or both. Yet when we get into specific cases, it is more difficult to outline the parallels. When a Good Samaritan saves a puppy from a burning building, we say

that they are praiseworthy for doing so.⁶¹ But when a professor believes his female students are excellent logicians it is less clear which feature of his belief is praiseworthy.

Concerns for the justification-praise relation arise in a similar vein to those for the justification-blame relation. One reason to doubt the distinctness of epistemic praise follows from the nature of spectrums in general. If there are praiseworthy and blameworthy beliefs, then the spectrum must also include ‘neutral’ beliefs. Not every action is praiseworthy or blameworthy (e.g., a sneeze, lifting your arm, opening the mail), so not every belief, or ignorant state for that matter, is going to be praiseworthy or blameworthy.⁶² But we can easily imagine that justified beliefs will be neutral (believing that it 8:54pm), so justification on its own will not indicate praise. These remarks do not mean that we are not responsible for our attitudes in a robust way. The justification analysis may be inapt in articulating a distinctive form of epistemic praise or blame, but it can still, nonetheless, articulate when an agent’s attitude is rationally sound or rationally problematic.

3. Ignorance: a new account

What kind of ignorance demands rationality? On the OG sense of ignorance where it is an epistemic absence, responsibility claims are paradoxical, or supererogatory at best. A new account of ignorance is required to address the rational standards that attach to ignorance. The following case, based on real-world events, provides a counterexample to the OG view.⁶³

⁶¹ The moral theorist who doesn’t want to reduce a view of praiseworthiness to ‘good’ and ‘bad’ actions would say that the Good Samaritan is praiseworthy for cultivating something like ‘good will’ (Johnson King 2021).

⁶² Also, not every morally good action is going to be morally praiseworthy. For example, someone isn’t morally praiseworthy for having failed to beat their children or having failed to sexually harass someone on a given day. See Johnson King (*ms*) on the ‘Slow Clap Phenomenon’ for more on this issue.

⁶³ For a detailed argument see Miller Larsen (2023a).

Who's the professor? On the first day of classes, Finn—a usually excellent student—is late to his intro to philosophy course. Thinking he has already missed valuable information; Finn frantically searches the room for the course professor. He considers several people in the room, including a male student in the back wearing Oxbridge-like attire. However, Finn fails to consider the woman at the board, despite her performing duties that would normally indicate she is the professor.

Let us suppose for now that the OG view is all that is required to explain the example, so that Finn is ignorant because he ‘doesn’t know’ that *the woman is the professor*. Finn can be said to also have a false belief that the other student wearing ‘Oxbridge-like attire’ is the professor. False beliefs count as ignorance, particularly when they operate as distributions of error or sources of misinformation. For example, Charles Mills’ view of white ignorance is based on the idea that ignorance is false belief.⁶⁴ Mills draws the distinction between ‘passive’ ignorance on the one hand and ignorance that ‘fights back’ on the other. White ignorance is the ‘fighting’ kind—it is resistant to belief revision and obstructs knowledge. Importantly, it is the product of social practices and systemic distributions of error that infiltrate individual views on race and racism (primarily in the United States). José Medina has also written about this kind of ignorance where ignorance involves “a whole battery of mechanisms of avoidance and resistance to know and to learn” (Medina 2016, 182). The arguments from Mills and Medina are part of a group of views, known as the *epistemologies of ignorance*, which outline the externalist nature of ignorance as cultural-psychological tools within social systems that afford individual

⁶⁴ See Mills (2007).

ignorance of racial injustice.⁶⁵ From these externalist accounts it is natural to wonder what kinds of mechanisms undergird ignorance of evidence uptake.⁶⁶

Of course, there are many times when things do not occur to us. What it means to be human just is to be the kind of creatures that have cognitive limitations that filter the amount of absorbable information. I am not proposing a normative project where agents must be omniscient. In fact, the problem of ignorance is not (e.g.) about Finn's fallibility. Rather, the problem of ignorance is that Finn omits obvious and relevant possibilities as evidence that *should be* legible to him. The OG view cannot address the problem, as I have described it, since it has no way of distinguishing between information we merely lack and information we fail to register.

The case of Who's the Professor? yields a new proposal for ignorance. The view that ignorance is an attitude which a characteristic profile of omitting. An epistemic agent is ignorant of the fact that *p* at *t* just in case the agent ignores that *p* at *t* by failing to consider obvious and relevant possibilities in inquiry. In the ordinary sense of the word 'ignoring' is the deliberate disregard of information. When we ignore an email or ignore a phone call, we manifest the active sense of ignoring. But the active sense of ignoring is not what I have in mind for the sense in which 'ignoring' is an attitude. When we ignore, in the epistemic sense, we unwittingly (non-consciously) omit obvious, and relevant, possibilities that *p* could be the case. Contrast such ignoring with examples of willful ignorance where someone suspects that *p* but *deliberately* disregard this possibility (Lynch, 2016). Instead, epistemic agents, in my cases, are

⁶⁵ For Mills (2007), Medina (2016), and Alcoff (2007) ignorance is sometimes, but not always, sub-intentional. They also discuss cases where people willfully perpetuate social structures that maintain and promote racial inequities.

⁶⁶ Miller Larsen (2023a) argues that ignorance is not mere lack of knowledge, nor is it simply the lack of knowledge plus other familiar psychological pathologies (i.e., dissonance, confirmation bias, or belief preservation).

sub-intentionally ignorant—they unwittingly omit the possibility that p precisely because they do not suspect that p .⁶⁷

The attitude view is an agent-centered account of ignorance. We are responsible for ignorance because ignorance is just as much of a ‘commitment’ as belief, where an important feature of this commitment consists in settling questions. The theory that I advance is underwritten by a re-orientation of the nature of ignorance. The approach finds company with Jane Friedman’s work on suspension of judgment where suspension is held to be a *sui generis* epistemic attitude (Friedman, 2013). Friedman thinks that suspension is not mere absence of belief, it is its own interrogative attitude. Similarly, ignorance is not mere absence of knowledge, it is also its own attitude. And, just as we are responsible for our beliefs, so too are we responsible for ignoring. Ignoring is ‘belief-like’⁶⁸ in the relevant sense because it inferentially interacts with other mental states in similar ways to belief by impacting higher order zetetic⁶⁹ propositions. For Student, being ignorant that the woman is the professor inferentially interacts with the question *Who is the professor?* in the sense that his ignorance limits the inferential connections that are drawn. Ignoring is more impactful than an absence because it bears on the nexus of one’s attitudes—influencing how we view and understand the world.⁷⁰

A final and crucial difference exists between the OG view and ignoring, which is a causal difference. When you unknowingly put arsenic in your friend’s cup of tea you do so because you are ignorant that the sugar contains arsenic—your ignorance is caused by your lack of

⁶⁷ Crucially, on my view ignoring is an attitude, albeit an implicit one, that epistemic agents take towards some content. The account I give follows the assumption that, as an attitude, ignorance can be modeled as a propositional attitude, which consists in a binary relation between an attitude A and a proposition p . The reason someone is ignorant, for example, that the president is lying (as some people were surrounding the January 6, 2021 riots on Capitol Hill) consists in their having an ignorance attitude A that is related to the proposition *the president is lying*. On my view, people are ignorant just in case they have ignored the possibility that the president *could have been* lying.

⁶⁸ The comparison can be understood as follows—Belief: Believing:: Ignorance: Ignoring. Just like belief has two modes, which include the state of having a belief and the mental activity of believing, so too does ignorance have two modes, where ‘ignoring’ is the mental activity of the mental state of ignorance (at least in those case where someone unwittingly omits a possibility).

⁶⁹ I use ‘zetetic’ to refer broadly to any state, attitude, or proposition related to inquiry.

⁷⁰ For an in-depth argument of the attitudinal view of ignorance see Miller Larsen (2023a).

knowledge. Note that causal ignorance is often one-off, so that we are not systematically ignorant of arsenic looming in tea.⁷¹ In fact, if someone were to be ignorant of such a fact, it seems more like a case of bad epistemic luck than a matter of epistemic failing on their part. In contrast, ignoring is constitutive of one's ignorance, not simply the result of it. When Finn is ignorant, his attitude is not *caused by* what he does not know. Rather, his ignorance is *constitutive* of belief forming processes that bear on the success of his inquiry.

4. Dissolving the paradox

I started with the paradoxical claim that we are never responsible for ignorance. The claim is based on two assumptions from literature on the ethics of belief and the epistemic dimension of the nature of ignorance.

The Old Guard: Ignorance is the lack or absence of knowledge.

Responsibility Norm: We are epistemically responsible for a state just in case it is the kind of attitude that is normatively evaluable.

Contra the OG view I argued that ignorance can be more than a mere absence—it is an attitude. I also addressed a preemptive objection that being responsible for ignorance necessarily entails that an agent is also blameworthy for that ignorance. Such a conclusion would be worrying given the nature of our cases as sub-intentional and unwitting manifestations of ignorance. *But how is an epistemic agent responsible for their ignorance?*

Consider one more case of ignoring.

⁷¹ You might be wondering about the epistemic status of someone who is always thinking about doomsday possibilities. I consider such example in §5.

Whodunit 1: In a stately mansion you find that someone has been murdered! The following people are at the residence at the time of the crime: Yvette the maid, Miss Scarlet, Colonel Mustard, Professor Plum, and Wadsworth the Butler. You have evidence that connects each person to the murder; you know that at least one person committed it; and you know that one of the suspects has an alibi. As the detective on the case, your task is to determine who did it. You consider some suspects yet unwittingly ignore the possibility that Wadsworth is one of them.

$\left\{ \begin{array}{l} A_1 \\ A_2 \\ A_3 \\ A_4 \end{array} \right.$	YVETTE THE MAID	✓	<table><tr><td>✓</td><td>CONSIDERED</td></tr><tr><td>✗</td><td>ELIMINATED</td></tr></table>	✓	CONSIDERED	✗	ELIMINATED
	✓	CONSIDERED					
	✗	ELIMINATED					
	MISS SCARLET	✓					
COLONEL MUSTARD	✗						
PROFESSOR PLUM	✓						

Like Who's the professor? this case features a type of ignorance where an epistemic agent omits obvious and relevant possibilities. If we were to model this case following the OG framework, we get the following set as an explanation for the suspects you consider and eliminate when deliberating over the question Q^{72} , "Who did it?"

What makes an attitude evaluable? Pamela Hieronymi (2008) has a view where our evaluations of attitudes consist in asking the reasons for which someone believes that p . Why-questions, on her view, are questions of responsibility, and beliefs are the 'kind of attitudes' that are subject to responsibility because they stand in the right relation to 'settle' why-questions. For Hieronymi, a belief settles the question whether to believe that p as part of the justification for why you believe that p . When I settle whether to believe that it will be six more weeks of winter, and I arrive at the belief that it will be six more weeks of winter because Punxsutawney Phil

⁷² I use an italicized 'Q' to denote a question in inquiry and a bold 'Q' to denote the conclusion of an inference.

sees his shadow, my belief is evaluable by evaluating the answer to why I settled the belief that way.

This approach can account for the evaluability of attitudes by providing an explanation of the relation between attitudes and inquiry. In whodunnit, each answer $\{A_1 \text{ to } A_4\}$ potentially settles Q , so that arriving at one of these answers is a way of fully settling Q . As it stands, Hieronymi's view does not count ignorance as an evaluable attitude. This is not an oversight of her view. The OG view simply cannot fit into this framework. To accommodate cases such as whodunit and who's the professor? I suggest that ignoring that p is part of our deliberative processes in the sense that it is part of how we settle questions.

With ignorance as an attitude, we can model your deliberative process in whodunit as follows:

A_1	YVETTE THE MAID	✓
A_2	MISS SCARLET	✓
A_3	COLONEL MUSTARD	✗
A_4	PROFESSOR PLUM	✓
A_5	WADSWORTH THE BUTLER	¬

✓	CONSIDERED
✗	ELIMINATED
¬	IGNORED

When you ignore that Wadsworth could have done it you have a fifth answer in the set that can potentially settle Q . Each of the answers partially settles Q , so that ignoring any one of them constrains the success of settling.⁷³

⁷³ Any particular stance is only a partial answer to Q if it turns out that there are multiple murders in the case. For instance, settling on A_1 may be true, but you would not want to do this and ignore that A_2 and A_4 could also be answers to Q . At the end of *Murder on the Orient Express*, there is a line in the movie where the detective says, "no one person committed the murder, yet someone committed the murder". The detective in the case strategically gives this line before he reveals that everyone on the train is the murderer.

Suppose you are asked “Why did you ignore that Wadsworth could have done it,” where we wanted you to justify your ignorance not just explain it. The reason we want to evaluate ignorance is because there is something epistemically wrong with it. As the detective, ignoring Wadsworth is irrational when there is evidence—right in front of your face—that connects him to the crime. If you have organized your mind such that obvious and relevant possibilities do not arise, then ignorance is not something that just happens to you, it is rationally problematic. The problem of ignorance lies in which answers are not considered, and how they impact those that are. And when certain obvious and relevant answers are omitted, the incomplete set of answers result in false conclusions that bear on *Q*. As a result, the attitude of ignoring structures the space of inquiry as an answer that is not considered yet is rationally significant.

How should we evaluate ignorance? I have also suggested that ignorance inferentially bears on *Q* in the sense that it limits the inferential connections that can be drawn. Through inference the epistemic ‘badness’ of ignorance transmits to the deliberative process when you have an open inquiry, so that the omission is an epistemic error.

Error of Omission: Even though you have an experience (**P**), you fail to infer **Q** from **P** so that you omit the conclusion (**Q**).

In whodunit you have an experience that Wadsworth is in the mansion at the time of the murder (**P**), yet you fail to infer that he is a suspect (**Q**) and omit the possibility altogether (as in *A*₅). As a result, ignoring as a ‘bad’ epistemic state negatively impacts inquiry given the failed inference (**P**—>**Q**).

In addition to your initial inquiry, ignoring Wadsworth influences questions downstream from *Who did it?* Your initial ignorance may cause you to fail to consider sub-questions such as, *where was Wadsworth at the time of the crime? Does Wadsworth have a motive? Is Wadsworth the primary killer or an accomplice?* Ignoring in one case ramifies to many cases of ignorance.

I have now made the case that ignorance as an attitude cannot be epistemically ignored. I have argued that the reason you are held responsible—on my modified Hieronymi view—is because you have settled whether to ignore that *p*. Such settling inferentially bears on *Q* so that it transmits the epistemic error of your ignorance to how you structure your inquiry.

How does the impact on the inferential structures of reasoning connect to a person's responsibility? Let us return to Who's the Professor? When Finn ignores that woman is the professor, we can ask for the reasons why he is ignorant when he takes an ignoring attitude towards that proposition. Although the attitude is non-conscious, sub-intentional, and unwitting, Student is responsible because he has failed to comply with a rational standard that requires epistemic agents to consider obvious and relevant possibilities. The proposition that the woman is the professor is 'relevant' to Finn because it is one (partial) answer to the question, *who is the professor?* The omitted answer is also 'obvious' because evidence that supports the answer is readily available to him.

Since ignoring the woman is a partial answer that Student takes to *Q*, it is an attitude for which he is responsible. The responsibility consists in a 'forward-looking' requirement to improve his belief states so that it is epistemically better for Student to consider obvious and relevant possibilities, the reason being that these possibilities directly bear on higher order zetetic propositions. So, not only are there practical reasons to know who the professor is when one is a student, there are also epistemic reasons for reviewing one's ignorance.

On my view, Student is responsible for his ignorance yet blameless. His epistemic blamelessness does not mean that there is no sense of blame to which he is subject. He may be morally blameworthy for ignoring the woman in this case and is especially morally blameworthy if he tends to ignore female authority figures. My argument has been that there is not a unique category of epistemic blame that is fitting for cases of ignoring.

Considering these arguments, I think ignorance is epistemically evaluable because it is operative in structuring our inquiry spaces. This means that it influences the kinds of sub-inquiries we undertake, which bear on the questions that occupy us.

5. Is there such a thing as rationally 'responsible' ignorance?

Throughout this paper I have been concerned with ignorance for which we are responsible. For many of us, the worry of responsibility arises because we view ignorance as a bad thing. In this section I consider examples of 'responsible' ignorance, where someone is rationally, or even praiseworthy ignorant.

Presumably, there are ready examples that cut across the moral, practical, and epistemic domains. For example, when we think of someone who considers every possibility we might wonder if they have gone too far. Pessimists often fit this category. They tend to have in mind the worst possible scenarios even when those scenarios are far-fetched— "What if the plane crashes?" "Maybe the cliff will crumble underneath us," "I bet we will get a flat tire." It is difficult to determine how we should think about the pessimist. On the one hand, they may be prepared if any of these negative scenarios were to transpire. On the other hand, they are filled with useless worry that is practically demanding.

For the moral domain, we usually think that it is a good thing if we are ignorant about how to be a serial killer or are ignorant about how to torture someone. We would only think that someone needed this knowledge if they were trying to satisfy a purely academic curiosity.⁷⁴ The other likely reason is that they want to know how to murder or torture someone so that they

⁷⁴ Suppose that someone is interested in these thoughts for purely academic reasons. We might argue that without any intent to use the knowledge, such know-how is tantamount to having junk beliefs. For an argument that junk beliefs are irrational, see Friedman (2018).

can carry out these acts. Since both acts are viewed as morally impermissible, we should think that if one has no desire to carry out the act of torture, then one should be ignorant of torture know-how.

The moral cases provide examples for when it may be good to be ignorant. Are there also cases where we are praiseworthy ignorant?

A final example.

Matriculation: Lacy is a teacher at a ‘second chance’ high school with a class that predominantly consists of Black students. Lacy ignores the statistic that 95% of Black students in the district do not matriculate and does not consider this possibility when she is forming a belief about how likely a particular student in her class is to matriculate.

Ignoring the statistical evidence in this case looks importantly different from other cases. For example, if a hiker knows that 95% of deer ticks in the area carry Lyme disease, it is straightforward that the statistical evidence bears on the question whether a particular tick that bites the hiker is carrying Lyme disease.⁷⁵ And if the hiker fails to consider the possibility that a particular tick that bites them is carrying Lyme disease, they would be ignorant on my account and responsible for that ignorance. Lacy, however, may not be responsible. We can think of three different versions of her situation and compare the epistemic ‘grade.’

Suppose the first Lacy, I will call her ‘Lacy 1’, fails to consider the statistical evidence. It seems plausible that Lacy 1 is being responsibly ignorant—arguably, her ignorance is not only permissible, but also praiseworthy. Sarah Moss (2018) has a view that supports at least the

⁷⁵ Thanks to Kate Nolfi for this example and discussion of Matriculation.

permissibility of Lacy's ignorance. Moss thinks we should treat statistical evidence that bears on *Q* differently when considering people as opposed to (say) ticks. Moss's 'rule of consideration' is the idea that "we should treat people as individuals" (Moss 2018, 223) because we do not want to create expected outcomes for other people. Lacy 1 does precisely this, although unwittingly since she does not realize that she is ignoring the statistical evidence. Consider a second version—Lacy 2—where she is aware of the evidence, in the sense that it is on her pedagogical radar, but she doesn't take it as evidence that bears on *Q*. Lacy 2 invokes Moss's rule, and she may also be an exemplar of how Renée Bolinger (2020) thinks epistemic agents should consider rational stereotypes in belief formation. Bolinger's view is that even if some racist beliefs are based on statistical evidence, there are times when we should not accept these beliefs once we consider the practical harms.⁷⁶ In the third version of the case, Lacy 3 fully considers the evidence and takes it to bear on *Q* to the extent that it informs her belief about the particular student. From the perspective of both Moss and Bolinger, Lacy 3 looks to be the least justified because she fully considers and accepts that the statistical evidence bears on *Q*.

So, what can we say about the epistemic 'grade' for each of the versions? First, it seems Lacy 3 is unjustified in taking the evidence to bear on *Q* because she violates Moss's rule, as well as Bolinger's 'acceptance' claim. Even so, we cannot jump to the conclusion that in the other two versions Lacy is justified. Lacy's ignorance in version one looks praiseworthy because she is not pigeonholing her student to an expected outcome, but she is also unaware of information that may be relevant. In fact, it may be harmful to the student if she does not consider the very real possibility that students at her school do not matriculate. For these reasons, I think Lacy 2 exhibits the best epistemic predicament. She considers the statistical so that she is aware of the pedagogical issues faced by her students, but she does not take the information to directly bear on any particular student.

⁷⁶ 'Accepting' a belief, for Bolinger (2020), is an endorsement that the belief is practically 'good' to have.

The idea of ‘responsible’ ignorance suggests that there is not a one-size-fits-all recipe for examples of ‘good’ ignorance. My argument has been that ignorance is the kind of attitude that is subject to epistemic responsibility, which can have an epistemic valence that is either positive or negative.

6. Conclusion

The account I have given is part of a larger project on the ethics of ignorance, which aims to provide an agent-centered view of responsibility. People are not just rationally responsible because they *owe it to others* to be decent epistemic agents. Demands for rational responsibility arise because we also *owe it to ourselves* to be rational agents.

REFERENCES

- Alcoff, L. (2007). Epistemologies of Ignorance: Three Types. *In Race and Epistemologies of Ignorance*, Sullivan, S. & Tuana, N. (eds.) New York: SUNY Press.
- Anscombe, G.E.M (1957). *Intention*. Oxford: Blackwell.
- Berker, S. (forthcoming). The Deontic, the Evaluative, and the Fitting. In *Fittingness*, Howard, C. M. & Rowland, R. (eds.) Oxford: Oxford University Press.
- Bolinger, R. (2020). The Rational Impermissibility of Accepting (Some) Racial Generalizations. *Synthese* 197 (6): 2415-2431.
- Brown, J. (2020). What is Epistemic Blame? *Noûs* 54 (2): 389-407.
- Fine, K. (2018). Ignorance of Ignorance. *Synthese* 195 (9): 4031-4045.
- Friedman, J. (2013). Suspended Judgment. *Philosophical Studies* 162(2): 165-181.
- Friedman, J. (2018). Junk Beliefs and Interest-Driven Epistemology. *Philosophy and Phenomenological Research* 97 (3): 568-583.
- Goldberg, S. (2017). Should Have Known. *Synthese* 194 (8) 2863-2894.
- Goldman, A. (1999). *Knowledge in a Social World*. Oxford: Clarendon Press.
- Goldman, A. & Olsson, E. J. (2009). Reliabilism and the Value of Knowledge. In *Epistemic Value*.

- Haack, S. (2001). 'The Ethics of Belief' Reconsidered. In *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Steup, M. (ed.) Oxford University Press, pp. 21-33.
- Haddock, A. Millar, A. & Pritchard, D.H. (eds.). Oxford University Press, pp. 182-204.
- Guerrero, A. A. (2007). Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution. *Philosophical Studies* 136(1): 59-97.
- Hieronymi, P. (2008). Responsibility for Believing. *Synthese* 161(3): 357-373.
- Hieronymi, P. (2011). Believing at Will. In D. Hunter ed., *Belief and Agency*. Calgary: University of Calgary Press, pp. 149-187.
- Johnson King, Z. (2021). Praiseworthy Motivations. *Noûs* 54 (2): 408-430.
- Lackey, J. (MS). Epistemic Reparations and the Right to Be Known.
- Le Morvan, P. (2011). On Ignorance: A Reply to Peels. *Philosophia* 39 (2): 335-344.
- Levy, N. & Mandelbaum, E. (2014). The Powers that Bind: Doxastic Voluntarism and Epistemic Obligation. In Jonathan Mathewson (eds.) *The Ethics of Belief*. New York: Oxford University Press, pp. 12-33.
- Le Morvan, P. & Peels, R. (2016). The Nature of Ignorance: Two Views. In *The Epistemic Dimensions of Ignorance*. Peels, R. & Blaauw, M. (Eds.). Cambridge: Cambridge University Press.
- Medina, J. (2016). Ignorance and Racial Insensitivity. In *The Epistemic Dimensions of Ignorance*. Peels, R. & Blaauw, M. (Eds.). Cambridge: Cambridge University Press.
- Miller Larsen, E. (2023a). Responsibility for Ignorance.
- Mills, C. (2007). White Ignorance. In *Race and Epistemologies of Ignorance*, Sullivan, S. & Tuana, N. (eds.) New York: SUNY Press.
- McCormick, M. (2011). Taking Control of Belief. *Philosophical Expectations* 14, 2: 169-183.
- Moss, S. (2018). *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Peels, R. (2010). What is Ignorance? *Philosophia* 38(1): 57-67.
- Peels, R. (2017). *Responsible Belief*. Oxford: Oxford University Press.
- Siegel, S. (2017). *The Rationality of Perception*. Oxford: Oxford University Press.
- Setiya, K. (2013). Epistemic Agency: Some Doubts. *Philosophical Issues* 23 (1) 179-198.

- Srinivasan, A. (2020). Radical Externalism. *Philosophical Review* 129(3): 395- 431.
- Van Woudenberg (2009). Ignorance and Force: Two Excusing Conditions for False Beliefs. *American Philosophical Quarterly* 35 (138): 1-31.
- Williams, B. (1973). Deciding to Believe. In *Problems of the Self*. Cambridge University Press, pp. 136-151.
- Williamson, T. (forthcoming). Justification, Excuses, and Sceptical Scenarios. In F. Dorsch and J. Dutant (eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Zimmerman, M. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press

Part III

Ignorance in Perception

“That invisibility to which I refer occurs because of a peculiar disposition of the eyes...A matter of the construction of their *inner* eyes, those eyes with which they look through their physical eyes upon reality.”

— *Ralph Ellison, Invisible Man*

1. Introduction

It is widely held among philosophers that ignorance serves as an excuse for *omissions* that would otherwise be responsible.⁷⁷ This position is usually illustrated by thinking about the role of ignorance in moral actions. Take for example a case where I unknowingly give a friend, who is vegan, a piece of chocolate cake containing dairy. My action is immoral, in the sense that I have caused my friend to imbibe food that contains ingredients they morally disagree with, but in another sense, we can question whether I am truly responsible since I did not know about the dairy in the cake. Typically, on this picture of ignorance, ignorance fully excuses my act because we are not thought to be responsible for those things that we do not know. For many of us, that includes an extensive list of facts, subject matters, and skills, such as the exact latitude of Boston, the whereabouts of all your colleagues at this very moment, and (perhaps) the entire body of literature that is Rawlsian political philosophy.

Turning to belief, we find a similar story about ignorance as prejudice: ignorance (sometimes) serves as an excuse for prejudiced *omissions*. Unlike the position in the previous paragraph, no philosopher explicitly endorses this view, but it is a view that is clearly operating in the

⁷⁷ The idea that ignorance serves as an excuse or renders one blameless for omission can be traced to Aristotle in the *Nicomachean Ethics* (NE III.i 13-27, v.7-12, V.viii.3-12). The idea can also be found in contemporary work by Brandt (1969), Fischer and Ravizza (1998), Goldman (1970), Rosen (2003), Smith (1983), and Zimmerman (2008). In some cases, ignorance acts only as a partial excuse and doesn't fully render a person blameless, see Peels (2014) for such an account. On the strict view of the partial view the primary role of ignorance in epistemic discourse has been as a way to excuse agents from full epistemic responsibility.

prejudice literature. The view is expressed in numerous ways. Sometimes, it is operative within a tension model between rational beliefs on the one hand, and the moral consequences of those beliefs on the other.⁷⁸ Take for example the case of a woman who mistakes a black guest at a country club as the coat check attendant.⁷⁹ She fails to see that he is a guest, and her failure is a type of omission that occurs due to a set of prejudiced background beliefs. Her omission, however, is often seen as epistemically excusable due to the overall social environment in which her prejudice functions. In particular, the view is often highlighted through thinking about the tension that arises due to the fact that epistemic agents live in a society that “on the one hand, [has] an explicit ideal of racial equality, and, on the other, bears the legacy of centuries of personal and institutionalized racial discrimination” (Gendler 2011: 44). What happens in this particular context is that epistemic agents are reasoning with evidence that is distorted by the very societies in which they live so it is less clear whether they should be held epistemically responsible for the kinds of information they lack.

At other times, ignorance as an excuse for prejudice undergirds explanations of prejudiced belief within probability theory. Epistemic agents are said to be no more than “culturally situated cognitive engines exhibiting (or failing to exhibit) certain checks and balances in the use of sensory evidence and priors” (Clark 2018: 747). The allocation of epistemic responsibility has no standing within this Bayesian schema because beliefs are simply the result of probabilistic calculations that are updated given information from the world. Simply put, in a society that is prejudiced, it is impossible to be without prejudice, or at least, if (and when) your thinking is contaminated by prejudice, it is not your fault.

This is the view that (at least) seems to be endorsed by Bayesian perceptual processing.⁸⁰ Prejudiced omission is simply a feature of our cognitive systems that are mostly viewed as

⁷⁸ See, for example, Gendler (2011).

⁷⁹ This case is also from Gendler (2011).

⁸⁰ Especially as endorsed by Clark (2018).

optional, thus, working fine, or as fine as they can be expected to work. I claim the entire setup of ignorance distorts the philosophical issue of prejudice. Moreover, the misstep originates from an understanding of ignorance as simply the lack of knowledge. This understanding cannot adequately account for the nature of ignorance for prejudiced omissions. Plants lack knowledge, rocks lack knowledge, and (a lot of times) so do I, but in these cases, lacking knowledge does not demand epistemic responsibility.

Recognizing how ignorance might demand epistemic responsibility requires a conceptual shift from the idea of ignorance as a type of epistemic *lack* to the adoption of “the kind of ignorance that involves a whole battery of mechanisms of avoidance and resistance to know and to learn” (Medina 2018: 182). These mechanisms of resistance are typically identified as social mechanisms, or structures of oppression and hegemony that arise directly from the social and material world. In this paper, I focus instead on psychological mechanisms of resistance. Psychological mechanisms that are operative in perception. I highlight a phenomenon that I call Perceptual Ignorance to serve as an example of the identified need for the conceptual shift from ignorance as the lack of knowledge, to a notion of ignorance that is *active* and normatively laden.

Perceptual Ignorance, I claim, occurs when perceivers fail to take in relevant information and fail to register that information when forming beliefs. In short, they omit a hypothesis.

Ignorance_{omission} : S omits *h*, where *h* is a hypothesis missing from S’s hypothesis set.

In perception, ignorance takes the form of a literal blind spot. So, the normative question, for the purposes of this paper, is the question of whether we are epistemically responsible for not seeing. Overall, this discussion aims to bring together many difficult and philosophically

contentious materials, i.e., ignorance, perception, prejudice, and epistemic responsibility, to illuminate a phenomenon that has serious import for our epistemic lives.

2. Problems of (in)sight

To sharpen the issue at hand, consider the following riddle. (Give the riddle a try, before reading on for the answer.)

Accountant	An attorney says, “That accountant is my brother”, and that is true—they really have the same parents. Yet that accountant denies having any brothers—and that is also true! How is this possible?
-------------------	--

The riddle above is an insight problem that generates a kind of cognitive block or blind spot to a correct answer. These blind spots arise from a reader’s inability to recognize and fill in missing details, usually due to underlying biases that inform dominant perspectives. In the Accountant, readers tend to identify any unmarked or underspecified protagonist as male. The evident gender bias then distracts readers from the possibility that the accountant is a woman.

Insight problems, like the one above, display an interesting cognitive effect. The important feature of insight problems is that readers are completely stumped, meaning they are typically unable to produce an answer at all! And if they do produce an answer, they often overlook the correct solution in favor of ones that are extremely implausible. For instance, to solve the riddle some readers claimed that the salutation ‘brother’ has more than one meaning, which is why one accountant can call the other a ‘brother’ while the other one denies having any brothers. Or readers solved the riddle by reasoning that the attorney must have *died* between utterances.⁸¹

⁸¹ See Bar-Hillel et al. (2018).

There are two explanations for what is happening at the psychological level when we are stumped by insight problems. The first explanation is that readers are resistant to disconfirming information.

2.1 Resistance to Disconfirming Information

Let us start with the response that readers are resistant to disconfirming information. On many accounts, such resistance is not synonymous with irrationality. Take for instance discussions about practical reasoning. There are those that argue that in some cases of practical reasoning resistance to disconfirming information is the rational response. It is rational to trust that your marriage will be successful despite statistics of the rising divorce rate.⁸² It is rational to believe that your friend is an excellent professor despite poor ratings from students in their courses. It is rational to believe that you will not die despite the fact that a doctor has just diagnosed you with a fatal illness.⁸³ In each of these cases it is argued that rationality, in particular, rationality in the face of practical considerations, can demand that we are resistant to disconfirming information.

Practical considerations are just that, considerations. They are reasons that epistemic agents consider when forming beliefs. The psychological mechanisms at work in insight problems, however, may not be considerations. In fact, it is likely that the mechanisms are non-conscious. This means that when someone is unable to solve the riddle and provide an answer that the account is a woman, it is not because they considered this point and resisted it. It is because the solution never came to them. So, resisting disconfirming information looks quite different from weighing considerations. It often takes the form of a type of dissonance where one is committed to preserving antecedent beliefs.

⁸² See Marušić, B. (2013).

⁸³ See Rinard, S. (2018).

The belief preservation theory does not hold up as an explanation for insight problems, however. For one, it is unlikely that readers are resistant to the very idea that a woman can be an accountant. Or that they are trying to preserve a long-held belief that accountants cannot be women. In many cases, readers report having antecedent beliefs that align with the solution. They believe that accountants can be women, and in some cases know women who are accountants.⁸⁴ But when presented with the riddle the right solution is still not activated. So, it is evident that readers are not resisting disconfirming information due to a need to preserve antecedent beliefs.

Another form of resistance that might be at work also arises because of one's background beliefs. Suppose that the reason readers are unable to solve the riddle is because in their experience few women are accountants. So, it is not the case that they do not think women can be accountants at all. Rather they have a degraded confidence in the belief. In fact, it is likely that they assign a low degree of confidence that women can be accountants that is robust because of the support it has garnered over the years. This means that when the person finally encounters a female accountant, they are resistant (non-consciously) since the new evidence goes against their robust priors.

There is reason, however, to reject the Bayesian explanation in this case since the distribution of confidence does not conform to the ecological background information with respect to female accountants. Therefore, we need another explanation to account for the resistance to new information.

2.2. Resistance to a Hypothesis

⁸⁴ In fact, 60% of all accountants in the United States are women.

Another possible explanation, which this paper will favor, is that readers are resistant to forming a hypothesis. This explanation proposes that we can have blind spots to information, not because we lack knowledge or are having a retrieval issue due to faulty memory, but because our minds actively resist certain hypotheses. Insight problems, in particular, suggest that we can actively resist correct solutions despite having otherwise rational background beliefs. When we are unable to produce the solution, it is because the hypothesis that the accountant is a woman has not occurred to us. Here, not-occurring does not mean that we have never considered it. It simply means that in this particular moment the hypothesis was not activated. And there is particularly good reason to question why activation does not occur for people who antecedently think that accountants can be women, since there is no anterior reason for why they should be resistant.

The rest of the paper is dedicated to unpacking these blind spots. I argue that the structure I have outlined for insight problems is similarly visible in cases of perception, and is, in fact, the paradigmatic structure for the phenomenon I call Perceptual Ignorance. Insight problems are a type of *mental* blind spot and Perceptual Ignorance is a type of *perceptual* blind spot. Both phenomena generate the normative question of responsibility. And both suggest that the epistemic agent or perceiver *should have known* or *should have seen*.

3. What is perceptual ignorance?

Like the general notion of ignorance, where I can be ignorant about any number of facts, experiences, theories, conditions, or state of affairs, ignorance in perception appears to be far reaching. We theorize about ignorance in perception as errors, such as illusions or hallucinations. Seeing a stick in the water as bent or seeing an oasis in the middle of the desert when none exists are ways that we can be ignorant in perception that is not necessarily Perceptual Ignorance. Much work has been dedicated to theorizing about both illusions and hallucinations, both of which are perceptual errors that can arise despite a perceiver's best intentions to view the world 'properly.' Take the stick in the water for example. I can know that

a stick in my hand is straight (or as straight as sticks can be), but when I place it in the water, I will still *see* it as bent. Like ignorance, illusions are often viewed as an excuse. They are examples where perception goes awry.

Perceptual Ignorance requires a unique suite of examples to illustrate how it is different from stick-in-the-water type illusions. I provide three such examples below.

Who's the Professor? On the first day of classes, Finn—a usually excellent student—is late to his intro to philosophy course. Thinking he has already missed valuable information; Finn frantically searches the room for the course professor. He considers several people in the room, including a male student in the back wearing Oxbridge-like attire. However, Finn fails to consider the woman at the board, despite her performing duties that would normally indicate she is the professor.

Capitol Hill: Tim Scott is an African American Republican Senator who is consistently denied access to many buildings on Capitol Hill. When Scott enters a building, Jon, a white male security guard, determines that he is not a senator and calls the police. Jon *sees* that Tim Scott is wearing a suit and has a senator pin but does not *register* that Tim Scott is a senator.⁸⁵

Bear with me for one more example. This example is critical for the overall characterization of Perceptual Ignorance as having a structure that is neutral with respect to the content.

⁸⁵ This example is derived from reports by Tim Scott about his experiences on Capitol Hill. Scott has been the victim of profiling and has been pulled over at least 7 times in one year. See the following articles for the full story: <https://www.npr.org/2016/07/14/485995136/watch-black-gop-senator-says-hes-been-stopped-7-times-by-police-in-a-year>; <https://www.cnn.com/2016/07/13/politics/tim-scott-police-racial-profiling/index.html>

Changing Shape: Beverly, a participant in a change blindness study, is tasked with sorting shapes by size. She knows that large shapes should be stacked in one pile and small shapes stacked in another. But when it comes to performing the task, Beverly cannot register the fact that shape sizes change mid-task, despite seeing them do so.⁸⁶

How are these examples like the insight problem from the previous section? In the Accountant, readers were provided with a riddle that was difficult for them to solve because of how they construed the vignette. The dominant construal relied on stereotypes that take the form of gender biases. In the examples above we can note a similar structure. Perceivers have an experience, and the dominant construal of the experience interferes with successfully completing a cognitive task. This time the task is not solving a riddle but answering a question about how one should see the world given their position within it. For Who's the Professor? the task is to identify the professor, for Capitol Hill the task is to let senators into a building, and for Changing Shape the task is to sort bricks based on size. But Finn, Jon, and Beverly are not successful at their tasks because they are relying on a dominant construal, yielding an omission that produces unstable patterns of reasoning.

No singular example by itself is meant to ground the phenomena of Perceptual Ignorance. But taken together we recognize how people exhibit perceptual omissions. In fact, the idea that people miss information in their visual environments is not radical. Countless articles on inattention and different forms of 'blindness' have suggested such omissions are possible and occur regularly for competent perceivers. The radical contribution from this paper is that sometimes perceivers omit information that seems the most relevant to their experiences in a way that is odd and puzzling. Whereas in classic inattention and perceptual blindness cases

⁸⁶ Based on the Triesch et al. (2013) study.

perceivers fail to pick up information due to some form of cognitive overload where they are asked to juggle multiple tasks.⁸⁷

In addition, one might think that what is happening in these cases is resistance to disconfirming information. For instance, in *Who's the Professor?* Certain socio-historical facts make it unlikely that the woman is the professor. Suppose Finn has a robust background belief that only men wearing Oxbridge attire are professors. Seeing the woman would present disconfirming information that would not be able to infiltrate such a robust background belief. *But what if Finn does not have any outright prejudiced beliefs?* Perceptual Ignorance can explain why Finn has blind spots even without maintaining an explicit belief. The cases of Perceptual Ignorance illustrate the following idea. In cases of prejudice perceivers are resistant to forming a hypothesis.

Ignorance_{omission}: S omits h , where h is a hypothesis missing from S's hypothesis set.

I will show how this resistance occurs, and highlight how Perceptual Ignorance, as an epistemic failing, is worrisome for epistemic accounts of normativity in general, and how it is particularly worrisome for the purported normative application of Bayesian perceptual processing.

4. A Bayesian explanation of perceptual ignorance

The simplest explanation of Bayesianism is that it is credence in a hypothesis. Bayesian perception is the idea that background hypotheses information perceptual experience through the combination of past ecological information and present distal stimuli. Taking the surge of research seriously that argues that perceptual reasoning approximated Bayesian probability⁸⁸,

⁸⁷ See §6 for a straightforward case of cognitive overload and attention.

⁸⁸ See Raheev & Denison (20218) for a brief survey of this research.

should make us wonder how exactly such reasoning accounts for social stimuli alongside other environmental features.

Habits of perception, according to Bayesian models, are governed by hyperpriors. A hyperprior acts as a general rule or hypothesis that underwrites, or governs, other priors. Take for example the hyperprior that there can only be one object, in one place, at any given moment—the spatiotemporal hyperprior. The spatiotemporal hyperprior as a rule governs the perceptual system by dictating how much information is visually processed at any one point in time. The spatiotemporal hyperprior also serves as a potential explanation for binocular rivalry. In the most well-known example of *binocular rivalry*, subjects in an experiment are simultaneously shown an image of a house and an image of a face. Rather than fuse the two images, the perceptual system alternates between them. So, one explanation may be that spatiotemporal hyperpriors govern hypotheses about houses and faces, limiting cognitive overload.

Another hyperprior is the rule that light comes from above. A reasonable phylogenetic explanation for this rule comes from our relationship to the sun dating back some 2.5 million years ago. An ontogenetic explanation includes our individual experiences with light—natural or fluorescent. So, hyperpriors, in their most basic state, are biases or prejudices. But they are ones that have evolutionary import and have proven useful for our development as organisms. Hyperpriors are, therefore, phylogenetically robust, but are also supported by experiences of individuals. They exhibit a type of stubbornness and recalcitrance that is neither seen as problematic nor irrational. In fact, hyperpriors are often thought of as optimizing features of our cognitive systems.

From here, one natural proposal is to think that hyperpriors explain the stubbornness in Perceptual Ignorance. Suppose that in *Who's the Professor?* Finn has priors about the look-profile of a professor, and that these priors drive him to ignore disconfirming information from his experience. Finn's hyperpriors could reasonably include racist or sexist beliefs that govern

which priors get represented in perception. Moreover, these hyperpriors, which have been enforced for generations, not just from Finn's own experience, but also from the experiences of his ancestors, boost the level of credence that he should hold in them.

Such a view would be mistaken. For one, the proposal that racist and sexist background beliefs are hyperpriors seems to lack the phylogenetic robustness of the hyperpriors previously mentioned. Stereotyping regarding the look-profile of a professor does not have the same long-standing precedent that spatiotemporal priors, or the light-from-above prior seem to carry. Finn's priors have a social origin, not a purely neurobiological one. It seems obvious that the kind of hyperpriors championed by Bayesians are in a completely different category from the kind of racist or sexist beliefs that are at issue in *Who's the Professor?* With this said, we should reject the proposal that the cases of Perceptual Ignorance mimic the kind of resistant-yet-optimal functionality of hyperpriors.

Instead, I suggest that the ways in which perceivers structure their hypothesis spaces illuminates the problem of Perceptual Ignorance. A hypothesis space is the set of all hypotheses as to what is in the world. Hypothesis spaces are an important feature of Bayesian models that express how relevant information is registered in experience. For our purposes, I argue that they can also express hypotheses that have been omitted.

For example, look at the following image:



Now, take a guess as to what the object could be.

We can characterize your guesses as individual hypotheses about what you are likely to see. These hypotheses operate as the set of non-conscious possible representations of the object. In this case, you might generate a set that includes the following:

$$H = \{H_1: \text{pliers}, H_2: \text{hammer}, H_3: \text{sickle}\}.$$

Within the set, hypotheses are given a probability weighting, which are assigned based on the available information. The Bayesian perceptual system then selects a hypothesis from the set by determining the hypothesis with the highest posterior probability. This hypothesis ‘wins out’ and is then represented in your experience.

Notice that ecological determinants are the main factors that restrict the set. Given the image, the set could not include hypotheses that represent (say) faces, or hypotheses that represent flora. The set is generated based on the available information, which serves as a foundation for the generative mechanism. In other words, a belief that the image is a hammer is the result of the combination of your past background beliefs and your immediate experience of the image. Also, notice that your guess is not an exercise in classification. It is one thing to attempt to *classify* information and another to *register* information. When you determine that the object is a hammer it is because you are seeing it and taking in information about the image’s shape, size, and form that generates the hypothesis that the image could be a hammer.

The Bayesian perceptual system can provide straightforward explanations for object-based cases like the hammer above. It is clear how the hypothesis set is generated, and what determinants bear on the hypotheses included in the set. But this paper is not concerned about objects. The realm of study is social, and the worry is (mostly) about hypotheses about human subjects. For the social realm, hypothesis sets are constrained by more than just ecological factors. I will suggest that our priors, as hypotheses, are not only restricted by their ecological *context*, but are also restricted by how we structure experiences. When that structure includes

perverse racial or sexist background beliefs, sets can be faulty by failing to include relevant hypotheses. In other words, perceivers exhibit an epistemic failure because they are resistant to forming a hypothesis.

5. Are missing hypotheses a form of ‘resistance’?

How are perceivers resistant to forming a hypothesis? To get a sense of how this failure occurs, suppose that in Capitol Hill Jon generates the following set of hypotheses when he sees Tim Scott walk into the building.

$$H = \{H_1: \text{terrorist}, H_2: \text{thug}, H_3: \text{disgruntled citizen}\}$$

Notice what is included in the set and what is missing. Notice H_1 and H_2 . These hypotheses are activated by homing in on skin color as the dominant information. In any experience there are many pieces of information that one can attune to, but we can still question why Jon focuses on skin color to the extent that he ignores the available sartorial information. In addition to race, he should have been aware of Tim Scott’s clothing and other accessories that indicate that Scott is a senator. For some reason, this information is not registered as relevant by Jon so that it cannot activate the appropriate corresponding hypotheses. Notice, then, that the hypothesis ‘senator’ is missing as a hypothesis in the set. If we suppose that Jon is a neurotypical perceiver, who tends to be successful at recognizing information in his environment, we can question why ‘senator’ is not a possibility that came to mind.

One response may be that Jon’s hypothesis set *does* include the possibility that Tim Scott is a senator, but the hypothesis is simply given an unfairly low weight. This is the response Gendler (2011), and Siegel (2017) endorse in their treatment of cases with stereotype-based belief. Gendler, for example, has pointed out that features, such as race, are often given biased probability assignments in cases of prejudice and stereotypes, resulting in strong one-off

predictions. Similarly, Siegel argues that imbalanced weighting of perceptual features results in faulty inference that transit poor warrant to perceptual experience. So, epistemic agents in Gendler and Siegel cases do not entirely omit information but instead fail to give certain information proper attention when assigning credence. These agents assign low base-rates and underweight priors, which then means that other beliefs, in this case, the belief that Scott is *not* a senator, is given a skewed weighting.

In some cases, it does seem to be the case that individuals underweight or overweight hypotheses, but such an explanation does not account for all cases of Perceptual Ignorance. For instance, one reason to think Capitol Hill is an example of hypothesis omission (rather than distorted weighting) is because updating does not occur. Proponents of Bayesian models tend to motivate the framework by relying on cases where epistemic agents update their beliefs. And the benefit of Bayesian models is that updating can occur by conditionalizing on new evidence. This means that Jon can start with a probability weight for hypotheses, such as H_i , and update the probability as he obtains new information given distal stimuli.

Updating can only occur, however, if a hypothesis *exists* in the set to start with. This is a familiar obstacle for Bayesian models and is known as the old evidence problem. For instance, before Einstein introduced his theory of general relativity, the theory, as a hypothesis H_{GR} was (reasonably) missing from the set that came to mind when scientists observed physical phenomena. So, any old evidence that a scientist might have about the physical world could not be taken as support for H_{GR} because the old, stored set of information cannot bear on how credence in Einstein's new theory is updated. Some have recognized the obstacle within the problem of abduction, and it has recently been proposed as a problem of 'epistemic imagination'.⁸⁹

⁸⁹ See Perez Caballo (ms) for details on epistemic imagination.

For our purposes, the obstacle puts pressure on the plausibility of updating. Jon cannot update if the hypothesis does not exist in his set. Unlike pre-Einsteinian scientists, Jon's lack of epistemic imagination is impermissible. No one prior to Einstein could have introduced Einstein's theory; this is logically impossible. But it is possible that other security guards could see and register that Tim Scott is a senator. In Capitol Hill the old evidence problem for Bayesian models does not seem to apply.

6. What are the normative demands on perceptual ignorance?

A true Bayesian might note that since Bayesianism, at its core, is about rational, ideal agents, I have already done something classically non-Bayesian by supposing that we can model irrationality on a Bayesian framework. Moreover, my concern about relevant hypotheses being included in hypothesis sets does not arise in orthodox Bayes because it is assumed that all relevant hypotheses are available to agents. It is typically taken for granted that in any belief calculation, information that bears on the solution can be retrieved. But this is just not how real, *fleshy*, humans operate. We are obviously bound by the information that we have been exposed to, and imperfect at deriving beliefs even when all available information is considered. In effect, Perceptual Ignorance puts pressure on the orthodox account to reconcile cognitive and epistemic imperfections that cannot be smoothed out by updating.

It is this pressure that illustrates how cases of Perceptual Ignorance generates the normative question of responsibility. One way that we can identify the unique epistemic error in cases like Who's the Professor? Capitol Hill, and Changing Shape is by specifying what the error is with respect to. The solution I offer is to evaluate the rationality of cases of Perceptual Ignorance by specifying a cognitive task and determining how well agents complete the task by reasoning with relevant hypotheses.

Case	Cognitive Task	Available Visual Cues	Hypothesis Set H_i-H_n	Missing Hypothesis H_M
Who's the Professor?	Identify the prof	prof duties, skin color, gender ¹⁸	H1: undergraduate H2: admin assistant	H_M : TA
Capitol Hill	Determine members of Congress	suit, senator pin, skin color	H1: terrorist H2: thug H3: disgruntled citizen	H_M : senator
Changing Shape	Track shape size	sizes of bricks	H1: small brick H2: large brick	H_M : brick sizes change

Once we specify the cognitive task for each case, the generated hypothesis set is strikingly *irrational*. Since the evaluation relies on how well an epistemic agent completes a cognitive task, epistemic agents are irrational when they ignore visual cues that are relevant to the task. We can also think of this evaluation method as one that outlines salience. If your cognitive task is to identify the professor, then when someone performs relevant professor duties, these perceptual acts should be salient to you. If your cognitive task is to determine who is a senator, then information that indicates someone is a senator should be salient. And, if your cognitive task is to track brick size, then changes in brick size should also be salient.

Epistemic agents who exhibit Perceptual Ignorance are, therefore, required in these cases to take certain information as salient. Notice that this rational demand does not apply in cases of inattention. A prime example is the famous gorilla study.⁹⁰ In the study, participants are asked to count the number of passes of a basketball. During this task, a person in a gorilla suit walks

⁹⁰ See Chabris & Simons (2010).

across the court. Surprisingly, most participants are unaware of the gorilla! While there is a type of 'blindness' that occurs in this study, demands of rationality do not arise as they do for examples of Perceptual Ignorance. For one, unlike Perceptual Ignorance, the study does not illustrate how someone has failed to recognize salient information. In fact, the participants perform well at the cognitive task they are given, since participants miss seeing the gorilla precisely because they are focused on counting the basketball passes. The gorilla study exemplifies the tradeoffs that occur when cognitive loads are high. It is reasonable for us to miss information when the cognitive demand exceeds our limited processing capabilities. But it is not reasonable for us to miss information when the *very* cognitive task at hand has not been given proper attention. So, when a perceiver is perceptually ignorant, they are irrational because they have failed to comply with an epistemic duty (i.e., the cognitive task) and have failed to see what they ought to see.

To close the paper, I would like to say more about my third case of Perceptual Ignorance. So far, I have said little about Changing Shape. This has been strategic. Changing Shape is a fascinating case because it does not have moral implications that can interfere with our intuitions.

However, like Finn and Jon, Beverly fails to epistemically do what she epistemically ought. She is meant to track brick size but it resistant to tracking the size. Naturally, we want to know if Beverly is responsible. We say that someone is *not* responsible for an action or belief if one of the following holds. (1) If we can think that someone with the same perspective would reason similarly, or (2) if we can think that the person could not have reached a different conclusion given their perspective.⁹¹

Starting with (1), it seems that, as the study illustrates, many perceivers would not register that the sizes of the bricks changed. This may be since we do not think that objects, on their own, can change size, so the virtual reality component of the study places perceivers in unfamiliar

⁹¹ Alvarez & Littlejohn (2017) for example, argue for these two conditions of responsibility.

territory. Beverly's response is also due in part to the fact that the brick sizes changed shape *after* she decides. Typically, people would not be prepared to think that objects can change size mid-task. (2) seems to follow from the considerations for (1). If a perceiver is unfamiliar with certain physical possibilities, then it seems unlikely that they could have reached different conclusions given their limited perspective.

We might wonder, however, whether participants like Beverly can, in fact, update their beliefs about how objects change size. Suppose all participants are told two facts and the virtual reality simulation. The first is the fact that the objects on their own could change size. The second is the fact that objects can change size mid-task. With this information, would participants be able to update their beliefs? Or would their experiences resemble Beverly's? Suppose that the idea that bricks do not change size on their own is a hyperprior. Updating would then be restricted, or impossible. We might also think that participants can experience seeing the bricks change size, yet still arrive at false beliefs about their experiences. Finally, it could also be the cases that participants still omit the hypothesis altogether, failing to update even with the two new pieces of information. This thought brings us back to the problem of a failure to update, which is a major defect of Bayesian perceptual processing.

7. The limitations of Bayesian perception for normative assessment

Suppose we take it as true, as the Bayesian says, that prejudice and Perceptual Ignorance are just the result of our cognitive machines mirroring a flawed and unjust world that is no way prohibits our cognitive systems optimization. Even so, we can put pressure on a perceiver's inability to update their background beliefs when they omit a certain hypothesis. There are two ways to understand hypothesis omissions in the Bayesian framework. One way to think of it is as having assigned the hypothesis a probability of zero. The other, as I have argued, is to omit the hypothesis altogether. Addressing the first understanding, suppose the hypothesis has been assigned a probability of zero. This would mean that a perceiver has, in fact, considered the

hypothesis but thinks that the possibility impossible— like the possibility of humans flying unaided, or the possibility of the existence of a round square. Perceptual Ignorance cases are not akin to these examples of logical impossibility. Some hypotheses are not logically impossible—like the possibility that women are professors or accounts or hold professional career in general—so, omitting them signifies a failure of rationality.

Here's what is left for a Bayesian theorist. Since updating is the fundamental component of an epistemic view that relies on probability theory, we must reject it as a useful means to talk about the rationality of these cases. Proponents of Bayesian perceptual processing can either concede that Perceptual Ignorance, and other examples of prejudice, present serious obstacles to the idea that Bayesian cognitive operations are always optimal. Or they can give up on probability theory as more than an ideal theory.⁹²

By focusing on optimization, bayes perceptual processing is not a useful normative tool. The purpose of this paper is to argue that there are rational norms that govern habits of ignorance. These norms would tell us when it might be wrong to be ignorant, and when we might be responsible for our ignorance. Much philosophy has focused on what we can know, what we can know we know, and even the degrees to which we can know. Now, it is time to turn a keen eye towards ignorance, not simply as the lack of knowledge, but as inimical to both reasoning and perceiving.

REFERENCES

Aristotle (2003). *Nicomachean Ethics*. Translated by Rackham, H. Cambridge MA: Harvard University Press.

Alvarez, M. & Littlejohn, C. (2017). When Ignorance is No Excuse. In P. Robichaud & J. W. Wieland (eds.) *Responsibility- The Epistemic Condition*. Oxford: Oxford University Press.

⁹² In other words, perhaps the domain of bayesian processing is for machine learning, artificial intelligence, and other program design but does not extend as a complete explanation for the psychological processes of humans who are both brilliant and ignorant at the same.

- Bar-Hillel, M. Noah, T., & Frederick, S. (2018). Learning Psychology from Riddles: The Case of Stumpers. *Judgement and Decision Making* 13 (1): 112-122.
- Brandt, R. B. (1969). A Utilitarian Theory of Excuses. *Philosophical Review* 78 (3): 337-361.
- Chabris & Simons (2010). *The invisible gorilla: And other ways our intuitions deceive us*. Crown Publishers/Random House.
- Clark, A. (2018) *Surfing Uncertainty*. Prediction, action, and the embodied mind. Oxford University Press
- Fischer and Ravizza (1998). Triumph of tracing. In J.M. Fischer *Deep control: an essay on free will and value*. Oxford University Press: pp. 206-234.
- Gendler, T. (2011). On the epistemic cost of implicit bias. *Philosophical Studies* 156: pp. 33-53
- Goldman, A. (1970). *A Theory of Human Action*. Princeton University Press.
- Marušić, B. (2013). Promising Against the Evidence. *Ethics* 123. 2 (2013): 292-317.
- Medina, J. (2018). Ignorance and Racial Insensitivity. In *The Epistemic Dimension of Ignorance* Peels, R & Martijn Blaauw (eds.). Cambridge University Press.
- Mills, C. (2007). White Ignorance. In *Race and Epistemologies of Ignorance*, Sullivan, S. & Tuana, N. (eds.). New York: SUNY Press, pp. 11-38.
- Peels, R. (2014). What Kind of Ignorance Excuses? Two Neglected Issues. *Philosophical Quarterly* 64 (256): 478-496.
- Perez Caballo, A. (ms). *New Boundary Lines*.
- Rahnev, D. and Denison, R. N. (2018). Suboptimality in perceptual decision making. *Behavioral and Brain Sciences* 41, e223: 1-66.
- Ramachandran, V. S. (1988). Perception of shape from shading. *Nature* 331, pp. 163-166.
- Rinard, S. (2018). Believing for Practical Reasons. *Nous*.
- Rosen, G. (2003). Culpability and ignorance. *Proceedings in Aristotelian Society, New Series* 103(1), 61-84.
- Smith, H. M. (1983). Culpable Ignorance. *Philosophical Review* 92 (4): 543-571.
- Triesch et al. (2003). What you see is what you need. *Journal of Vision*.
- Zimmerman, M. J. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.