

# How Could You Be So Oblivious?

## Positive Epistemic Duties and Oppressive Ignorance

DRAFT – March 2025 – Annette Martín

### 1. Introduction

Zarak is sitting in a coffee shop trying to finish a problem set. He has been puzzling over a derivation for hours but keeps hitting a dead end, so he once again tries running it through from the beginning. Meanwhile, the coffee shop has filled with smoke and panicked customers scramble to exit the building. Eventually one of the workers shakes him and tells him they need to evacuate immediately. Zarak, shocked, quickly gathers his papers and rushes out the door.<sup>1</sup>

Now take another case: Valerie considers herself passionate about social justice issues, and has been working to improve the gender balance in the philosophy department for years. However, it is not until she hears someone from history talk about what it is like to be the only person of color in their department that Valerie realizes that there are no people of color in the philosophy department.<sup>2</sup>

The focus of this paper is a class of phenomena illustrated by these cases, which I will term *obliviousness*.<sup>3</sup> Zarak is oblivious of the fact that the cafe has filled with smoke; Valerie has been oblivious of the fact that all the members of her department are white. The most distinctive feature of this class is, roughly, that there is some fact that should have been obvious to the

---

<sup>1</sup> I will refer to this case as *Smoky Cafe*.

<sup>2</sup> I will refer to this case as *Racial Blindness*.

<sup>3</sup> To clarify, my goal is not to give an analysis of the term or concept of obliviousness. Rather, my interest is in the class of phenomena illustrated by the cases highlighted in the paper, which I think can be aptly termed “obliviousness.” This phenomenon is in some respects narrower than the range of phenomena the word “obliviousness” is standardly used to refer to. For example, we might say that someone with noise-cancelling headphones on is oblivious to the office cacophony. This is not, however, a case of obliviousness in my sense, precisely because the headphones render the noise outside of the person’s perceptual range in this context. At the same time, the class of phenomena that I am interested in is quite broad in certain respects, and may extend beyond standard use of the term “oblivious.” In particular, it includes many specific kinds of cognitive failures (e.g. a failure of perceptual awareness in Zarak’s case vs. a failure of belief closure in Valerie’s case). I think obliviousness is worth studying as its own topic in spite (and because) of this breadth because (1) the distinctive sense that there is something that should have been obvious to the agent is useful for exploring questions relating to epistemic responsibility and (2) it presents an instructive contrast to the kinds of cases often discussed in the epistemologies of ignorance literature. Thank you to an anonymous referee for pressing me on this point.

relevant agent that the agent nevertheless genuinely “misses.” In this paper, I will analyze obliviousness in more detail, suggesting that cases of obliviousness are ones in which an agent fails to take a rational route to belief that is immediately available to them that they should have taken, although this failure is not deliberate on the part of the agent. In most cases, obliviousness results in ignorance – in not taking the relevant route to belief, the agent fails to come to know that which should have obvious to them. However, it could happen that an agent comes to know in some other way, say because someone informs them that *p*. Nevertheless, cases of obliviousness typically manifest a kind of ignorance that is distinctive, again, in that the agent “misses” something that should have been obvious to them.<sup>4</sup>

Obliviousness is, I think, familiar to our everyday lives. I suspect that many will have experienced the embarrassment, guilt, incredulity, amusement, irritation, or anger that can come in the wake of realizing that someone – whether oneself or another – has failed to grasp that which should have been obvious to them. These common (though not definitive) emotional responses indicate that a charge of obliviousness typically comes loaded with epistemic criticism. In wondering, e.g., how Zarak could be so oblivious as to not realize that the cafe was filling with smoke around him, or how Valerie could be so oblivious as to not recognize the total lack of people of color in her department, we seem to be pointing to an epistemic failure on behalf of the oblivious agent. On my analysis, this epistemic failure will correspond to the claim that in cases of obliviousness, there is a rational route to belief that the agent *should have taken*.

However, the claim that a charge of obliviousness comes loaded with *epistemic* criticism is controversial. In particular, there has been skepticism regarding the idea that we have any positive epistemic duties; it has been thought that while there are things we are epistemically permitted to believe, there is nothing that we epistemically *ought to believe*.<sup>5</sup> In contrast, the

---

<sup>4</sup> As a result, cases of obliviousness will often, but not always, fall under the scope of what Goldberg calls the *should have known* phenomena. See (Sanford C. Goldberg 2017). However, even if some cases of obliviousness do not technically count as “should have known” cases, I think that obliviousness helps demonstrate that we don’t always just care about whether or not someone knows something; sometimes, we also care about how they come to know. This is why I still think about obliviousness as a form of ignorance, even if there are cases when the agent has the knowledge that *p* – even though they know, they do not know it in the way they should have. For instance, they should have figured out or for themselves that *p*, without needing to be told. But again, many cases of obliviousness will be ones where the agent does not know that *p*, and so will fall under the scope of the *should have known* phenomenon as traditionally understood.

<sup>5</sup> For arguments against positive epistemic norms, see (Littlejohn 2012; Nelson 2010; Wrenn 2007). For a helpful overview of the literature, see (Ichikawa 2022).

claim that there are rational routes to belief that agents ought to take, and that agents are oblivious for not taking, suggests that there are things that we ought (epistemically) to believe.

One prominent reason for doubting that we have positive epistemic duties comes from the fact that we are constantly inundated with information, and that there are infinitely many beliefs that are justified by our evidence at any given moment. As it would be impossible to form all of these beliefs, it has been thought untenable to accept that we could have positive epistemic duties.<sup>6</sup> Instead, any apparent positive epistemic duties are explained away as merely moral or pragmatic duties to believe, rather than being genuinely epistemic in nature.

Against this line of reasoning, and in the course of spelling out more precisely when someone counts as oblivious, I argue that we should embrace a middle ground between thinking that we either have zero or infinitely many positive epistemic duties.<sup>7</sup> We can grant that the skeptic about such duties is right in noting that there are infinitely many rational routes to belief that are available to an agent at any given moment. But note that while we do not take Zarak, for instance, to be oblivious for not noticing that the napkin he has been using at the cafe is blue, we do take him to be oblivious for not noticing that the cafe is filling with smoke, even though he could just as readily come to know both. This asymmetry suggests that there is some further normative condition that distinguishes cases of obliviousness, in which there is a duty to take a rational route to belief, from similar cases where there is no such duty. In spelling out this normative condition, I will suggest that, for limited epistemic agents such as ourselves, moral and pragmatic factors play a role in determining which of the infinitely many positive epistemic duties that we could in principle have, we actually have. I will also argue that allowing moral and pragmatic factors to play such a role does not undermine the claim that there is an *epistemic* duty involved, nor does it instrumentalize epistemic rationality.

Developing a more detailed analysis of obliviousness that points to the existence of positive epistemic duties and that posits a role for moral and pragmatic factors in specifying those duties will be one of the central aims of the paper. Another aim will be to highlight the social and political significance of obliviousness by connecting obliviousness to epistemologies of ignorance. While not all cases of obliviousness are socially or politically significant, I will

---

<sup>6</sup> (Nelson 2010)

<sup>7</sup> In so doing, I take myself to be building on and adding to recent defenses of the idea that we have positive epistemic duties. As I will discuss later, I particularly see myself as building on (Simion 2024), (Sanford C. Goldberg 2017; Sanford C Goldberg 2018; Sanford C. Goldberg 2022), and (Ichikawa 2022).

suggest that a study of obliviousness can give us insight into some of the important ways that our social environments shape us as epistemic agents.

Obliviousness is particularly interesting for this purpose because of the contrast it presents to many of the cases that are standardly considered in the epistemologies of ignorance literature. Much of the focus has been on cases where the, say, white ignorant agent either takes deliberate steps to avoid coming to know that  $p$  (e.g. refusing to believe, explaining away evidence, avoiding gathering evidence, or not seeking out hermeneutical resources), or where there is some structural impediment to their coming to know that  $p$  (e.g. there is a lack of information, information is hard to come by, or there are no adequate hermeneutical resources). Cases of obliviousness do not feature either of these things— the oblivious agent does not deliberately avoid coming to know that  $p$  at  $t$ , and they have all of the information, resources, and abilities needed to know that  $p$  at  $t$ . By directing our attention to cases where an agent not only should have known that  $p$  at  $t$ , but in which it *should have been obvious* to them that  $p$ , a study of obliviousness helps to illustrate just how far agential insensitivity to the facts of one's environment can go.<sup>8</sup> For instance, in discussing the ways that agents' white ignorance is supported by social mechanisms, Woomer notes that agents have "ready-made excuses to look away. They have ready-made narratives to explain away evidence."<sup>9</sup> However, a study of obliviousness shows that agents don't even need excuses to "look away" in order to remain ignorant— they can look directly at a situation and still not grasp what is happening. In this way, attending to obliviousness— and particularly to the diversity of cognitive failures that can be found within the class— can provide insights into the various ways in which social and political factors shape our (in)sensitivities as epistemic agents. In particular, I will suggest that social conditions help shape key cognitive dispositions in ways that often escape our deliberate control and awareness. In this way, these social conditions shape the beliefs that agents form and, conversely, the things that they "miss," even when these things should have been obvious to them.

The paper will proceed as follows. After providing an expanded set of core cases to work from (Section 2), in Section 3, I will precisify what I take to characterize obliviousness as a class of phenomena. In Section 4, I will turn to specifying the normative principle that serves to demarcate cases of obliviousness from informationally symmetric cases. I will suggest an aim-

---

<sup>8</sup> Woomer characterizes agential insensitivity as a form of epistemic insensitivity that is due to the agent's failure to use the epistemic tools that they have. See (Woomer 2019)

<sup>9</sup> (Woomer 2019: 15)

centered account that takes obliviousness to depend on the aims that the agent is or should be actively pursuing at *t*. In Section 5, I will argue that, despite the reference to aims in this account, obliviousness centrally involves an epistemic failure. This argument will involve situating obliviousness, and the normative principle that I propose for demarcating obliviousness, within a non-ideal approach to epistemology. I will sketch a picture on which practical (including moral) considerations help determine our positive epistemic duties from a non-ideal perspective, and show that this picture retains a distinctively epistemic concern for forming beliefs in a truth-conducive manner and does not instrumentalize epistemic rationality. Finally, in Section 6, I will show that obliviousness can function as oppressive ignorance in ways that help showcase how the (oppressive) social conditions that we inhabit significantly shape us as epistemic agents.

## 2. Core Cases

Before I present my account of obliviousness in the next section, here are three additional cases that I take to illustrate the class of phenomena that I am interested in.

**The Landlady:** Gina and Kristen recently moved to a new city. They rent a one-bedroom apartment, and every month they send their landlady a check with both of their names on it. They have different last names and do not resemble each other. One day Gina runs into their landlady, who asks how her sister is doing. “My sister?” “Yeah, I saw her the other day with some crutches. I hope she’s alright.” “Ohh, actually, Kristen is my wife. But she’s doing better! Thanks for asking.” The landlady turns red and apologizes for the confusion.

**I Don’t Work Here:** Kiara is a fashionably dressed Black woman who, after a day of shopping, is loaded with bags full of her day’s purchases. Before heading home, she steps into Bath & Body Works to test some of the scented hand lotions, only to be interrupted by another shopper, Lauren, who asks, “Excuse me, do you have any more of the Vanilla Sugar hand soap in stock?” Kiara tiredly responds, “I’m sorry, I don’t work here. Try asking the lady in the uniform over there.”

**I Can’t Operate:** Ana is doing her psych reading when she turns to her friend Bianca: “I’ve got a riddle for you. *A father and his son get into a car accident. The father is killed instantly, and the son is rushed to the hospital in critical condition and prepared for surgery. The surgeon comes in, sees him, and exclaims, ‘I can’t operate, that’s my son!’ How can this be?*” Bianca thinks for a few moments and replies, “It’s the boy’s stepfather.” “Ok. Any other possibilities?” “Oh! His dads could be gay!” “Anything else?” “Uh... It’s his dead father’s ghost...?” “Well, you’re in good company. Only 30% of participants in this study realized that the surgeon could be the boy’s mother, even after

being given multiple tries.”<sup>10</sup> Bianca buries her face in her hands and declares herself a horrible feminist.

### 3. Characterizing Obliviousness

As a starting point for characterizing obliviousness, I want to highlight both similarities and differences across the cases I have presented. With respect to the latter, note that different things seem to go wrong in rendering the person oblivious in different cases – for instance, whereas some cases involve a failure to attend to information in the agent’s vicinity, others involve a failure to reason from information that the agent already has. Thus, it is not the particular kind of cognitive failure that unifies the class of obliviousness; obliviousness can involve multiple kinds of cognitive failures.

Rather, on an intuitive level, the cases are unified in that the oblivious agent genuinely “misses” something that should have been obvious to them. For Zarak, it is that the coffee shop was no longer safe; for Valerie, that there are no people of color in her department; for Lauren, that Kiara does not work there; for the landlady, that Gina and Kristen are romantically involved; and for Bianca, that the surgeon could be the boy’s mother. It is not that these agents are intentionally avoiding or suppressing their knowledge in the situation. Indeed, each has good reason to want to know: Zarak does not want to endanger his life; Valerie wants to be a good social justice advocate; Lauren wants to find out about the soap, not offend Kiara; the landlady is just trying to be nice; and Bianca wants to solve the riddle and be a good feminist.

In the next sections, I will spell out the intuitive idea that the oblivious agent genuinely “misses” something that should have been obvious to them in more precise terms. The “should have been obvious” notion will be spelled out in terms of there being a rational route to belief that the agent fails to take but ought to have taken. To contrast cases of obliviousness with cases where an agent could readily have known that  $p$ , but where it does not seem that the fact that  $p$  should have been obvious in the same way, I will specify that the rational route is one that is *immediately* available to the agent. I will also precisify the sense in which the agent’s obliviousness is *genuine* at  $t$  through comparison to an account of willful ignorance.

#### 3.1. Immediate Availability

One of the markers of obliviousness is the sense that there is some fact that should have been obvious to the agent. This is part of what distinguishes obliviousness from other kinds of cases more frequently discussed in the epistemologies of ignorance literature. In order to

---

<sup>10</sup> (Belle et al. 2021)

provide a more precise characterization of obliviousness, I will begin by capturing when it is that the fact that  $p$  should be obvious to  $S$  at  $t$ . I will suggest that this is the case when there is a rational route to the belief that  $p$  that is immediately available to  $S$  at  $t$ .

Whereas a *route to belief* is any causal pathway that leads, or would lead, to the formation of a belief, a *rational route to belief* is one that respects traditional epistemic norms of doxastic justification.<sup>11</sup> Rational routes to belief do not involve, for instance, wishful thinking or fallacious reasoning. Rather, a rational route to belief involves, for instance, using reliable processes of evidence gathering, forming beliefs on the basis of one's evidence, drawing on one's background knowledge, and making logical inferences.<sup>12</sup>

Next, I want to introduce a notion of availability. I will use availability to refer to a practical sense of possibility. Following Woomer, I will understand practical possibilities here in terms of not requiring an inordinate amount of time, effort, or resources in order to be accessed or achieved.<sup>13</sup> A rational route to some belief that  $p$  is *available* to an agent when each of the components that make up the route (e.g. the evidence, the background knowledge, the inference skills) is available to the agent. That is, it must be practically possible for the agent to attain the relevant information, the logical inference skills, the hermeneutical resources, etc. required to be able to form a belief in the way specified by the rational route in question.<sup>14</sup> It would not be practically possible for the agent to attain the relevant information, for instance, if that information had been destroyed, if it would require technologies that do not currently exist in order to be accessed, or if it would require visiting an archive on the other side of the world and the agent lacks the funds to do so. In contrast, it would be practically possible for the agent to attain the relevant information if this required finding a book in their local library or doing a Google search on their internet-connected smartphone.

---

<sup>11</sup> Others who use the language of a "route to belief" and a "rational" or "epistemically legitimate" or "justification-conferring" route to belief include (Barkasi 2019; Marušić and White 2018; Siegel 2016). The fact that these routes are *rational* routes that respect traditional epistemic norms of doxastic justification will become important when it comes to the argument about positive epistemic duties, and in particular my claim that the aim-centered view I will offer does not instrumentalize epistemic rationality.

<sup>12</sup> Technically, because rational routes to belief are causal pathways, there will be many epistemically equivalent rational routes (e.g. different ways of causally implementing the inference from 'P' and 'P implies Q' to 'Q'). Consequently, when referring to a rational route I will actually mean any representative of a class of epistemically equivalent rational routes.

<sup>13</sup> Woomer is more specifically interested in the notion of available evidence, rather than of rational routes to belief. See (Woomer 2019)

<sup>14</sup> I will often use "information" instead of "evidence" throughout the paper in an effort to avoid orthogonal debates about the ontology of evidence (e.g. concerning whether evidence is object-like or proposition-like). See (Kelly 2008) for an overview.

Note that availability will be vague, context-sensitive, and agent-relative. For instance, whether it is practically possible for an agent to visit an archive on the other side of the world to access some relevant information – and therefore, whether that information is available to the agent – will depend on an agent’s social and economic situation. Likewise, whether such an effort would count as inordinate will depend on the context in question. If the agent is writing a book on the topic and the information in the archive is central to establishing their argument, it may not be inordinate; if the agent is trying to satisfy their child’s curiosity about the culinary habits of ancient Egyptians, it may be.

Note, too, that availability will come on a spectrum. Two rational routes to belief can both be available to an agent, in that neither would require an *inordinate* amount of time, effort, or resources to gather the relevant information, acquire the necessary skills, etc., and yet one available route may require a lot more effort, etc. than the other. Because my interest is in obliviousness – where there is the sense that the relevant facts should have been obvious to the agent – and because I am interested in the ways that obliviousness contrasts with other cases of ignorance with respect to the relative lack of barriers to knowledge, I want to introduce some further distinctions within this spectrum.

First, I will describe those rational routes to belief that are on the end of the spectrum requiring relatively little time, effort, and resources as being *readily available*.<sup>15</sup> Many of the cases discussed in the epistemologies of ignorance literature fall into this camp. That is, they are cases where the agent is ignorant of *p* despite there being a rational route to *p* that is readily available to them.<sup>16</sup> For instance, Woomer notes that many white Americans “[fail] to change their beliefs [about the existence of racial disparities in policing] to reflect readily available evidence,” citing publicly accessible reports by the ACLU and the Department of Justice as examples of sources of evidence that are readily available to white Americans regarding these disparities.<sup>17</sup> These agents remain ignorant despite having a rational route to belief about these disparities readily available to them, which Woomer diagnosis as “not a matter of skeptical white Americans’ being unable to access this information, but [of] being unwilling to do so.”<sup>18</sup> Similarly, Pohlhaus calls out cases of what she calls willful hermeneutical ignorance.<sup>19</sup> These are cases in which a

---

<sup>15</sup> As with availability, ready availability will also be vague, context-sensitive, and agent-relative.

<sup>16</sup> I will abbreviate “a rational route to the belief that *p*” to “a rational route to *p*.”

<sup>17</sup> (Woomer 2019: 3-4)

<sup>18</sup> (Woomer 2019: 10). This is not all that Woomer has to say about what underlies agential insensitivity, but it is an important piece of it.

<sup>19</sup> (Pohlhaus Jr 2012)



dominant agent lacks the hermeneutical resources that would allow them to properly make sense of a situation, but in which their lacking those resources is due to a failure on behalf of the agent. In particular, these are cases in which marginalized knowers in the dominant agent's community have developed the necessary hermeneutical resources – thereby rendering the relevant resources readily available to the agent – and yet the agent refuses to give uptake to those resources. They thereby remain willfully ignorant despite the relevant rational route to belief being readily available to them.

I want to draw two contrasts between these cases and the kinds of cases that I am interested in. First, both Woomer and Pohlhaus emphasize an unwillingness to know on behalf of the ignorant agent that does much to explain why they fail to take the rational route to belief that is readily available to them. In contrast, the cases that I am interested in lack this feature; this goes back to the idea that the oblivious agent genuinely “misses” that which they are ignorant of – I will say more about this contrast later.

The second contrast, which is relevant here, is about where the cases stand on the spectrum of availability – cases of obliviousness are even further towards the end of the spectrum than Woomer's and Pohlhaus's cases. In Woomer's case, for instance, it is certainly true that it would require little in the way of time, energy, or effort to look up the publicly accessible ACLU or DOJ reports. But the agent would still need to take the steps to search for the report, pull it up, and read it. In contrast, Zarak in *Smoky Cafe* just needs to attend to information that is already perceptually accessible to him in the situation, and likewise for Laura (*I Don't Work Here*). Meanwhile, Valerie, the landlady, and Bianca, do not even need to attend to new information in their environment; they already have the evidence, background knowledge, and inference skills needed to rationally form the relevant beliefs.

To capture this contrast, I will further distinguish *immediate availability* from *ready availability*. As with the notion of availability more generally, these notions are vague, agent-relative, and context-sensitive. For instance, information on a flashing neon sign may be immediately available to one person (because it is within their field of vision and visually accessible to them), and yet be merely readily available to the person right next to them who is also looking in the direction of the sign (because they would need to put on their glasses to be able to make out the words). Drawing on this distinction, I take a rational route to  $p$  to be immediately available to an agent at  $t$  when the agent satisfies the background conditions at  $t$  (i.e. they already have the relevant background knowledge, hermeneutical resources, inference

skills, etc.) and the relevant information is immediately available to the agent at  $t$ .<sup>20</sup> Note that this is more demanding than my characterization of (readily) available rational routes when it comes to the satisfaction of background conditions. That is, whereas satisfaction of the background conditions must be (readily) available to the agent in order for the rational route to be (readily) available to them, a rational route is only immediately available if the background conditions are already satisfied by the agent. This is to do justice to the sense that the fact that  $p$  should have been obvious to the agent at  $t$ . If the agent lacked the background knowledge and skills needed to come to know that  $p$  at  $t$ , then the fact that  $p$  could not have been obvious to them at  $t$ .

Finally, note that the claim that some rational route is immediately available to an agent is not the same as saying that it will be “easy” for them to gain the corresponding knowledge. The notion of immediacy points to the absence of barriers to knowledge. In a real sense, the agent is capable of coming to know that  $p$  at  $t$ : they possess the relevant skills and abilities, and the relevant information is either in their possession or perceptually accessible to them at  $t$ . But agents can still make mistakes and be prone to systematic errors. In this sense, it might not be “easy” for an agent come to know that  $p$ , despite being fully capable of doing so at  $t$ . Immediacy does not entail ease.

To summarize, the first condition that characterizes obliviousness is that there is a rational route to the belief that  $p$  that is immediately available to the agent at  $t$  that the agent does not take. This condition helps to more precisely capture the sense that the fact that  $p$  should have been obvious to the agent at  $t$ ,<sup>21</sup> while also helping to distinguish oblivious ignorance from other cases in the epistemologies of ignorance literature where knowledge is readily available to the agent, but where the agent would need to take additional steps to seek out information or other resources in order to be capable of coming to know that  $p$ .

### ***3.2. Non-deliberate***

---

<sup>20</sup> If the agent has the relevant evidence, I will also take this to count as the relevant information being immediately available to them. For instance, the landlady already has the relevant evidence concerning the nature of Gina and Kristen’s relationship; I would count this information as being immediately available to her. In the case of Zarak, the relevant information (that the cafe is filling with smoke) is immediately available to him because he has the perceptual capacity to see and smell the smoke. Depending on how one thinks about evidence, however, one might not want to say that Zarak has the evidence that the cafe is filling with smoke. As noted before, I am trying to remain neutral regarding debates about the ontology of evidence and what it takes to possess evidence.

<sup>21</sup> Though, as we shall see when discussing the normative condition, it is not sufficient on its own.

The first condition in my characterization of obliviousness is that the agent does not take the rational route to  $p$  that is immediately available to them at  $t$ . Zarak does not attend to the sights, sounds, and smells around him, Valerie does not draw the conclusion that the department is all-white, and so forth.

Importantly, however, none of these agents *deliberately* elects not to take the rational route to  $p$  that is immediately available to them at  $t$ ; rather, the agent “genuinely misses” the rational route to  $p$ . To clarify this second condition, I will compare obliviousness to standard cases of willful ignorance.

On Lynch’s (2016) account, the willfully ignorant agent suspects that  $p$  and deliberately performs or avoids performing certain actions in order to avoid knowing that  $p$ . Lynch uses the example of Albert Speer, a high-ranking Nazi official. In his memoirs, Speer writes that he avoided inspecting the concentration camps, because, based on a tip from a friend, he suspected that atrocities were being committed there and wished to remain ignorant. Supposing Speer writes honestly, this is a case in which someone avoids  $\phi$ -ing because they know it would lead to finding out whether  $p$  and they do not want to know that  $p$ . Alternately, consider someone who wants to avoid hearing what they suspect will be an uncomfortable truth, and so covers their ears to block out what is said. This person performs an action  $\phi$  in order to avoid knowing that  $p$ . The willfully ignorant agent, on Lynch’s account, deliberately takes or avoids taking certain actions in order to prevent themselves from taking a rational route to believing (indeed, knowing) that  $p$ .

In contrast, the oblivious agent neither deliberately takes nor avoids taking certain actions at  $t$  in order to prevent themselves from taking the rational route to  $p$  that is immediately available to them. They do not close their eyes or plug their ears or somehow distract themselves away from drawing simple inferences in order to avoid coming to know that  $p$  at  $t$ . In this sense, obliviousness is not deliberate: the oblivious agent does not deliberately choose, at  $t$ , not to take the rational route to  $p$  that is immediately available to them at  $t$ .<sup>22</sup>

These first two conditions are not yet sufficient for obliviousness, however, because there will be very many, perhaps infinitely many, rational routes to belief that are immediately available to an agent at any given time, and we do not take agents to be oblivious of all of these

---

<sup>22</sup> However, I do think that there can still be cases of willful obliviousness, if the agents takes actions before  $t$  in order to dispose themselves to be oblivious of facts like  $p$ .

things.<sup>23</sup> Accordingly, I will next suggest that for an agent to be oblivious of the fact that  $p$  at  $t$  it must also be the case that the agent *should* take the rational route in question. This is compatible with saying that there are many other rational routes that are immediately available to  $S$  at  $t$  that  $S$  need not, or perhaps even should not, take.<sup>24</sup>

### 3.3. Failure

The last key feature of obliviousness is that the oblivious agent *should have taken* the rational route to  $p$  that is immediately available to them at  $t$ . Not taking this route must constitute a normative failure, and not merely a descriptive failure.

To motivate this further, compare the following cases:

**Lost Tourist #1:** Alberto and his sister are visiting New York for the first time and are trying to figure out how to get to the Brooklyn Bridge. After unsuccessfully puzzling over a map, Alberto goes over to a man at the corner and asks him for directions. When Alberto comes back and tells his sister that he didn't know, she grumbles back, "Well, yeah. Did you not notice the big map in his hand and the camera around his neck?" "Oh. No. I didn't," he replies sheepishly.

**Lost Tourist #2:** Alberto and his sister are visiting New York for the first time and are trying to figure out how to get to the Brooklyn Bridge. After unsuccessfully puzzling over a map, Alberto goes over to a man at the corner and asks him for directions. When Alberto comes back and tells his sister that he didn't know, she replies, "Darn. Well, did you notice that watch he had on?" "Ooh, no, was it fancy?" "Very fancy!"

In both cases, there is a rational route that is immediately available to Alberto that he non-deliberately fails (descriptively) to take. In both cases, he overlooks immediately available information: in the first case, he doesn't notice the map and the camera; in the second, he doesn't notice the fancy watch. Despite these parallels, I want to say that Alberto is oblivious in the first case, but not the second.

Note that while it would be reasonable for Alberto to feel embarrassed and for his sister to be annoyed with him for not noticing the relevant item(s) in the first case, such reactions do not seem appropriate in the second case. This suggests that the relevant difference is a normative one. While Alberto might have enjoyed seeing the fancy watch and commenting on it with his sister, the fancy watch is not something that he *should have* attended to. His descriptive failure to notice the fancy watch is not a normative failure. In contrast, it seems that Alberto

---

<sup>23</sup> This point is central to Nelson's argument against positive epistemic duties. See (Nelson 2010).

<sup>24</sup> In this way, obliviousness is compatible with the idea that we should avoid cluttering our minds with junk beliefs, or beliefs about subjects that are of no interest to us, which many epistemologists have been sympathetic to. (Friedman 2018) develops this idea and provides an overview of the literature.

*should have realized* that the man was not a good person to approach for directions – he should have noticed the map and the camera and thereby concluded that the man was likely also a tourist.

Notice that the failure to take a rational route to belief can involve multiple, distinct kinds of failure that it has been controversial to take as being epistemic in nature: (1) a failure to form a belief, and (2) a failure to engage in epistemic activities that help constitute the route in question (e.g. failing to gather the relevant evidence).<sup>25</sup> Different cases will vary in these respects. For instance, while Alberto fails to gather relevant evidence in failing to notice the map or camera (which helps explain why he does not form the belief that the man is a fellow tourist), there isn't any relevant evidence that Bianca fails to gather in *I Can't Operate*.<sup>26</sup> Given that the failure to take a rational route to belief in cases of obliviousness includes both of these kinds of failures, then if this is indeed an epistemic failure, as I will argue in Section 4, then this suggests that we have epistemic duties to believe, as well as epistemic duties to engage in a host of epistemic activities that underlie responsible belief formation, such as gathering evidence.

Going back to the intuitive description of obliviousness that I have been trying to capture more precisely, notice that the normative condition on obliviousness works with the immediacy condition to capture the sense that the fact that  $p$  should have been obvious to the agent at  $t$ . That is, it is not enough for a rational route to be immediately available to an agent in order for it to be the case that the relevant fact should have been obvious to the agent, as the Lost Tourist cases illustrate. Rather, our sense of what should be obvious to someone in a situation is sensitive to normative concerns that go beyond the immediate availability of a rational route to belief.

### 3.4. *The Account*

Putting these pieces together, I characterize the class of obliviousness as follows –

**Obliviousness:** An agent  $S$  is oblivious of the fact that  $p$  at  $t$  just in case there is a rational route  $R$  to the belief that  $p$  such that:

1.  $R$  is immediately available to  $S$  at  $t$ ,
2.  $S$  should take  $R$  at  $t$ ,

---

<sup>25</sup> See (Flores and Woodard 2023) for defense of the view that there are epistemic norms on evidence-gathering and an overview of the literature opposing such a view. See (Singer and Aronowitz 2021) for defense of the view that we have epistemic reasons to perform ordinary actions like eating a sandwich, and not just to believe.

<sup>26</sup> In this way, obliviousness will include both “epistemically interesting” and “epistemically uninteresting” cases of the should have known phenomenon, as termed by Goldberg. This also means that some, but not all, cases of obliviousness can be understood as cases of normative defeat. See (Sanford C. Goldberg 2017; Sanford C Goldberg 2018; Sanford C. Goldberg 2022; Lackey 2020)

3.  $S$  does not take  $R$  at  $t$ , and
4.  $S$  does not deliberately avoid taking  $R$  at  $t$ .

This description spells out more precisely the class of cases I am interested in. In so doing, it helps distinguish obliviousness from cases that are popular in the epistemologies of ignorance literature. This, I will show, helps provide additional insight into the ways that we are shaped as epistemic agents by our social environments such as to perpetuate oppressive ignorance. In particular, by filtering out cases of ignorance that are explained by motivational factors or other traditional barriers to knowledge, the account draws attention to the role of cognitive processes that are under less deliberate control and awareness. If, as I will suggest, the functioning of these processes is shaped by our social environments, then this points to the epistemic as well as political significance of considering how these environments are structured.

This account also immediately raises questions about the scope of our epistemic responsibilities— in particular, when is it that an agent should take a rational route to belief that is immediately available to them at  $t$ , such that their not doing so constitutes a normative failure? I will take up this question in the next section, ultimately proposing an aim-centered account on which moral and pragmatic considerations play a central role in making these determinations. But while one might think that this renders the failure in question a moral or pragmatic failure, and not an epistemic failure, I will argue that the failure is both epistemic and practical. This indicates that we have positive epistemic duties, and outlines a role for practical considerations in determining which of the infinitely many positive epistemic duties that we could have, we actually have.

#### **4. When should one take a route to knowledge that is immediately available to one?**

The oblivious agent is one who fails, normatively speaking, to take a rational route to  $p$  that is immediately available to them, although they did not, at  $t$ , deliberately avoid taking it. But when is it that an agent *should* take a rational route that is immediately available to them?

The first possible answer is one that we have already discarded—

*Immediate:* Whenever a rational route  $R$  to  $p$  is immediately available to  $S$  at  $t$ ,  $S$  should take  $R$  at  $t$ .

*Immediate* is far too demanding of a principle, and does not capture judgments about obliviousness. There are many things that one can attend to at any given moment: the paint on the walls, the lint on the floor, the conversation at the next table, the smell of the pasta, the

placement of napkins, the trajectory one's hand traces through space while eating, the reflection of light off the fork, etc. Supposing there is nothing special about them, it does not seem as though you must attend to these things; they are of no interest to you. Indeed, others have plausibly supported the idea that you *should not* form beliefs about these things – to do so would be to clutter your mind with junk.

Another possible answer is suggested by Simion (2024). While Simion does not focus specifically on obliviousness, she is interested in a class of cases that significantly overlaps with obliviousness, in which agents exhibit what she terms resistance to evidence.<sup>27</sup> At first glance, Simion appears to endorse a principle like *Immediate*, proposing that "S has an epistemic duty to form a belief that *p* if there is sufficient and undefeated evidence for *S* supporting *p*."<sup>28</sup> However, Simion's knowledge-first account of evidence holds that something (specifically, a fact) will only count as evidence for an agent if the agent is in a position to know that fact. And further, Simion's account of when an agent is in a position to know something takes into account that agents are quantitatively, qualitatively, and environmentally limited. Thus, for Simion, the evidence that an agent has is more limited than the class of facts that are accessible to them:

[T]he *most easily available subset* of facts that I can take up delivers the set of evidence I have: in the case of visual perception, for instance, facts located right in front of me, in the centre of my visual field, which are the brightest, and clearest etc – in general facts that are most easily available to the cognitive processes of a creature like me." (Emphasis added) (Simion 2024: 209).

Thus, Simion's proposal is better translated into my framework roughly as follows:

*Easy*: If a rational route *R* to *p* is immediately available to *S* at *t* and the facts that help constitute *R* are the most easily available to the cognitive processes of a creature like *S*, then *S* should take *R* at *t*.

However, it does not seem that the ease of availability of information – corresponding to such factors as where something is located in one's visual field – will do the work needed in sorting out cases of obliviousness. Consider again Lost Tourist #1 and Lost Tourist #2. We can suppose that the camera and the watch are just as easily available to Alberto in each respective case. And

---

<sup>27</sup> Many cases of obliviousness will exhibit resistance to evidence, but not all – in particular, not the cases where the agent already has all of the relevant evidence. On the other hand, many but not all of the cases that Simion discusses seem to be cases of obliviousness. One difference is that Simion does not stipulate that the resistance to evidence must be non-deliberate.

<sup>28</sup> (Simion 2024: 205)

yet, again, it seems that whereas Alberto should notice the camera, it's fine if he does not notice to watch. This indicates that Simion's proposal is also too demanding.

#### 4.1. *Aim*

Reflecting on the difference between the watch and the camera for Alberto, what seems to be important is their relation to the aim that he is pursuing: the camera is relevant to identifying a good person to ask for directions, while the watch is not. Following this insight, I want to turn to an aim-centered account to capture when it is that an agent ought to take a rational route that is immediately available to them. As a first pass, consider the following proposal—

*Aim:*  $S$  should take a rational route  $R$  to  $p$  that is immediately available to  $S$  at  $t$  when knowledge of  $p$  is directly relevant to pursuing one of  $S$ 's aims.

The aim-centered account requires some finessing. On one hand, as it stands, *Aim* is overly demanding: it would unreasonably require agents to abandon aims that they are actively pursuing whenever new rational routes that are relevant to aims that they are not currently pursuing become available.<sup>29</sup> To illustrate, suppose Maggie is in a math class where, while reviewing concepts at the start of class, the professor says something that provides the key to finishing the proof that Maggie got stuck on last night. Since Maggie still has the aim of completing the proof, *Aim* would say that she now *ought to* take that route and work on completing the proof— but, she will miss important new material if she does so. In this context, it does not seem that Maggie fails if she does not stop paying attention in order to complete the proof; indeed, it seems that she *should not* do that. We are not required to switch the focus of our attention whenever new information pertinent to one of our aims becomes available. This suggests that we must also take into account which aims the agent is actively pursuing *at t*.

However, simply specifying that the rational route must be directly relevant to the aim that the agent is actively pursuing at  $t$  is also not sufficient.<sup>30</sup> This is because the aim that an

---

<sup>29</sup> I see the distinction here between an aim that one has and an aim that one is actively pursuing as a generalization of Friedman's distinction between "opening" a question (or putting it on one's "research agenda") versus "actively investigating" a question. In the former case, one has become curious about the answer to a question and determined that one would like to know the answer to the question (in the zetetic case), or one has identified an aim as one that one would like to achieve. But one is not actively pursuing the aim or actively investigating the question unless one is performing actions that are intended to move one towards answering the question or achieving the aim. See (Friedman Forthcoming: 2-3)

<sup>30</sup> I here specify "direct relevance" rather than mere relevance in an effort to bracket what might be termed "sandwich cases," after a case attributed to Sophie Horowitz. For instance, suppose Maggie is trying to complete the proof and she might do better if she raised her blood sugar by eating a sandwich. We might say that the knowledge that there is a sandwich shop next door is thereby relevant to the aim of



agent *should be* actively pursuing at  $t$  also matter.<sup>31</sup> Consider a teacher who is a mandated reporter trained to recognize signs of child abuse. If a child in their care presents with signs of abuse and the teacher fails to notice this, then the teacher is oblivious and culpably fails to fulfill their role as a mandated reporter. It doesn't matter if the teacher is callous and hasn't in fact taken an interest in the children's well-being; by virtue of their role, the teacher *should* have adopted the aim of looking out for the children's well-being, and so *should* notice and report. This indicates that it is not only the aims one has in fact adopted that matter for obliviousness, but also the aims that one should have adopted.<sup>32</sup> Likewise, recall Zarak. Although he is actively pursuing the aim of completing the proof at  $t$ , in Zarak's case, the aim of preserving his life normatively trumps the aim of completing the proof. And so he while he was not actively pursuing the life-preserving aim at  $t$ , he *should have been*.

#### 4.2. IAP

Factoring in these insights, I will suggest the following principle—

***Immediate Availability Principle (IAP):***  $S$  should take a rational route  $R$  to  $p$  that is immediately available to  $S$  at  $t$  when knowledge of  $p$  is directly relevant to pursuing an aim that  $S$  is or should be actively pursuing at  $t$ .

I take IAP to plausibly capture when it is that an agent should take a rational route that is immediately available to them. Because it is restricted to cases where one is or should be actively pursuing some aim, IAP does not require agents to form junk beliefs and generally allows them to focus their epistemic energies on subjects of interest. As such, I do not think that IAP is too demanding. Further, it is able to make sense of the cases we have been discussing.

To return to the core cases, each agent fails to take a rational route to knowledge that is directly relevant to an aim that they are or should be actively pursuing at  $t$ . Lauren wants to know whether the Vanilla Sugar hand soap is in stock; the landlady activated and quickly

---

completing the proof, but this does not seem like the kind of thing that is directly relevant here. When something will count as being directly relevant versus merely relevant will be vague and context-sensitive. For discussion of “sandwich reasons” (i.e. whether we can have epistemic reason to eat a sandwich), see (Falbo 2021; Singer and Aronowitz 2021; Thorstad 2022)

<sup>31</sup> Ichikawa puts forward, but does not settle between, two plausible candidate positive epistemic norms: “QPo+ If one considers the question whether  $P$ , and is in a position to know that  $P$ , one ought to believe that  $P$ .

□ QPo+ If one considers or should consider the question whether  $P$ , and is in a position to know that  $P$ , one ought to believe that  $P$ ” (Ichikawa 2022: 18).

This speaks in favor of the more normative variant, □ QPo+.

<sup>32</sup> (Friedman Forthcoming: 12) asks, “Might there be norms for opening questions that are sensitive to thoroughly non-epistemic considerations? ... Could we ever be morally required to be curious?” This indicates an affirmative answer.

settled inquiry into the nature of Gina and Kristen's relationship; Bianca is actively inquiring into how it is possible for the surgeon to be the boy's parent; there are several points at which Valerie was working to make the department more equitable, and in which recognition of the racial composition of the department would be directly relevant to fulfilling that aim; and Zarak should be actively pursuing his safety in the smoky cafe.

Despite these virtues, some may still think that IAP is either too lenient or too demanding.

Some may think that IAP is too lenient because it lets agents off the hook in certain cases where an inquiry has been paused.<sup>33</sup> However, I think it is a feature of the account that just as it allows agents to decide that certain subjects are not of interest to them, it allows them to decide that certain subjects are not of interest just now.

Alternately, one may worry about cases where we would expect information to be salient independently of an agent's aims. For instance, suppose a family member chides someone for not noticing that the person they were talking to for 10 minutes had rainbow-colored hair.<sup>34</sup> If the rainbow hair is irrelevant to any aim that the agent is or should be actively pursuing, they would not count as oblivious on my account. This may seem like the wrong verdict.

In response, I think we should distinguish between descriptive and normative expectations. I agree that in many cases we expect that most people would notice their interlocutor's rainbow-colored hair. This is a descriptive expectation. However, in line with IAP, my intuitions track whether there is a further normative expectation. If I take the chiding family member to be concerned about the agent's tendency to overlook unusual features of their environment because this awareness is important for survival, then I can see the agent as oblivious. But then noticing the hair would be connected to the agent's aims, and the case would be covered by the account. On the other hand, if noticing the hair really is irrelevant to the agent's aims or responsibilities, then it seems fine for the agent not to attend to it, and they don't strike me as oblivious for not doing so.

---

<sup>33</sup> For example, suppose that after spending hours searching for a top hat, Nalini decides to stop and resume her search another day. Just then, she walks right past a store with top hats in the display, but she doesn't notice. Some may think that Nalini should notice the top hat, right there in the display, that she has spent so much time looking for. However, it also seems fair for Nalini to be able to pause her search. IAP captures the latter sentiment.

<sup>34</sup> Thank you to Tez Clark and to an anonymous referee for pressing me on this kind of example.

The rainbow-colored hair case also gestures at the possibility of “good cases” of obliviousness.<sup>35</sup> Allowing that different kinds of normativity can deliver conflicting verdicts as to whether an agent should be pursuing a particular aim at *t*, an agent could be oblivious for not taking a rational route that they, say, pragmatically should take at *t*, but morally should not – e.g. noticing a weakness that they can exploit, or noticing someone’s deviance from social norms. These normative conflicts allow that an agent may be virtuously (if foolishly) oblivious for not following the demands of self-interest, or viciously oblivious for overlooking the demands of morality.

In the other direction, some may think that IAP is too demanding. Consider, for instance, a trivia case in which someone is asked to name two countries that start with a ‘C’ and their mind just comes to a blank, despite having the relevant knowledge – it may seem too harsh to call them oblivious.<sup>36</sup> While I think it is fine to judge this person to be oblivious, the sense of harshness here can be captured by the fact that the charge of obliviousness typically functions as a form of critique. Thus, it may be too harsh to *call* someone oblivious even if they are oblivious, especially if it is about a minor issue.<sup>37</sup>

## 5. What kind of failure does obliviousness involve?

A central feature of obliviousness is that it involves a normative failure to take a rational route to belief – but what kind of failure is this? I claim that obliviousness centrally involves both an epistemic and a practical failure. This has the controversial implication that we have some positive epistemic duties.

To some, it will seem obvious that obliviousness centrally involves an epistemic failure. After all, obliviousness consists in a failure to know what one should have known because one fails to take a rational route that one should have taken. Obliviousness thus centrally concerns an evaluation of our epistemic lives.

---

<sup>35</sup> Thank you to an anonymous referee for pushing me to address this issue.

<sup>36</sup> Thank you to an anonymous referee for this example.

<sup>37</sup> We can also allow that obliviousness comes in degrees, and say that this person is less oblivious than other agents we have been considering. For instance, how oblivious this person seems changes if they are a Canadian or from some other ‘C’ country. This suggests that a person’s social role will make a difference to the degree of obliviousness, perhaps in part by modulating how salient certain kinds of information should be to them. Further, since ‘immediate availability’ is a vague and context-sensitive threshold notion, some immediately available routes may be more readily available than others, which will plausibly also influence the degree of obliviousness.

For others, however, this will seem too quick. In particular, because IAP makes reference to agents' aims, some may worry that the relevant "should"s are practical in nature, and that obliviousness is thus ultimately not epistemically interesting. Goldberg distinguishes between the epistemic *contents* of an expectation, and the *grounds* of that expectation.<sup>38</sup> The worry here is thus that the seeming moral and practical grounds of IAP means that the failure to take the relevant rational route is not a genuinely epistemic failure, despite the epistemic contents of what IAP demands. Indeed, skepticism around the idea that we have positive epistemic duties has taken for granted that a positive epistemic duty, to be a genuine epistemic duty, must be "grounded wholly in one's epistemic circumstances and not also in the aims, desires, moral duties, etc. of the agent."<sup>39</sup> In making reference to an agent's aims, IAP would seem unable to deliver any kind of positive epistemic duty; instead, the thought goes, cases of obliviousness only show that we can have practical (including moral) duties to believe.

Recently, Simion (2024), Goldberg (2017, 2018, 2022), and Ichikawa (2022) have each provided arguments that push back on this general line of thought. Simion is interested in cases where agents fail to take up evidence that is easily available to them, and argues that this resistance is an epistemic failure. As part of defending this claim, she appeals to a widely agreed upon epistemic condition on moral responsibility in cases involving ignorance. This condition says that an agent cannot be morally blameworthy unless they are epistemically blameworthy for their ignorance.<sup>40</sup> Thus, one cannot explain away the failure to take up evidence and form the relevant belief by saying that it is a moral rather than an epistemic failure – given the epistemic condition, it cannot be a moral failure unless an epistemic failure is involved. While Simion's argument does not directly address cases where the failure is alleged to be pragmatic rather than moral, it bolsters the case for the existence of positive epistemic duties, and undermines the suggestion that an epistemic failure can be explained away by appealing to other (non-epistemic) kinds of duties. Goldberg also argues against attempts to explain away the sense of epistemic failure by appealing to other kinds of normativity. Goldberg notes that what is at issue is irresponsibility in belief formation, which appears to be an epistemic matter.<sup>41</sup>

---

<sup>38</sup> (Sanford C. Goldberg 2022)

<sup>39</sup> (Nelson 2010: 89). Ichikawa provides a helpful discussion of this trend. See (Ichikawa 2022; Littlejohn 2012; Wrenn 2007).

<sup>40</sup> (Simion 2024: 205)

<sup>41</sup> (Sanford C. Goldberg 2022: 92). Goldberg further appeals to our testimonial practices to strengthen this argument, while also noting that the view that attempts to "explain away" the seeming epistemic flaw is not able to do justice to all cases. See also (Sanford C. Goldberg 2017; Sanford C Goldberg 2018).

Ichikawa, in turn, rejects the constraint that positive epistemic duties must be grounded in purely epistemic considerations by observing that such a constraint is not universally accepted in other cases.<sup>42</sup> On one hand, he notes that this kind of constraint is controversial when it comes to negative epistemic duties – in particular, some take practical factors to help determine whether it is permissible to believe that  $p$ , as in cases of moral and pragmatic encroachment. He further notes that we do not take parallel constraints to apply in, for example, the moral domain – whether or not one has a moral duty to  $\phi$  depends in part on non-moral facts about one's circumstances. Given that this kind of constraint is not generally accepted, it is not clear why we should be bound by it when it comes to positive epistemic duties.

I would like to build on these ideas by taking a closer look at Nelson's (2010) argument against positive epistemic duties, and then using it to sketch a picture of positive epistemic duties within a non-ideal approach to epistemology.

Nelson's argument that we have no positive epistemic duties centers on what he terms the infinite justificational fecundity of evidence – the claim that any piece of evidence can epistemically justify infinitely many beliefs. The thesis that there are positive epistemic duties, Nelson argues, would entail that one should form every belief that one's evidence epistemically justifies. However, given the infinite justificational fecundity of evidence, this would entail that we should form infinitely many beliefs at any given time, which is not humanly possible.<sup>43</sup> Significantly, Nelson draws on a parallel to ethics in order to emphasize that the problem here is not merely one of demandingness, but of impossibility:

Suppose I agree with the utilitarian that I have a duty to maximize utility. Does it follow that I have a duty to produce a large positive net amount of utility? Of course not. I may be so situated in life that the best I can do is to produce some small net gain. If, on the other hand, I am a person of great power, resources, and good fortune, I may be able to produce a lot of net utility - in which case I have a duty to do so. Might I then have a duty to produce an infinite amount of utility? Even the most enthusiastic utilitarian should say 'No', for the reason that we are not obliged to do the truly impossible, and... it is impossible for any human being to produce an infinite amount of utility. By parity of reasoning, we might suppose that we cannot have a duty to add an infinite number of beliefs at any given moment, because (given similarly reasonable assumptions about our limited psychological capacities) this is not humanly possible, and we do not have a duty to do the impossible. Hence, we have no positive epistemic duties" (Nelson 2010: 97-98).

---

<sup>42</sup> (Ichikawa 2022)

<sup>43</sup> (Nelson 2010: 96-97).

In short, Nelson notes that, because we cannot have a duty to do the impossible, we cannot, in the moral domain, have a duty to produce infinite utility. He thereby concludes “by parity of reasoning” that we cannot, in the epistemic domain, have positive epistemic duties.

However, Nelson does not actually apply parity of reasoning here. After all, the utilitarian (and Nelson) does not conclude from the impossibility of producing infinite utility that there are no positive moral duties; and yet Nelson concludes from the impossibility of forming infinitely many beliefs that there are no positive epistemic duties. Rather, the implicit conclusion in the moral case is that our moral duties must be sensitive to practical factors and limitations – as Nelson indicates, someone with a lot of resources would have a duty to produce a large amount of net utility, whereas someone with fewer resources would have a duty to produce a smaller amount of net utility. Practical considerations help determine one’s moral duties, given the impossibility of living up to the ideal. Parity of reasoning then, would suggest *not* that there are no positive epistemic duties, but that our positive duties are similarly circumscribed and shaped by practical considerations.<sup>44</sup> Just as ethics takes our limitations into account when determining our positive duties (rather than simply taking these limitations to entail that we have no positive duties whatsoever), so, too, should epistemology.

Drawing out this idea, I suggest that we should think about obliviousness, IAP, and the claim that these indicate that we have positive epistemic duties, from the perspective of non-ideal epistemology. Following Carr, whereas *ideal epistemology* is concerned with how “perfectly rational, cognitively idealized, computationally unlimited” epistemic agents should comport themselves, *non-ideal epistemology* focuses on norms that real humans can satisfy.<sup>45</sup> Recall that IAP seeks to identify when an agent ought to take a rational route that is immediately available to them. For an idealized agent, the answer is clear – they should gather all of the available information, make all of the inferences, and increase their knowledge accordingly. The question is only interesting because we are not ideal agents. Our human limitations constrain what we are able to achieve epistemically. Our positive epistemic duties, from the perspective of non-ideal epistemology, are sensitive to these limitations.

There are different ways of developing this idea, in combination with the insights of IAP, into a more precise view. For instance, one might think that whether one ought to believe that  $p$  is dependent on a conjunction of epistemic and practical factors, such that there is no duty

---

<sup>44</sup> This aligns with Ichikawa’s suggestion that non-epistemic factors help trigger our positive epistemic duties.

<sup>45</sup> (Carr 2021: 1132)

to believe that  $p$  unless both kinds of factors are in place. Alternately, one could think that while we have an ideal duty to believe in accordance with our evidence, practical factors determine the non-ideal scope of our epistemic duties – that is, practical factors (in particular, our aims and responsibilities) determine which of those epistemic duties we are beholden to in practice, and exempt us from epistemic obligations that fall outside of that scope. Note that on either alternative, the normative expectations captured by IAP do not merely have epistemic contents, but also have epistemic grounds.<sup>46</sup> Without settling the details of the view here, the broad picture that I am suggesting is that, in light of our limitations, practical (including moral) factors help determine our positive epistemic duties, as understood from a non-ideal perspective. On this broad picture, practical factors can play both a diachronic role – helping determine how we should be directing our epistemic energies across time (e.g. what inquiries we should pursue, and how we prioritize and regulate those inquiries) – as well as a synchronic role – helping determine what, if anything, one epistemically ought to believe at  $t$  given one’s circumstances at  $t$ .<sup>47</sup>

On this picture, a violation of IAP constitutes both a practical and an epistemic failure. It is a practical failure insofar as it is a failure to properly further one’s aims and responsibilities. But it is also an *epistemic* failure because it is a failure to comport oneself as one ought to epistemically and to form beliefs as one epistemically ought to. It is not a coincidence that the epistemic failure is accompanied by a practical failure, since the practical factors help determine the scope of one’s positive epistemic obligations (or which positive epistemic obligations one is beholden to in practice), but this does not eliminate or diminish the epistemic failure.

Before moving on, I want to address a potential worry concerning whether this picture reduces epistemic rationality to instrumental rationality. The picture that I am proposing is not a reductionist view. We can still think that epistemic normativity is centrally aimed at believing truth and avoiding falsehood, and, as Kelly emphasizes in arguing against instrumentalism,

---

<sup>46</sup> I take this to be an advantage over Goldberg’s view. This view will still capture Goldberg’s observation that legitimate social expectations shape epistemic norms, without (as Goldberg does) taking epistemic norms to be *grounded in* these legitimate social expectations. On my view, epistemic norms can still be seen as being grounded in traditional epistemic values (e.g. believing truths, avoiding falsity), even while our non-ideal epistemic duties are sensitive to what is legitimately socially expected of us. See (Sanford C Goldberg 2018).

<sup>47</sup> On this picture, we can see obliviousness and the notion of junk beliefs as reflecting opposing pressures of practical factors – junk beliefs emphasize the ways in which practical factors push us away from certain subjects, whereas obliviousness emphasizes the ways in which practical factors push us toward other subjects; both reflect the idea that practical factors help determine where we should direct our epistemic energies.

that agents “have epistemic reasons to believe propositions even in cases in which it is clear that one’s believing those propositions holds no promise of advancing any goal which one actually possesses.”<sup>48</sup> The picture that I am proposing suggests that practical (including moral) factors help determine when (from a non-ideal perspective) acting on those epistemic reasons amounts to an epistemic duty; it does not deny that the epistemic reasons exist independently of practical factors. Further, recall that IAP is about when one should take *epistemically rational* routes to belief. The picture that I am proposing – and IAP as a non-ideal epistemic norm in particular – does not license one to believe other than on the basis of one’s evidence.<sup>49</sup> Rather, it preserves a distinctively epistemic concern for only forming beliefs in an epistemically responsible, truth-conducive manner.

This picture offers a middle ground between thinking that we either have zero or infinitely many positive epistemic duties. It takes seriously our human limitations and proposes that, in light of those limitations, practical factors help determine the practical, non-ideal scope of our epistemic duties.

## 6. Obliviousness as Oppressive Ignorance

Having explored dimensions of obliviousness that intersect with fairly traditional epistemic concerns about the nature and scope of epistemic normativity, in this last section I want to return to the relevance of obliviousness for social epistemology. In particular, I want to highlight how obliviousness can function as oppressive ignorance in ways that help showcase how we are shaped as epistemic agents by the (oppressive) social conditions that we inhabit in ways that often escape our deliberate control or awareness.

To make the case that obliviousness can function as oppressive ignorance, I want to build on Martín’s (2020) account of white ignorance. On Martín’s account, white ignorance is a kind of ignorance that is both systematically produced by and that plays a systematic role in sustaining, perpetuating, or giving rise to the social structures of white racial domination.<sup>50</sup> We can generalize from this notion of white ignorance to a notion of oppressive ignorance, where

---

<sup>48</sup> (Kelly 2003: 630).

<sup>49</sup> One might nonetheless worry that an appeal to practical factors will license agents to avoid believing things that they don’t care about, even if it does not license them to believe whatever they want to believe. This is true to some extent, but recall that IAP will not let one off the hook if the matter is one that the agent *should be* pursuing. The view thus does not permit objectionable cases of avoidance.

<sup>50</sup> See (Martín 2020).



oppressive ignorance refers to ignorance that is both a systematic product of and that plays a systematic role in sustaining diverse forms of oppression.

Cases like *Racial Blindness*, *I Don't Work Here*, and *The Landlady* – cases that are all too familiar to many people with marginalized identities – exhibit oppressive ignorance. In one direction, the obliviousness in these examples helps to sustain or give rise to oppressive social conditions. Valerie in *Racial Blindness*, for example, illustrates how agents, and particularly dominantly-situated agents, often manage to remain ignorant about oppressive social conditions even in the absence of traditional barriers to knowledge, such as a lack of information or deliberate avoidance. Valerie's obliviousness protects the injustice of the status quo from being recognized, and thereby potentially addressed, even when Valerie has all of the information she needs to know that the status quo is unjust. *I Don't Work Here* and *The Landlady*, in contrast, demonstrate how obliviousness can enact microaggressions that contribute to subjects' experiences of oppression.<sup>51</sup> Both by enacting microaggressions and preventing injustices from being recognized as such, these forms of obliviousness help to give rise to and sustain oppressive social conditions.

In the other direction, these cases – as well as cases like *I Can't Operate* – are also a product of oppressive social conditions. To make this case, I want to consider diverse kinds of failures that may underlie cases of obliviousness, and suggest that oppressive social conditions make it more likely that agents, and particularly dominantly-situated agents, will make these kinds of systematic errors.<sup>52</sup>

First, many cases of obliviousness seem to involve a salience failure: information (either in the environment or stored as part of the agent's background knowledge) that should be salient isn't, or information that shouldn't be salient is. *I Don't Work Here*, for instance, seems to involve both of these problems: information that more reliably indicates that Kiara doesn't work there is not salient to Lauren, whereas information about her racial categorization is very salient to Lauren. This demonstrates what Whitely (2023) has characterized as “relative attentional surplus on the wrong property” – an excess of attention to an individual's membership in a

---

<sup>51</sup> For discussion of the harms of microaggression and the role they can play in oppression, see (Freeman and Stewart 2018; Friedlaender 2018; McClure and Rini 2020). Further, note that my argument that obliviousness centrally involves an epistemic failure adds to Perez Gomez's case that some cases of microaggression are epistemically impermissible, without needing to appeal to moral encroachment. See (Perez Gomez 2021).

<sup>52</sup> To be clear, I do not intend to suggest that these are the only kinds of failures that might underlie obliviousness.

marginalized group that is often accompanied by a deficit of attention to other important properties.<sup>53</sup> Munton (Forthcoming) has argued that this misattribution of salience, as part of more general attentional practices, is a product of socialization that systematically prioritizes certain kinds of information in line with the values and norms of a social milieu.<sup>54</sup> This suggests that oppressive social structures will tend to produce salience structures that probabilify oppressive forms of obliviousness.

However, misattributed salience is not the only thing that goes wrong in *I Don't Work Here*; the case also exhibits a mishandling of inquiry – specifically, premature settling. We can suppose that Lauren gives relative attentional surplus to Kiara's being Black, and that she takes this fact to provide statistical evidence for thinking that Kiara is an employee. But even supposing that it is legitimate to take statistical facts about a group to provide some evidence about a particular member of the group, it does not seem that Lauren should take such statistical evidence to settle the question of whether Kiara works there, or at the very least not without first assessing whether there is other, more determinate evidence (e.g. whether Kiara is wearing the employee uniform) that might settle the question.<sup>55</sup> But Lauren does not do this, and prematurely concludes that Kiara is an employee based solely on her identification of Kiara's race. This indicates that not only is there a relative attentional surplus given to the wrong property, but that undue epistemic weight is also assigned to this property.<sup>56</sup> Premature settling is not the only kind of zetetic failure that we see in these cases. Valerie's failure in *Racial Blindness*, for instance, seems to result from a failure to raise a question in the first place. Which questions get taken up, as well as how inquiry into those questions is handled, seems to be influenced by social norms, schemas, and values. In this way, oppressive structures also help give rise to obliviousness by shaping whether or how inquiry is conducted.

Another kind of failure that seems to underlie cases of obliviousness is a failure to consider relevant possibilities. Drawing on results from three related studies, Philips and Cushman found that implicit, default representations of possibility are constrained by

---

<sup>53</sup> (Whiteley 2023)

<sup>54</sup> (Munton Forthcoming) Munton has also argued that salience structures that are inappropriately organized around demographic categories constitute a form of prejudice. See (Munton 2021)

<sup>55</sup> For views that question the assumption that beliefs formed using such statistical generalizations are epistemically unproblematic, see (Basu 2019; Bolinger 2020; Munton 2019).

<sup>56</sup> We can also understand this as a case of normative defeat: the evidence that Lauren fails to gather in her mishandling of the inquiry, normatively defeats the justification of her belief (if we allow that the belief would have been justified in the absence of the other evidence). See (Sanford C. Goldberg 2022; Lackey 2020).

normative considerations; in particular, participants tended to implicitly rule out possibilities that did not conform to prescriptive norms.<sup>57</sup> Similarly, studies in which participants were given either the same riddle as Bianca in *I Can't Operate*, or a gender-swapped version, show that participants' performance tracks conformity with gender schemas in ways that suggest that the schemas are regulating the possibilities that participants consider:

Despite being encouraged to think of more than one way to solve the riddle, only 30% (n = 45) of participants given the schema-inconsistent form of the riddle who had never heard the riddle before (n = 152) transcended stereotyped gender schemas to realize that the surgeon in the riddle could be the child's mother... Most of the participants who named the mother in one or both of their responses did so on the first response (n = 41, 91%). Participants who provided more than one response to the riddle were no more likely than participants who provided only a single response to the riddle to realize that the surgeon could be the child's mother.... [In contrast,] respondents were significantly more likely to realize that the surgeon could be the child's other parent when given the schema-consistent form of the riddle in which the surgeon is the father (n = 113, 78%) as opposed to the schema-inconsistent form of the riddle in which the surgeon is the mother (n = 45, 30%),  $\chi^2(1, n = 297) = 69.61, p < .001, \Phi C = .48$ . (Belle et al. 2021: 6-7)

These results suggest that our social norms and schemas— including oppressive norms and schemas— shape the possibilities that we consider when reasoning about the world. By implicitly excluding possibilities that one has the background knowledge to know can obtain, these norms and schemas make it more likely that we will be oblivious in schema-conforming ways. Notably, Belle et al.'s study also suggests that marginalized knowers are less likely to be susceptible to this kind of error: the only factor that they analyzed that predicted a greater likelihood of responding that the surgeon in the original version could be the boy's mother was identifying as a woman; factors like identifying as a feminist, exposure to female physicians, having a mother who worked outside the home, and political orientation were not significant predictors.

Lastly, hermeneutical failures can also underlie cases of obliviousness. The landlady, for instance, fails to adequately interpret her evidence about Gina and Kristen's relationship as evidence of a romantic relationship, despite having the resources that would enable her to do so. We can explain this by supposing that the landlady is not accustomed to viewing relationships through a hermeneutical lens that admits of those kinds of relationships. Such a failure, I suggest, is influenced by unjust hermeneutical practices. The problem here is not a gap

---

<sup>57</sup> (Phillips and Cushman 2017)

in the landlady's hermeneutical resources, as has been the focus in the hermeneutical injustice literature – again, the landlady has the relevant resources.<sup>58</sup> Rather, cases like *The Landlady* suggest that it is not enough to merely *possess* a concept; one must also cultivate the habit of routinely and appropriately *applying* the concept when making sense of one's experiences. Insofar as unjust hermeneutical practices continue to marginalize certain hermeneutical resources, it seems plausible that they can support systematic application failures even after agents have acquired the relevant resources.<sup>59</sup> In this way, (1) hermeneutical injustice might support obliviousness about certain matters, but also (2) thinking about obliviousness points to a distinctive manifestation of hermeneutical injustice that consists in these systematic application failures.

This discussion highlights that social conditions – including social norms, values, and schemas – shape us as epistemic agents in significant ways that often escape our deliberate control or awareness, at least in the moment. Social conditions help determine, e.g., what questions we take up, what information we attend to, what possibilities we consider, how we weigh our evidence, and when we take our questions to be settled. In so doing, social conditions help determine whether we take the rational routes that are immediately available to us, or whether we overlook these routes or take irrational routes instead. Oppressive social conditions in particular shape us in pernicious ways that incline us – especially along dimensions of privilege – toward oppression-sustaining obliviousness. This suggests that there are, thereby, epistemic as well as moral and political reasons to consider how our social environments are structured. By focusing on cases where more traditional barriers to knowledge are filtered out, a study of obliviousness draws attention to some of the more subtle factors that underlie patterns of ignorance, and the ways in which these are shaped by our social environments.

## 7. Conclusion

In this paper I have been interested in obliviousness – a familiar class of cases in which an agent “misses” something that should have been obvious to them. I have suggested that an

---

<sup>58</sup> See, for instance, (Fricker 2007; Pohlhaus Jr 2012).

<sup>59</sup> At the same time, there is also room for individual agency here that allows for willful obliviousness. In particular, if the landlady becomes aware of the issue, then she can take steps to practice applying the concept, such that it eventually becomes part of how she interprets the world. If she refuses to do so, then her obliviousness in other, similar cases can be considered willful.

agent is oblivious of the fact that  $p$  at  $t$  when they non-deliberately fail (normatively speaking) to take a rational route to  $p$  that is immediately available to them at  $t$ . Further, I have offered the Immediate Availability Principle as a norm that captures when agents ought to take a rational route that is immediately available to them: they should take those routes that are directly relevant to pursuing some aim that that they are or should be actively pursuing at  $t$ .

While I think that obliviousness is philosophically interesting for a number of reasons, I have here focused, first, on connections to debates about epistemic responsibility and normativity, and, second, implications for the epistemologies of ignorance literature.

The claim that obliviousness involves a failure to take a rational route to belief that one ought to have taken raises the question of what kind of failure obliviousness involves. I have defended the claim that obliviousness centrally involves an epistemic failure in addition to a practical failure, indicating that we have positive epistemic duties. I defended this claim by challenging a popular argument from the infinite justificational fecundity of evidence that appeals to our human limitations in order to conclude that we cannot have positive epistemic duties. In response to this argument, I sketched out a picture of non-ideal epistemic normativity that takes seriously our human limitations without thereby denying that we have any positive epistemic duties. Instead, I have suggested that in light of our limitations, practical (including moral) factors help determine the scope of our epistemic duties. This picture offers a middle ground between thinking that we either have zero or infinitely many positive epistemic duties, and brings practical considerations into play without instrumentalizing epistemic rationality.

Lastly, I have argued that obliviousness can be a form of oppressive ignorance, and that, by filtering out traditional barriers to knowledge, it showcases important ways in which social conditions shape us as epistemic agents. In particular, social conditions – including social norms, values, and schemas – help shape key cognitive dispositions, such as what questions we take up, what information we attend to, what possibilities we consider, how we weigh our evidence, and when we to take our questions to be settled. In so doing, social conditions shape our functioning as epistemic agents in ways that often escape our deliberate control or awareness in the moment, and that are thereby easily overlooked. A study of obliviousness helps to illuminate these factors, thereby highlighting epistemic as well as practical reasons to attend to these cognitive dispositions, as well as the social conditions that help shape them.

## Bibliography

- Barkasi, M. (2019), 'Are there epistemic conditions necessary for demonstrative thought?', *Synthese*, 1-28.
- Basu, R. (2019), 'The wrongs of racist beliefs', *Philosophical Studies*, 176/9, 2497-515.
- Belle, D., et al. (2021), "'I Can't Operate, that Boy Is my Son!': Gender Schemas and a Classic Riddle', *Sex Roles*, 1-11.
- Bolinger, R.J. (2020), 'The rational impermissibility of accepting (some) racial generalizations', *Synthese*, 197/6, 2415-31.
- Carr, J.R. (2021), 'Why Ideal Epistemology?', *Mind*, 131/524, 1131-62.
- Falbo, A. (2021), 'Inquiry and Confirmation', *Analysis*, 81/4, 622-31.
- Flores, C. and Woodard, E. (2023), 'Epistemic Norms on Evidence-Gathering', *Philosophical Studies*, 180, 2547-71.
- Freeman, L. and Stewart, H. (2018), 'Microaggressions in Clinical Medicine', *Kennedy Institute of Ethics Journal*, 28/4, 411-49.
- Fricker, M. (2007), *Epistemic Injustice: Power and the Ethics of Knowing*: Oxford University Press.
- Friedlaender, C. (2018), 'On Microaggressions: Cumulative Harm and Individual Responsibility', *Hypatia*, 33/1, 5-21.
- Friedman, J. (2018), 'Junk Beliefs and Interest-Driven Epistemology', *Philosophy and Phenomenological Research*, 97/3, 568-83.
- (Forthcoming), 'Zetetic Epistemology', in A. Flowerree and B. Reed (eds.), *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*: Routledge.
- Goldberg, S.C. (2017), 'Should have known', *Synthese*, 194/8, 2863-94.
- Goldberg, S.C. (2018), *To the best of our knowledge: Social expectations and epistemic normativity*. New York, NY: Oxford University Press.
- Goldberg, S.C. (2022), 'Normative Expectations in Epistemology', *Philosophical Topics*, 49/2, 83-104.
- Ichikawa, J.J. (2022), 'You ought to have known: positive epistemic norms in a knowledge-first framework', *Synthese*, 200/5, 400.
- Kelly, T. (2003), 'Epistemic Rationality as Instrumental Rationality: A Critique', *Philosophy and Phenomenological Research*, 66/3, 612-40.
- (2008), 'Evidence: Fundamental Concepts and the Phenomenal Conception', *Philosophy Compass*, 3/5, 933-55.
- Lackey, J. (2020), 'FALSE CONFESSIONS AND TESTIMONIAL INJUSTICE', *The Journal of Criminal Law and Criminology (1973-)*, 110/1, 43-68.
- Littlejohn, C. (2012), *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
- Martín, A. (2020), 'What is White Ignorance?', *The Philosophical Quarterly*, 71/4.
- Marušić, B. and White, S. (2018), 'How can beliefs wrong? – A Strawsonian epistemology', *philosophical topics*, 46/1, 97-114.
- McClure, E. and Rini, R. (2020), 'Microaggression: Conceptual and scientific issues', *Philosophy Compass*, 15/4, e12659.
- Munton, J. (2019), 'Beyond accuracy: Epistemic flaws with statistical generalizations', *Philosophical Issues*, 29/1, 228-40.
- (2021), 'Prejudice as the misattribution of salience☆', *Analytic Philosophy*, 64/1, 1-19.
- (Forthcoming), 'All the Things I Do Not Know and Refuse to Learn', *Oxford Studies in Epistemology*.
- Nelson, M.T. (2010), 'We have no positive epistemic duties', *Mind*, 119/473, 83-102.

- Perez Gomez, J. (2021), 'Moral encroachment and the epistemic impermissibility of (some) microaggressions', *Synthese*, 199/3-4, 9237-56.
- Phillips, J. and Cushman, F. (2017), 'Morality constrains the default representation of what is possible', *Proceedings of the National Academy of Sciences*, 114/18, 4649-54.
- Pohlhaus Jr, G. (2012), 'Relational knowing and epistemic injustice: Toward a theory of willful hermeneutical ignorance', *Hypatia*, 27/4, 715-35.
- Siegel, S. (2016), *The rationality of perception*: Oxford University Press.
- Simion, M. (2024), 'Resistance to evidence and the duty to believe', *Philosophy and Phenomenological Research*, 108/1, 203-16.
- Singer, D.J. and Aronowitz, S. (2021), 'What Epistemic Reasons Are For: Against the Belief-Sandwich Distinction', in B. Dunaway and D. Plunkett (eds.), *Meaning, Decision, and Norms: Themes From the Work of Allan Gibbard*: Maize Books.
- Thorstad, D. (2022), 'There are no epistemic norms of inquiry', *Synthese*, 200/5.
- Whiteley, E.K. (2023), 'A Woman First and a Philosopher Second: Relative Attentional Surplus on the Wrong Property [Open Access]', *Ethics*, 133/4, 497-528.
- Woomer, L. (2019), 'Agential Insensitivity and Socially Supported Ignorance', *Episteme*, 16/1, 73-91.
- Wrenn, C.B. (2007), 'Why There Are No Epistemic Duties', *Dialogue*, 46/1, 115-36.