

Obliviousness

DRAFT - Sep 2023 - Annette Martín

1. Introduction

Put yourself in my shoes for a moment. You're on the subway train back home. You've had a long day and the next days are going to be even longer. You exit at your station. You tell yourself it's time to stop thinking about work, so you put your earbuds in and turn on some music. A few minutes later, you stop— *where am I? Am I lost? How did I get lost on my way home?* It takes you a moment to orient yourself, and you realize that you missed a turn after exiting the station. You turn back, annoyed at yourself for obviously walking in the wrong direction for several minutes.

This paper is about obliviousness. Obliviousness is, I think, a familiar phenomenon, and one that many of us come across in our everyday lives. Perhaps you, like me, have taken an embarrassingly long time to realize you made a wrong turn in a familiar area; or perhaps you, like me, have deviously waited to see how long it would take an oblivious loved one to put two and two together. I suspect that many will have experienced the embarrassment, guilt, incredulity, amusement, irritation, or even anger that can come in the wake of obliviousness, whether one's own or others'. But despite its familiarity, an account of obliviousness is missing from the literature. My goal in this paper is to give an account of this familiar phenomenon and to show that it is worthy of sustained philosophical attention.

To highlight the major contours of the phenomenon, obliviousness, as I understand it, is a kind of ignorance— the oblivious agent fails to know or understand in a particular way. Apart from this, there are 3 main aspects of obliviousness, which can be loosely glossed as follows: the agent who is oblivious of p at t *should have known, could readily have known, and does not deliberately avoid coming to know that p at t* . Recent work in epistemology has increasingly broadened what is considered to be the traditional scope of epistemology, and obliviousness's distinctive features make it a fruitful tool for pushing forward questions in these spaces. For instance, thinking about obliviousness lends further support to the idea that the scope of epistemic evaluation extends beyond whether we believe in accordance with our evidence by pushing us to evaluate how we gather evidence, how we distribute our attention, what possibilities we consider, and how we go about managing and prioritizing our inquiries.

My aim in this paper is twofold. In the first part of the paper, I will elaborate on the key features of obliviousness that I have glossed above so as to provide an account of when it is that an agent is oblivious. Then, in the second part of the paper, I will take up some questions that are raised by the account. Here my aim will be to both deepen our understanding of obliviousness and to illustrate some of the ways in which engaging with this phenomenon can inform and enrich a number of debates in different areas of philosophy. In particular, I will highlight connections to discussions about the nature of epistemic normativity, inquiry, salience and norms of attention, non-ideal vs. ideal epistemology, claims that an agent *should have known*, moral evaluations of ignorance, and the insidiousness of oppressive social conditions.

2. Core Cases

In order to provide an account of obliviousness, I will begin by introducing some additional cases that I take to be paradigmatic of the phenomenon. Here are four core cases:

Smoky Cafe: Zarak is sitting in a coffee shop trying to finish the problem set due later that day. He has been puzzling over one 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.

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. Not long after, she's interrupted by a woman, 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."

The Landlady: Gina and Kristen recently moved to a new city together. 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. Gina is confused – "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, apologizes for the confusion, and hastily walks away.

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 are out driving and are involved in a terrible accident. The father is killed instantly, and the son is in critical condition. The son is rushed to the hospital and prepared for an operation that could save his life. The surgeon comes in, sees the patient, and exclaims, 'I can't operate, that boy is my son!' How can this be?*" Bianca thinks for a few moments and replies, "It's the boy's

stepfather.” “Can you think of any other possibilities?” “Oh! His dads could be gay!” “Anything else?” “Uh... It’s his dead father’s ghost...?” “Alright. 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.”¹ Bianca stares at Ana, buries her face in her hands, and declares herself a horrible feminist.

3. Characterizing Obliviousness

Looking at these core cases of obliviousness, there are some key features that stand out. In all of these cases, there is something that the oblivious agent could readily have known and seemingly should have known— for Zarak, it’s that the coffee shop was no longer safe; for Lauren, that Kiara did not work there; for the landlady, that Gina and Kristen were romantically involved; and for Bianca, that the surgeon could be the boy’s mother. Yet in each instance, the agent has no idea that this is the case. None of them are intentionally avoiding the relevant knowledge, and indeed they all have good reason to want to know: Zarak does not want to endanger his life; 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.

At the same time, there are also some important differences across these cases. In particular, different sorts of things seem to go wrong— whereas some of these cases involve a failure to perceive or attend to information in the agent’s vicinity, others involve a failure to reason from the evidence and background knowledge that the agent already has.

Attending to these differences, I will build up an account of obliviousness, making more precise the loose idea that obliviousness is a form of ignorance in which the agent both could readily have known and should have known that p , but also does not deliberately avoid coming to know that p at t .

3.1. Immediate Availability

One of the most striking features of obliviousness is the sense that knowledge was so readily within the oblivious agent’s reach. To spell this out, I want to begin with the notion of information that is available to an agent. Information is available to an agent if the agent could, practically speaking, obtain that information.² Availability will come on a spectrum: on one end of this spectrum, agents face considerable barriers to accessing that information— perhaps

¹ (Belle et al. 2021)

² I use “information” instead of “evidence” to avoid orthogonal debates about the ontology of evidence (particularly whether evidence consists in physical objects, sense data, propositions, etc.) See (Kelly 2008) for an overview. This notion of available information is similar to Woomer’s notion of available evidence. Woomer characterizes evidence as being available to an agent when it would not require an inordinate amount of effort to access the evidence (Woomer 2019: 9).

they would need to raise funds, travel to the other side of the world, and spend hours digging through an archive in order to obtain the relevant information; they could do this, but it would take quite a bit of time, energy, and resources. In other cases, there are very slight barriers to accessing the information – it may just require a simple Google search on the phone that is in their pocket. When thinking about obliviousness, we are on the far end of this spectrum of availability, where there are no significant external barriers to accessing the relevant information, and the agent has the abilities needed to access the information. Consider *Smoky Cafe*, for example: Zarak does not need to go anywhere, read anything, consult anyone, or even do a Google search to obtain the information indicating that there is a fire in the cafe; he is able to see and hear the scrambling customers, and to see and smell the smoke filling the cafe. We can say that the relevant information is not only available to him, it is *immediately* available to him.³

I do not want to stop at the notion of immediately available information to capture the sense in which knowledge is so readily within reach, however. This is because obliviousness does not just have to do with whether an agent gathers information that is immediately available to them. For instance, recall *I Can't Operate*. Bianca actually *has* all of the information she needs, so the problem isn't that she doesn't gather information that is immediately available to her. Rather, the problem is that she doesn't draw a simple conclusion (viz. that the surgeon could be the boy's mother) from the information she has – a conclusion that she is fully capable of drawing. To capture this, we need to consider all that is required for Bianca to arrive, in an epistemically happy way, at the realization that the surgeon could be the boy's mother, and not just whether she gathers (or indeed, possesses) all of the information she needs. In Bianca's case, she not only needs the premises provided by the riddle, but also the background knowledge that (indeed, most) children have mothers as well as fathers, and the logical skills needed to draw the conclusion that the surgeon could also, therefore, be the boy's mother.

To capture this, I want to generalize from the notion of (immediately) available information to the notion of *rational routes to belief* that are (immediately) available to an agent. 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

³ I take it that what counts as immediately available will be vague and context-sensitive.

justification; a rational route to belief leads to beliefs that are supported by an agent's evidence.⁴ In Bianca's case, the relevant rational route involves using her logical inference skills to draw the conclusion that the surgeon could be the boy's mother based on the information from the riddle and background knowledge about children and parents. This rational route is not just available to Bianca, in that she could feasibly gather the necessary information and develop the requisite skills, it is *immediately* available to her at the time she hears the riddle. She already has the information and skills; she just needs to draw the conclusion.

Drawing on this framework, then, the first key characteristic of obliviousness is that, for some proposition p , there is a rational route to the belief that p that is immediately available to the agent at time t . Now because rational routes to belief are causal pathways, there will be many epistemically equivalent rational routes – for instance, different ways of causally implementing the inference from 'P' and 'P implies Q' to 'Q'. Because of this, when referring to a rational route I will actually mean any representative of a class of epistemically equivalent rational routes.

For shorthand, we can use $R(p,S,t)$ to refer to a rational route to the belief that p that is immediately available to an agent S at time t . For such a route to exist, the agent must already have the relevant background knowledge and skills at t , and the relevant information must either already be in their possession or be immediately available to them at t .⁵

We can see how this will apply to the other core cases. Zarak (from *Smoky Cafe*) is able to see and hear the panicked customers, and to see and smell the smoke; further, he knows that smoke indicates that there is a fire, and that panicked people rushing to exit an establishment generally indicates some kind of emergency. The relevant information is thereby immediately available to Zarak, and he has the relevant background knowledge and cognitive abilities. Thus, a rational route to the belief that the cafe is not safe for him to stay in is immediately available to him. Likewise, Lauren (from *I Don't Work Here*) can see Kiara's shopping bag and fashionable clothing, and she has the background knowledge that these are indicators that Kiara doesn't work there. In *The Landlady*, the information is not only immediately available to the landlady,

⁴ Others who use the language of a "route to belief" and moreover of a "rational" or "epistemically legitimate" or "justification-conferring" route to belief include (Barkasi 2019; Marušić and White 2018; Siegel 2016).

⁵ Note that this means that obliviousness is going to be agent-relative, in the sense that not all agents will be oblivious in the same situation. For instance, there will be more information that is readily available to someone with good eyesight as compared to someone who isn't wearing their glasses, such that the former may be obliviousness in situations where the latter wouldn't be.

she actually has the relevant evidence: she knows that Gina and Kristen share a one-bedroom apartment, have different last names, have a joint checking account, and do not resemble each other. Moreover, she has the concept of a lesbian relationship and she has the background knowledge that these factors most likely indicate marriage, or at least a serious romantic relationship.

Thus, in all of these cases the agent has the relevant background knowledge and skills, and the relevant information is either in their possession or immediately available to them – in other words, a rational route to the belief that p is immediately available to each agent in these cases. It is in this sense that the oblivious agent could readily have known that p at t .

3.2. *Non-deliberate*

Again, the first part of my characterization of obliviousness is that for an agent to be oblivious of p at t there must be some rational route to the belief that p that is immediately available to the agent at t .

The next key feature is that the agent does not take this rational route that is immediately available to them: Zarak does not attend to the sights, sounds, and smells around him; Lauren does not attend to Kiara's clothing or shopping bags; the landlady does not draw on her background knowledge to conclude it is likely that Gina and Kristen are lovers; and Bianca does not realize that the surgeon could be the boy's mother.

Importantly, however, none of them *deliberately* elects not to take the rational route to p ⁶ that is immediately available to them at t (again, call this $R(p,S,t)$). To clarify this, it is helpful to compare obliviousness to standard cases of willful ignorance.⁷

On Lynch's (2016) account of willful ignorance, the willfully ignorant agent suspects that p and deliberately performs or avoids performing certain actions in order to avoid knowing that p . Lynch draws on the example of Albert Speer, a high-ranking Nazi official during World War II, as a paradigmatic case of willful ignorance. In his memoirs, Speer describes that, based on a tip from a friend, he avoided inspecting the concentration camps and seeing what went on there. He did this because he (correctly) suspected that atrocities were being committed in these camps and wanted to remain ignorant of this. Supposing that what Speer wrote in his memoirs was true, this is a case in which someone avoids ϕ -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

⁶ I will abbreviate "a rational route to the belief that p " to "a rational route to p ."

⁷ Note though that, as I will discuss later, there can be cases where obliviousness is actually willful.

wants to avoid hearing what they suspect will be an uncomfortable truth, and so covers their ears and starts singing loudly to block out what their partner has to say to them. This is a case in which someone performs an action ϕ because they do not want to know that p . In either kind of case, the willfully ignorant agent deliberately takes or avoids taking certain actions in order to prevent themselves from taking a rational route to believing (indeed, knowing) that p .

Now compare this with cases of obliviousness. The oblivious agent neither deliberately takes nor deliberately 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 otherwise take actions at t that would prevent them from accessing the information that is immediately available to them because they do not want to know that p . Nor do they somehow choose to distract themselves at t in order to prevent themselves from making the inference to p that they are in a position to make because they recognize that, if they do, they may come to know that p . It is in this sense that obliviousness is not deliberate: the oblivious agent does not deliberately choose, at t , not to take $R(p,S,t)$.⁸

To see why these first two conditions are not yet sufficient for obliviousness, note that there will be very many, perhaps infinitely many, rational routes to belief that are immediately available to an agent at any given time. As such, it is simply impossible for an agent to take all of them, whether deliberately or non-deliberately. Nelson, for instance, has appealed to the “infinite justificational fecundity of evidence” – that is, the claim that the evidence we possess at any given time can epistemically justify infinitely many beliefs – to argue that there cannot be positive epistemic duties, as that would entail that we should form infinitely many beliefs at any given time, which is not humanly possible.⁹ For similar reasons, many epistemologists have been sympathetic to the idea that we should avoid cluttering our minds with junk beliefs – that is, that we should avoid forming beliefs about subjects that are of no interest to us.¹⁰ It seems right that there will be very many rational routes to belief that are immediately available to an agent at any given time, and it cannot be the case that accepting obliviousness as a phenomenon means thinking that agents are required to take all of these routes.

Accordingly, I will next suggest that for an agent to be oblivious of p at t it must also be the case that the agent *should*, at t , take $R(p,S,t)$. This is compatible with saying that there are

⁸ As indicated above, this is still compatible with there being cases of willful obliviousness, but these will involve a diachronic element. I discuss this later on in the paper.

⁹ (Nelson 2010). See also (Nottelmann 2021) for further discussion.

¹⁰ (Friedman 2018) develops this idea and also provides an overview of the literature.

many other rational routes that are immediately available to S at t that S need not, or perhaps even should not, take.

3.3. Failure

The last key feature of obliviousness on my account is that the oblivious agent *should have taken* $R(p,S,t)$. It is not enough for S to be oblivious of p at t that there is some rational route to p that is immediately available to S at t and that S non-deliberately fails, in a purely descriptive sense, to take. Rather S 's not taking $R(p,S,t)$ must constitute a normative failure.

To see this, 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. More specifically, in both cases, there is some information he overlooks, despite its being immediately available to him: in the first case, he doesn't notice the map and the camera; in the second case, he doesn't notice the fancy watch. Yet despite these parallels, I want to say that Alberto is oblivious with respect to whether the person is a tourist, but not with respect to whether the man is wearing a fancy watch.

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 between these cases 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 does not constitute a normative failure.

In contrast, in the first case it seems that Alberto *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. In both cases there is a rational route that is immediately available to Alberto that he does not take, and in both cases he does so non-deliberately. Yet it is only in the first case – the case in which he should have realized that p – that Alberto is oblivious. Again, what this all suggests is that there is a normative dimension to obliviousness: the oblivious agent is one who fails, normatively speaking, in not taking a rational route that is immediately available to them at t , even though their failure to take that route is not deliberate.

3.4. *The Account*

At this point I have said enough to lay out the account of obliviousness –

Obliviousness: An agent S is oblivious of 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 should¹¹ take R at t ,
3. S does not take R at t , and
4. S does not deliberately choose not to take R at t .

This fulfills the first major aim of the paper: to provide an account of this familiar, but under-theorized, phenomenon.

But this account also raises some questions. When is it that an agent fails in not taking a rational route that is immediately available to them? What kind of failure does this involve – is the failure merely moral or prudential, or is it an epistemic failure? Why do agents fail to take these rational routes that are immediately available to them? What is it that goes wrong?

In the second part of the paper, I will take up these questions in turn. My goal will not be to provide complete answers or full discussions of all these topics; as we shall see, in some cases the answers raise further questions, and answering them all is beyond the scope of this paper. Rather, my more modest aim throughout the next few sections will be to provide a deeper understanding of obliviousness, while also highlighting some interesting avenues for further thinking and demonstrating that this is a rich phenomenon that can inform ongoing discussions in epistemology, ethics, and social and political philosophy.

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

¹¹ I'll discuss when an agents should take R and what kind of "should" is involved here in the second part of the paper. For now, I will say that I take this "should" to have an important epistemic dimension.

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 implicitly discarded on the basis of clutter avoidance. This possibility can be formulated as follows –

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

As already suggested, *Immediate* is far too demanding of a principle. 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. At any given moment, therefore, there are many rational routes to many different beliefs that are immediately available to one. Supposing there is nothing special about the paint on the walls, the lint on the floor, etc., it does not seem as though you must attend to these things; they are of no interest to you. And indeed, as previously discussed, 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.

4.1. *Aim*

Given that it seems permissible, or even obligatory, to forego rational routes into subject matters that are not of interest to one, I next want to consider an account that restricts when an agent should take a rational route to cases where it is relevant to pursuing one of the agent's aims. There are many aims, practical as well as theoretical, that require paying attention to certain things and pursuing certain lines of inquiry in order to be fulfilled. Someone who has the aim of driving safely to work, for instance, must attend to road signs, cars, cyclists, pedestrians, and other objects as part of successfully pursuing that aim. This suggests 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.

While *Aim* succeeds in filtering out junk, and is thus an improvement over *Immediate*, it also will not do. *Aim* is still too demanding while also, in another way, not being demanding enough.

First, there are cases in which *Aim* is too demanding because it would unreasonably require agents to abandon aims that they are already actively pursuing. Suppose Maggie is in a math class where the professor first reviews concepts before moving on to new material. During

the review, the professor says something that provides the key to finishing the proof that Maggie got stuck on last night; she now has everything she needs to complete the proof. Since Maggie still has the aim of completing the proof inquiry, *Aim* would say that she should now go through the mental exercise of completing the proof. But Maggie is currently absorbed in pursuing another aim – learning the class material – and if she does this, she will miss important new information. Under these circumstances, is Maggie required to take the rational route to completing the proof that is now immediately available to her? Does she fail in not doing so?

I do not think that Maggie fails if she does not stop paying attention to the lesson in order to complete the proof; indeed, it seems that she *should not* do that, and should continue to focus her epistemic energies on the lesson. Maggie's case suggests that what matters in cases of immediate availability is not just whether it is relevant to *some* aim that the agent has, but also whether it is directly relevant to the aim that the agent is *actively pursuing* in that moment. We are not required to switch the focus of our attention whenever new information pertinent to one of our aims becomes available.

On the other hand, there is also a sense in which *Aim* is not demanding enough, because it does not take into account objective normative considerations. For instance, consider a teacher who is a mandated reporter and has been trained to recognize signs of child abuse. If a child in their care clearly presents with signs of potential abuse, the teacher should activate inquiry into the child's safety and fulfill the mandated reporting requirements. It doesn't matter if the teacher is callous and doesn't take an interest in the well-being of the children in their care; the teacher *should* care and so *should* notice and report on the signs of potential abuse. It is not only the aims one has in fact adopted that matter, but also the aims that one should have adopted and should be actively pursuing right now.¹²

Relatedly, normative considerations can also have a trumping effect, such that the pursuit of some aims should take precedence over an aim that one is currently pursuing. For instance, consider Zarak in *Smoky Cafe* again. Although he is actively pursuing the aim of completing the proof inquiry, this is a case in which strong prudential considerations make it urgent for him to attend to the smoke. The aim of preserving his life takes precedence over the

¹² (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.

aim of completing the proof. Thus, while one is not always required to switch tasks when a rational route becomes immediately available to one, sometimes one should do so.

This suggests the following principle:

Normative Aim: *S* should take a rational route *R* to *p* that is immediately available to *S* at *t* when there is some aim, *A*, such that knowledge of *p* is directly relevant to pursuing *A* and *S* is or should be actively pursuing *A* at *t*.

While I take *Normative Aim* to be essentially correct, there is one more complication to consider: cases where there are competing rational routes to *p* that would each further the aim that *S* is or should be actively pursuing at *t*.

4.2. *Competing Rational Routes*

Consider the following case:

Diagnosis: Harry is a junior medical resident who has been learning about biases in medicine, including differences in “typical” presentations of heart attacks among men and women. Harry is presented with a new case for which he is supposed to provide a diagnosis. As he opens the file and starts to look over the chart, a reliable colleague tells him, “It’s a myocardial infarction.”

In *Diagnosis*, Harry is attempting to provide an accurate diagnosis for the patient whose chart he is given. At *t*, there are two rational routes to the belief that the patient is having a heart attack that are immediately available to him: (1) the route in which Harry assesses the information presented in the chart and, drawing on his medical knowledge, concludes that the patient had a heart attack; and (2) the route in which Harry listens to his colleague’s testimony and thereby comes to know that the patient had a heart attack.

There are three claims I want to make about this case. First, this seems like a case in which, at *t*, Harry can permissibly take either of these routes. If he listens to his colleague’s testimony rather than analyzing the chart, he has not done something wrong; likewise, he does nothing wrong if he tunes out what his colleague is saying because he is busy analyzing the chart. Some cases are not like this – an army scout, for instance, plausibly has a responsibility to tune out what her buddies are saying and to keep her eyes and ears peeled for signs of the opponent. In Harry’s case, however, there does not seem to be any clear protocol that would rule out the permissibility of attending to the testimony at *t*. The testimony could, after all, provide some helpful context.

Second, however, it seems that Harry’s work is not done if he listens to his colleague’s testimony at *t*. It is not enough for Harry, as a medical resident who is both furthering his medical training and treating patients, to merely know via testimony that this patient had a heart attack. Rather, Harry has a responsibility to also understand the diagnosis and the details

of the case, and his colleague's testimony does not provide this understanding. Thus, while at t_1 it seems fine for Harry to take either route to p , if he takes the testimony route, he should then take the other route (which is still immediately available to him) at t_2 and analyze the chart for himself.

Third, following the idea that Harry has a responsibility to understand the diagnosis and analyze the chart for himself at t_2 , I want to suggest that Harry is oblivious in the following continuation of the case:

Diagnosis (continued): At t_1 , Harry comes to know that the patient had a heart attack through his colleague's testimony. At t_2 , he analyzes the chart for himself. But upon doing so, Harry is confused — this doesn't look like a myocardial infarction to him. Seeing his confusion, the colleague tells him, "It's an atypical presentation."

When Harry goes to analyze the chart for himself, he is much like Bianca in *I Can't Operate*: he has the relevant information in the chart, and he has the background knowledge and skills needed to draw the right conclusion — after all, he has been learning about atypical presentations of myocardial infarctions — but he fails to apply this knowledge to the case at hand. While at t_1 Harry was not oblivious, at t_2 he is.

These observations give us insight into our immediate question — when it is that an agent is required to take a rational route that is immediately to them — but they also have broader implications for our understanding of obliviousness, as well as for how we should think about charges that an agent *should have known* that p .¹³

First, what the *Diagnosis* case suggests is that knowledge is not the only thing that is relevant for obliviousness — an agent can sometimes be oblivious even if they know that p . This is because sometimes the fulfillment of an agent's aims and responsibilities requires not only that they know that p , but also that they have knowledge of certain base facts and thereby come to *understand* that p . This is the case for Harry in *Diagnosis*. Coming to know that p via the colleague's testimony does not provide knowledge of the base facts and does not provide Harry with the requisite understanding. To gain this understanding he must look at the chart, gather the information in the relevant figures, and apply his background knowledge to that information to arrive at the correct diagnosis. Thus, even though Harry knows that the patient had a heart attack, he is still oblivious for failing to put the pieces together for himself at t_2 .

¹³ See (Sanford C. Goldberg 2015; Sanford C Goldberg 2018) for discussion of the *should have known* phenomenon.

Obliviousness is thus not merely a matter of whether one comes to know that p , but is also sensitive to how one comes to know that p . Recall that on my account, one is oblivious when one fails to take a rational route that one should have taken – it is centrally a matter of taking or failing to take certain rational routes. To return to the question of when an agent should take a route that is immediately available to them, this case illustrates that sometimes an agent is required to take $R(p,S,t)$, even though S already knows that p , because taking $R(p,S,t)$ is needed to achieve an understanding of p that S should have.

Further, this suggests that we should refine how we think about claims that an agent *should have known* that p . In particular, cases like *Diagnosis* and the army scout case indicate that we should take into account when it is that we expect someone not just to know that p , but also to take a particular kind of rational route to knowing that p . To use a distinction drawn by Lackey, in some cases we care not only about the *quantity* of an agent's support (whether that support is sufficient for them to know that p), but also the *quality* or kind of epistemic support they have for their belief.¹⁴ Further, sometimes the demand is not only for knowledge, but also for understanding – the agent not only *should have known*, but also *should have understood* that p .

To summarize what we have learned about cases where there are competing routes: sometimes the particular route one takes at t does not matter; just take one of them. Other times the particular choice of route does matter – one may have a responsibility to observe a protocol that specifies which rational routes one should take, in which case one should take the route thus specified. Further, some routes will give one understanding that one needs to properly fulfill one's aims and responsibilities. In such cases, one could know that p and still be oblivious if one closes the inquiry without taking such a route.

4.3. The Immediate Availability Principle

The question that we have been pursuing is when is it that an agent should take a rational route that is immediately available to them, such that they are oblivious for (non-deliberately) failing to do so.

¹⁴ (Lackey 2011)'s argues that norms of assertion are not merely sensitive to whether an agent knows that p , but also to what it is that grounds this knowledge. Lackey argues that in some cases, the quantity of epistemic support that an agent has may be sufficient for them to know that p , and yet the quality of that support (the kind of epistemic support they have) is such that they should not assert that p . This notion of the quality of one's support seems to align with thinking about the kind of rational route to belief or knowledge that an agent takes. Lackey does not, however, make a connection to understanding, though it does seem that this is what is missing in many of the cases she discusses.

It seems there will often be complex moral, social, prudential, and epistemic considerations that factor into when and how we should adopt and prioritize certain aims; likewise, various kinds of factors are likely to determine when one should prioritize a particular route among a set of competing rational routes. Given the scope of this paper, I will refer to the factors that would determine a unique route to be prioritized as a “protocol,” and leave full discussion of these matters to another time. For now, by adding what we have learned about competing rational routes to the insights previously captured in the formulation of *Normative Aims*, we have enough to suggest the following principle –

Immediate Availability Principle (IAP): If there is some aim, *A*, that *S* is or should be actively pursuing at *t*, and there is a set of rational routes that are immediately available to *S* at *t* such that each R_i in the set is a route to some p_j , where knowledge of p_j is directly relevant to pursuing *A*, then

- a. If there is only one such route in the set, or there is a protocol that specifies that *S* should prioritize a particular route (call this unique route R_1), then *S* should take R_1 at *t*;
- b. Or, if there are multiple such routes and no protocol that specifies a particular route that *S* should prioritize at *t*, then *S* should take any member of this set of routes at *t*.

I take the Immediate Availability Principle, as specified above, 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 agents 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 initial cases, in *I Don't Work Here*, *The Landlady*, and *I Can't Operate!*, each agent fails to take a rational route to knowledge that is directly relevant to an aim that they are actively pursuing at *t*. Lauren wants to know whether the Vanilla Sugar hand soap is in stock; the landlady activated and quickly settled the question of the nature of Gina and Kristen's relationship; and Bianca is actively inquiring into how it is possible for the surgeon to be the boy's parent. In *Smoky Cafe*, on the other hand, Zarak fails to take a rational route to knowledge that is directly relevant to an aim that he should be actively pursuing at *t*: he should attend to the smoke filling the room because this will further his aim of survival, which he should be actively pursuing in the midst of a threat to his safety.

IAP also makes sense of the difference between *Lost Tourist #1* and *Lost Tourist #2*. In both cases, Alberto is actively inquiring into how to get to the Brooklyn Bridge. Attending to the map and camera advances that inquiry, and so Alberto should take this rational route.

Attending to the fancy watch, in contrast, is irrelevant to this inquiry, and so Alberto need not take that rational route. This explains why Alberto is oblivious in the first case, but not the second.

Further, IAP can also capture Harry's obliviousness in the continuation of *Diagnosis*. When he goes to analyze the chart for himself, he is actively pursuing the aim of making and understanding the correct diagnosis. There is a rational route to that diagnosis that would give him knowledge of the base facts and allow him to understand the case, and this further knowledge is necessary for fulfilling his aim. He should thereby take this rational route that is immediately available to him. He fails to do so, and is therefore oblivious.

4.4. IAP: Too lenient?

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

One reason some may think that IAP is too lenient is because it lets agents off the hook in certain cases where an inquiry has been paused. For instance, suppose that Nalini has spent hours searching for a top hat that she can wear to a costume party in a few weeks. Eventually she gets sick of looking and decides to get some ice cream and resume her search another day. On the way to the ice cream shop she walks right past a store with top hats in the display – but, no longer being on the lookout, she doesn't notice it. On one hand, some may think that Nalini should pay attention to the top hat, *right there in the display*, that she has spent so much time looking for, and that she is oblivious for overlooking it. But on the other hand, it seems fair for Nalini to pause her search and not be constantly on the lookout for the top hat. I want to allow that, just as one may designate certain subjects as being junk subjects for one, so one may, in some cases, decide that subjects of interest are not of interest just now. Moreover, for those with interests in *a priori* subjects, there may always be knowledge that is immediately available to them if they just continue thinking about the logical implications of their current knowledge. But it does not seem right to say that they should always be pursuing these rational routes unless something more pressing comes along. Again, it seems like people are allowed to put a pause on pursuing certain aims, even if relevant knowledge is immediately available to them. So I take this more lenient aspect of IAP to be a good feature.

Alternately, one may worry about cases where there is information that we would expect to be salient to an agent even though it is not relevant to any aim that the agent is or should be actively pursuing. For instance, suppose that someone completely misses the fact that

the person they were talking to for 10 minutes had bright, rainbow-colored hair.¹⁵ We can imagine a family member who chides the agent for being oblivious. But, supposing that the rainbow hair is irrelevant to any aim that the agent is or should be actively pursuing, they would not count as being oblivious on my account. Some will worry this makes IAP too lenient.

In response, I think it is important to distinguish between descriptive and normative expectations of salience. I agree that in many – but not all – contexts, we would expect that most people would notice their interlocutor’s rainbow-colored hair.¹⁶ This is something that we descriptively expect to be salient to many people. However, in line with IAP, my own intuitions about the case track whether we can further interpret this expectation as being a normative expectation. For instance, if I consider that the chiding family member is concerned about the agent’s tendency to overlook important features of their environment, I can talk myself into thinking that perhaps the agent is oblivious. But this is only insofar as a plausible case can be made that attending to surprising features of one’s environment (like the rainbow-colored hair) is part of maintaining a basic awareness that is important for survival. But if that’s right, then the case would be covered by the account after all.

On the other hand, it also seems plausible that the color of one’s interlocutor’s hair is not very important for ensuring one’s survival, at least in our social context. If it really is irrelevant to the agent’s aims, then it seems fine for the agent not to attend to this feature, and they don’t after all strike me as oblivious for not doing so. Again, my intuitions as to whether the person is oblivious in this case directly track whether I interpret the expectation that the rainbow hair would be salient to the agent as being normative or not, just as IAP would indicate.

4.5. *Degrees of Obliviousness*

Heading now in the other direction, some may fear that IAP is too demanding. For instance, suppose that you ask someone to name two countries that start with a ‘C,’ and the only country that comes to mind for them is Colombia, although they know that Canada and China are also countries whose name starts with a ‘C’ – call this the *Country Trivia* case.¹⁷ On my account the person would be oblivious, but this may strike some as being too harsh.

While I think it is fine to say that this person is oblivious, I can agree that this person seems less oblivious than, say, Bianca in *I Can’t Operate*, or Laura in *I Don’t Work Here*. In light of

¹⁵ Thank you to [redacted] and an anonymous referee for pressing me on this kind of example.

¹⁶ For example, in contexts where many people have bright and multi-colored hair, I think there is far less of an expectation that the rainbow-colored hair would be salient.

¹⁷ Thank you to an anonymous referee for this example.

this, it is useful to introduce the idea that obliviousness can come in degrees – an agent can be more or less oblivious.

What might determine the degree of obliviousness in a particular case? I will not attempt to offer a complete story here, but I will point to some relevant considerations. For one, insofar as ‘immediate availability’ is a vague and context-sensitive threshold notion, some immediately available routes may be more readily available than others. Generally speaking, we can expect that the more readily available the route, the more oblivious the agent.

Further, thinking about a case like *I Don’t Work Here*, it seems notable that the information that Laura overlooks is precisely the kind of information that she should be seeking to help fulfill her aim. That is, if Laura is trying to identify an employee that she can direct her question to, she should be on the lookout for the kind of information that would reliably indicate whether the person is an employee (are they wearing the store uniform? Do they have a shopping cart?), but this is exactly the information that she overlooks. In contrast, suppose someone is trying to meet up with someone that they’ve never seen before at a crowded park, and the person informs them via text that they are standing right across the street from them. In this case it is less clear what exactly the agent should be looking for to identify the person waving at them across the street. Thus, it seems that an agent will be more oblivious to the extent that they are reasonably expected to seek out the particular kind of information that they overlook in the process of pursuing the relevant aim.¹⁸

To return to the *Country Trivia* case, I think that how oblivious the person seems will also depend on facts about their social role. For instance, it seems to make a difference if we imagine that the person in the case is Canadian (in which case they seem more oblivious) versus if they are Brazilian. We can explain this by saying that, by virtue of their status as a Canadian, Canada should be more salient to the Canadian than it is to the Brazilian. Likewise, by virtue of their respective roles, we might expect accessibility concerns (e.g. the lack of a microphone in a large lecture hall) to be more salient to someone who works in the office of Diversity, Equity, and Inclusion (DEI) than to a faculty member who is not connected to the DEI office.¹⁹ Thus,

¹⁸ Social practices can play a role in determining whether the agent is reasonably expected to seek out the kind of information in question. For instance, suppose your partner left a note on the fridge saying you are out of milk. If you and your partner have a practice of leaving such notes, that will make it the case that you are reasonably expected to notice the note, as compared to a case where you don’t have such a practice.

¹⁹ This is not to say that other faculty members do not have a responsibility to ensure accessibility in their classrooms, workshops, etc. On the contrary, I think faculty do have such a responsibility. It simply seems

depending in part on an agent's roles and responsibilities, there may be differing expectations concerning how salient certain information should be to an agent, and this too can shape an agent's degree of obliviousness.

There is certainly more to be said about what helps determine an agent's degree of obliviousness in any given case. For the purposes of this paper, I merely want to say that obliviousness can come in degrees. This allows us to say that cases like *Country Trivia* are ones where an agent is only mildly oblivious.

5. What kind of failure does obliviousness involve?

Moving on to a new question, I next want to consider the kind of failure that obliviousness involves. I claim that obliviousness centrally involves both an epistemic and a practical failure.

To some, the claim that obliviousness centrally involves an epistemic failure will seem obvious. After all, obliviousness consists in a failure to know what one should have known (or understood) because one fails to take a rational route that one should have taken. Obliviousness thus centrally concerns an evaluation of our epistemic lives – of what we (fail to) know, and of how we (fail to) arrive at that knowledge. Obliviousness is, therefore, an epistemic phenomenon that centrally involves an epistemic failure.

For some, 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 an interesting epistemic phenomenon. In short, how can obliviousness constitute an epistemic shortcoming if it is for practical reasons that an agent should have taken the rational route in question?

To respond to this worry, I want to identify IAP as a non-ideal epistemic norm. 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 epistemic norms that real humans can satisfy.²⁰ Recall that the question that IAP seeks to answer is when it is that an agent ought to take a rational route that is immediately available to them. If we imagine an idealized agent, the answer seems clear – such an agent should take all of the rational routes to knowledge that are immediately available

that the DEI officer has an even greater responsibility to ensure accessibility in these spaces. It seems even worse for the DEI officer to fail on this front.

²⁰ (Carr 2021: 1132)

to them. An epistemic agent who is not limited in the ways that humans are should gather all of the information available to them in their surroundings, make all of the inferences that follow from their existing body of knowledge, and increase their knowledge accordingly. The question is only an interesting one precisely because we are not ideal agents. Charges of obliviousness take into account our human limitations – as we have seen, we do not take someone to be oblivious every time they overlook something in their environment.

IAP thus answers the question of when it is that agents ought to take rational routes that are immediately available to them while taking into account the fact that we are non-ideal epistemic agents. Our limitations constrain what we are able to achieve epistemically, such that it becomes important to consider how we ought to assign priority to different inquiries. Practical considerations thus enter into IAP because these factors guide how we ought to make those rankings.

In this way, IAP presents a picture on which practical factors (our moral responsibilities, the projects we have taken on, etc.) help determine how we should direct our epistemic energies, given our various human limitations.²¹ Violating IAP thus constitutes both a practical and an epistemic failure. It is a practical failure insofar as it is a failure to direct one's epistemic energies in the way that will best 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. It is a failure to gather evidence that one should gather, or to interpret one's evidence correctly, or to draw the inferences one should draw (etc.), and thereby a failure to expand the quantity or quality of one's knowledge.²² Obliviousness thus centrally involves both an epistemic and a practical failure.

Note that this picture does not commit one to thinking that epistemic rationality just is instrumental rationality.²³ In particular, we can still think that epistemic normativity is centrally aimed at believing truth and avoiding falsehood, and that epistemic rationality guides agents toward this regardless of whether they have it as one of their goals to believe true propositions and not believe false ones. IAP is compatible with this traditional conception of epistemic

²¹ While more should be said to develop and defend this picture in detail, I do not have the space to do so in this paper. My much more modest aim to briefly sketch out a picture that helps reconcile the practical and epistemic in a way that preserves what I take to be the epistemic core.

²² I will address what I mean by the quality of one's knowledge in the next section

²³ Here I respond to those like (Kelly 2003) who argue against an instrumentalist conception of epistemic rationality on which epistemic rationality is reduced to instrumental rationality.

normativity because it centrally concerns when one should take *epistemically rational* routes to belief. IAP does not license one to believe other than on the basis of one's evidence.

With IAP, the way in which practical considerations make a difference to the beliefs one should form is in helping determine, for example, which of the countless inferences that one could make one should make, or precisely which information in the sea of readily available information one should attend to. IAP thus preserves a distinctively epistemic concern for only forming beliefs in an epistemically responsible, truth-conducive manner. As such, we can think of obliviousness as providing a complementary perspective to clutter considerations: one should generally direct one's epistemic energies towards higher-priority subjects and away from lower-priority subjects. We can still take those epistemic energies, however, to be centrally aimed at pursuing truth in a way that avoids falsehood within the relevant domains. From this perspective, it becomes clear that obliviousness is an interesting phenomenon that helps illuminate important considerations for a non-ideal epistemology.

6. What explains obliviousness?

The last question I will consider is how we might explain instances of obliviousness. In particular, one might wonder why it is that individuals fail to take rational routes that are both immediately available to them and that are relevant to aims that they are or should be actively pursuing. To motivate this question in the form of a worry, some might think that my account makes obliviousness too mysterious – after all, if it is so easy for the oblivious agent to take the relevant rational route, then why don't they just do it? What is it that goes wrong in these cases?

In response to the worry, I would like to distinguish between the claim that some rational route is immediately available to an agent and the idea that it will be "easy" for them to gain the corresponding knowledge.

The notion of immediacy, as I use it in this paper, points to the absence of certain kinds of barriers. In a real sense, the agent is capable of coming to know that p at t – they possess the relevant skills and abilities and have ready access to the information that they need to rationally conclude that p . But even when agents are epistemically capable, it is still possible for them to make mistakes. Further, agents can be prone to certain kinds of systematic errors even when they are capable of doing better. In this sense, it might not be "easy" for an agent come to know that p , although they are fully capable of coming to know that p at t . Immediacy does not necessarily amount to ease.

Indeed, my interest in obliviousness originally stems from an interest in white ignorance and other forms of socially supported, oppressive ignorance. In many of the cases of oppressive ignorance standardly discussed, there is some impediment (whether willfully placed there by the agent or not) that prevents the agent from being in a position to know that p at t . Sometimes this impediment takes the form of structural barriers to information and background skills.²⁴ For instance, the underfunding of woman's health research often means that the information that doctors would need to offer adequate care has not been produced and is therefore not available (let alone readily available) to them. Similarly, hermeneutical injustice means that agents sometimes lack the concepts that would allow them to adequately interpret the information that is available to them.²⁵ Other times agents deliberately erect barriers: they choose not to do a simple Google search or to seek out the appropriate hermeneutical resources because they would rather remain ignorant.²⁶

Yet sometimes, even when there are no such barriers – the agent has the relevant concepts, the relevant background skills and knowledge, and they do not even need to do a simple Google search because the information is immediately available to them in their perceivable environment – agents somehow still manage to remain ignorant, often (though not always) about precisely the sorts of things that obscure or reinforce oppressive social conditions. Cases like *I Don't Work Here*, *The Landlady*, and *I Can't Operate* are all too familiar to many people with marginalized identities. Given that these are familiar phenomena, what might explain why these kinds of systematic errors persist even when agents are capable of knowing better and, further, do not deliberately avoid coming to know that p at t ? What further insulates the oblivious agent even once these more standard barriers are removed?

I will not try to provide a comprehensive list of what might go wrong in cases of obliviousness. Instead, my more modest aim is to canvas some interesting possibilities that might help us make sense of the kinds of cases I have presented in this paper, and to consider how social factors might play a role in cultivating obliviousness.

Many cases of obliviousness seem to involve some kind of salience failure: either information (whether 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 a combination of both of these problems: information that

²⁴ See, for instance, (Alcoff 2007; Martín 2020; Mills 2007; Tuana 2006)

²⁵ (Fricker 2016)

²⁶ (Medina 2013; Pohlhaus Jr 2012; Woomer 2019)

more reliably indicates that Kiara doesn't work there is not salient to Lauren, whereas information about her racial categorization seems to be very salient to Lauren.²⁷

The *I Don't Work Here* case also suggests that obliviousness can occur as a result of mishandling inquiries – in this case, prematurely settling a question. We can suppose that Lauren identifies Kiara as Black and takes this to provide statistical evidence for thinking that Kiara is an employee. Even granting that such an inference is legitimate, it does not seem that Lauren should take this to settle the matter without at least first assessing whether there is other, more determinate evidence. Instead, she takes this to be sufficient for determining that Kiara is an employee, and so prematurely settles the question. While *I Don't Work Here* involves prematurely settling a question, other cases might involve a failure to raise a question in the first place. The DEI officer, for instance, should ask themselves whether there are barriers to accessing the workshop. If they raised the question, they would immediately realize that there should be a microphone in a large lecture hall, but they fail to do so.

Another kind of error that might underlie obliviousness concerns the possibilities that an agent considers. Consider *I Can't Operate*, for instance. It seems that Bianca implicitly codes the surgeon as male, thereby implicitly (and groundlessly) ruling out the possibility that the surgeon is a woman. Having made this mistake early on, she fails to consider the possibility that the surgeon be the boy's mother.

Lastly, obliviousness can also be connected to hermeneutical failures. For instance, while the landlady has the concept of a lesbian relationship, she may not be accustomed to viewing relationships through a hermeneutical lens that admits of those kinds of relationships. She thus fails to adequately interpret her evidence as evidence of a romantic relationship. Or consider the case of a narcissistic boss who, because he thinks so highly of himself, misinterprets his employees' exaggerated obsequiousness as a sign of respect – even though, were he to observe this behavior directed toward someone else, he would instantly recognize it as mockery. His inflated ego leads him to systematically misinterpret the evidence available to him, rendering him oblivious of his employees' open disdain.

Without, again, claiming to have provided a comprehensive list, I have highlighted some possibilities for explaining what goes wrong in cases of obliviousness: something may be awry with the agent's salience structure; agents may fail to consider certain possibilities,

²⁷ Munton argues that these kinds of salience failures can constitute prejudice. I also take this to support the Watzl's idea that we should develop an ethics of attention. See (Munton 2021; Watzl 2022)

adequately interpret their evidence, or raise certain questions; or agents may prematurely settle a question.

6.1. *Social Influences on Obliviousness*

As I observed, cases like *I Don't Work Here*, *The Landlady*, and *I Can't Operate* represent familiar patterns of experience. It does not appear to be just random individual error that leads to the kinds of mistakes that underlie these cases, but rather, given the pattern, there seems to be something more systematic at play. I will highlight some hypotheses as to how social factors could play a role in generating these systematic errors.

Turning first to cases involving misinterpretation, it seems reasonable to expect that hermeneutical injustice will play a role here. But while discussions of hermeneutical injustice have primarily focused on hermeneutical gaps (including cases in which agents intentionally avoid acquiring certain hermeneutical resources), this does not apply in cases of obliviousness – the oblivious agent possesses the relevant hermeneutical tools. As such, obliviousness suggests another way in which hermeneutical injustice can manifest. In particular, 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. These hermeneutical resources need to shape the lens that one actually uses to interpret the world, but cases like this one suggest that there can be a systematic failure to internalize these resources. 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. In this way, we see that (1) hermeneutical injustice might support obliviousness about certain matters, but also (2) thinking about obliviousness helps us to identify a distinctive form of hermeneutical injustice that consists in these systematic failures of application.

Next I want to consider ways in which social norms might impact both the relative salience of certain kinds of information, as well as the possibilities that one considers. For one, it seems that violations of or deviations from social norms are often salient to us. Recall, for instance, the person with the rainbow-colored hair. The fact that many people seem to expect that the rainbow-colored hair will be salient to an interlocutor seems connected to the fact that having such bright, multi-colored hair represents a deviation from mainstream social norms. But imagine instead that this interaction occurs in a context where lots of people have brightly colored and multi-colored hair. Within those spaces, I would not expect the rainbow-colored

hair to be very salient; instead, it seems more likely that “mainstream” hairstyles will be salient. This suggests that the norms of a given social context help shape which information is salient in that context, and in particular, that deviations from a norm tend to be relatively salient.

Norms also seem to make a difference to the possibilities that one considers, though in the opposite way – people seem more likely to consider possibilities that conform to a social norm, and to rule out possibilities that fail to conform to the norm. There is some attenuated empirical support for this idea. For instance, drawing on results from 3 related studies, Phillips and Cushman found that implicit, default representations of possibility are constrained by normative considerations; in particular, participants tended to implicitly rule out possibilities that did not conform to prescriptive (specifically, moral and rational) norms.²⁸ While social norms can be distinguished from moral norms, strong social norms could have similar force to moral norms. If that’s right, then social norms could help explain why oblivious individuals fail to consider certain possibilities that they would, upon further reflection, recognize as possibilities²⁹.

Lastly, Whiteley (2023) has highlighted ways in which members of marginalized groups often suffer from “relative attentional surplus on the wrong property” – in particular, an excess of attention to their membership in a marginalized group that is often, but not always, accompanied by a deficit of attention being paid to other important properties.³⁰ This phenomenon, illustrated in the *I Don’t Work Here* case, seems clearly connected to our social practices. Further, as previously discussed, in this case it seems to be not just the excess of attention given to Kiara’s race that plays a significant role, but also the fact that this is taken to be sufficient for inferring that she is an employee. Strong ideological associations between being a member of a certain group and occupying certain kinds of social roles could play a role in the prevalence of such inappropriate inferences.

In these ways (and surely also in others), social forces may plausibly play a role in systematically generating the kinds of errors that underlie obliviousness.

6.2. *Space for Individual Agency*

²⁸ (Phillips and Cushman 2017)

²⁹ Part of Phillips and Cushman’s study was to show that there is a difference between default representations of possibility and what people take to be possibilities upon further reflection.

³⁰ (Whiteley 2023)

While I have been highlighting ways in which social forces might create conditions that support systematic errors of the sort that underlie obliviousness, we should not think that this leaves no room for individual agency on my account.

For one, agents can be willfully oblivious. This might seem surprising, given that it is built into the account that the agent does not deliberately avoid taking the relevant route to p at t . However, this is compatible with the agent having, at some previous time, willfully cultivated habits of performing or not performing certain actions that would dispose them to be oblivious of certain things at t .³¹

For example, suppose that, over time, someone trains themselves to ignore homeless people sitting on the sidewalks of a busy city. At first they deliberately avert their eyes to avoid the feelings of guilt and discomfort that result from these interactions, but over time they develop a habit of focusing their attention straight ahead, such that eventually we can imagine that they genuinely do not even notice the homeless people they are passing. If they manage to cultivate this habit, then they are now in a position to be oblivious to the homeless people on the sidewalks – they have shaped themselves to be such that they do not have to deliberately avert their eyes at t in order to remain ignorant. This shows that first, obliviousness can be willful even if it is not deliberate in the particular moment;³² and second, that individuals can play a role in determining whether they are likely to be oblivious of certain matters.

Further, even if it is as a result of social forces, and not the result of a conscious decision, that an agent has developed certain habits of thought or attention that dispose them to be oblivious of certain matters, there is still quite a bit of space for individual agency. For instance, if the agent becomes aware that they have such habits, then they have a choice as to whether they work on correcting those habits. If they choose to do nothing then they can also reasonably be said to be willfully oblivious when the time comes. Or an agent can have good reason to believe that they, too, are likely to be susceptible to the relevant social forces, and so should keep an eye out for cases where they are likely to get things wrong. Even if individuals have limited control in the moment as to whether they are oblivious, they do have some control over managing the mental habits that might underlie obliviousness.

³¹ Individuals may also unconsciously or subconsciously develop self-protecting habits that dispose them to be oblivious of certain matters. These, too, can be cases of obliviousness, so long as the agent is not aware of the fact that knowledge that p is immediately available to them at t and so does not, at t , deliberately avoid coming to know that p . Thank you to [redacted] for pressing this point.

³² Note that this indicates that Lynch's account of willful ignorance needs some tweaking to allow for this diachronic, self-shaping, dispositional variety of willful ignorance.

This discussion highlights some of the ethical and political significance of obliviousness. On an individual level, individuals may be blameworthy both for their obliviousness and for any harms or wrongs that stem from it. This is most clear when the obliviousness is willful or the result of negligence on the part of the agent, as discussed above. But more generally, while ignorance is often taken to be an excuse for moral responsibility, this is only if the ignorance is not itself culpable, and some have thought that the ignorance is culpable only if it is unjustified.³³ Thus, insofar as obliviousness centrally involves an epistemic failure, it does not provide an automatic excuse for wrongdoing.

On a social and political level, I have suggested that obliviousness can systematically result from social forces, and particularly oppressive social forces, that encourage certain kinds of errors. As such, obliviousness can constitute an oppressive form of ignorance that both results from and helps reinforce oppressive social conditions.³⁴ A major idea in the epistemologies of ignorance literature is that there are social practices that cultivate ignorance. I take this discussion to provide further support for this idea, and to thereby provide further reason to think that epistemologists as well as political philosophers should be concerned with how social forces can shape our epistemic agency in both deep and subtle ways.

7. Conclusion

I had two aims in this paper. The first was to provide an account of obliviousness. I have fulfilled this aim by providing an account on which an agent is oblivious of p at t when they fail (normatively speaking) to take a rational route to p that is immediately available to them at t , although they do not do so deliberately at t .

My second aim was to explore some questions raised by this account that would give us a deeper understanding of the phenomena while also highlighting ways in which thinking about obliviousness can inform ongoing discussions in epistemology, ethics, and social and political philosophy. I considered when agents should take rational routes that are immediately available to them; what kind of failure obliviousness involves; and what goes wrong in cases of obliviousness such as to explain why agents fail to take these routes that are immediately

³³ See (Biebel 2018). For a more general overview of the debate around the epistemic condition for moral responsibility, see (Rudy-Hiller Fall 2018). For an argument that we can be morally appraised for failures to notice, consider, remember, and other “passive” mental states of the kind that obliviousness frequently involves, see (Smith 2005).

³⁴ This aligns with Martín’s account of white ignorance, which I think can be generalized to a broader notion of oppressive ignorance. See (Martín 2020).

available to them. In answer to these questions, I have offered the Immediate Availability Principle (IAP) to spell out when agents fail in not taking a rational route that is immediately available to them: in short, agents should take those routes that are directly relevant to pursuing some aim that that they are or should be actively pursuing at *t*. I have argued that obliviousness involves both an epistemic failure and a practical failure, and I have highlighted a number of errors that might underlie obliviousness – including failures of salience, failures of interpretation, failures to consider certain possibilities, and failures to raise or properly settle certain questions.

In answering these questions, I have touched on a number of different topics of interest across areas of philosophy, including the nature of epistemic normativity, inquiry, salience and norms of attention, non-ideal vs. ideal epistemology, claims that an agent *should have known*, moral evaluations of ignorance, and the insidiousness of oppressive social conditions. To highlight some upshots, within epistemology, obliviousness suggests that epistemic assessability is not limited to whether an agent believes in accordance with their evidence, and that factors specific to agents' aims, obligations, and circumstances can make a difference to how they should comport themselves epistemically, given real-world constraints. Moreover, sometimes we care about *how* agents come to know and whether they come to *understand* that *p*, and not merely whether they know that *p*. Within moral and political philosophy, thinking about the role of mental habits in giving rise to obliviousness highlights both that individuals can be straightforwardly culpable for their obliviousness, as well as the ways in which oppressive social practices can cultivate obliviousness in ways that reinforce and support oppression.

Given the rich connections that obliviousness has to multiple areas of philosophy, there is good reason to think more deeply about the phenomenon. In providing a clear account of obliviousness and outlining some avenues for further exploration, I have laid the groundwork for doing so.

Bibliography

- Alcoff, L.M. (2007), 'Epistemologies of ignorance: Three types', in S. Sullivan and N. Tuana (eds.), *Race and Epistemologies of Ignorance*, 39-58. Albany, NY: State University of New York Press.
- Barkasi, M. (2019), 'Are there epistemic conditions necessary for demonstrative thought?', *Synthese*, 1-28.
- Belle, D., et al. (2021), "'I Can't Operate, that Boy Is my Son!': Gender Schemas and a Classic Riddle', *Sex Roles*, 1-11.
- Biebel, N. (2018), 'Epistemic justification and the ignorance excuse', *Philosophical Studies*, 175/12, 3005-28.
- Carr, J.R. (2021), 'Why Ideal Epistemology?', *Mind*, 131/524, 1131-62.
- Fricker, M. (2016), 'Epistemic injustice and the Preservation of Ignorance', in R. Peels and M. Blaauw (eds.), *The Epistemic Dimensions of Ignorance*, 160-77. Cambridge, UK: Cambridge University Press.
- 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. (2015), '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.
- 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. (2011), '25111 Assertion and Isolated Second-Hand Knowledge', in J. Brown and H. Cappelen (eds.), *Assertion: New Philosophical Essays*, 0: Oxford 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.
- Medina, J. (2013), *The epistemology of resistance: Gender and racial oppression, epistemic injustice, and the social imagination*. New York, NY: Oxford University Press.
- Mills, C. (2007), 'White ignorance', in S. Sullivan and N. Tuana (eds.), *Race and Epistemologies of Ignorance*, 26-31. Albany, NY: State University of New York Press.
- Munton, J. (2021), 'Prejudice as the misattribution of salience☆', *Analytic Philosophy*, 64/1, 1-19.
- Nelson, M.T. (2010), 'We have no positive epistemic duties', *Mind*, 119/473, 83-102.
- Nottelmann, N. (2021), 'Against Normative Defeat', *Mind*.
- 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.
- Rudy-Hiller, F. (Fall 2018), 'The Epistemic Condition for Moral Responsibility', in E.N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*.
- Siegel, S. (2016), *The rationality of perception*: Oxford University Press.
- Smith, A. (2005), 'Responsibility for Attitudes: Activity and Passivity in Mental Life', *Ethics*, 115/2, 236-71.
- Tuana, N. (2006), 'The Speculum of Ignorance: The Women's Health Movement and Epistemologies of Ignorance', *Hypatia*, 21/3, 1-19.

- Watzl, S. (2022), 'The Ethics of Attention: an argument and a framework', in S.A. Archer (ed.), *Salience: A Philosophical Inquiry*: Routledge.
- 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.