

Rumination and Wronging: The Role of Attention in Epistemic Morality

CATHARINE SAINT-CROIX

University of Minnesota, Twin Cities

Last Update: August 9, 2022

The idea that our epistemic practices—from the ways we regard others' credibility, to our evidence-gathering practices, to the beliefs we harbor—are morally significant has been the core observation driving the growing literature on epistemic injustice, doxastic wronging, and moral encroachment. But, one element of our epistemic practice has been starkly absent from this discussion of epistemic morality: attention.

While the topic of attention has been a mainstay of empirical literature in psychology and neuroscience, it has only recently become commonplace in philosophy of mind and related fields. The main goal of this article is to demonstrate the importance of attention within epistemology, and especially within epistemic morality.

To see this, we will focus on doxastic wronging, which is the idea that our beliefs can wrong others in and of themselves. (*Basu and Schroeder, 2018*, p. 182) provide this example in their recent article “Doxastic Wronging”:

Party. Suppose that you have struggled with an alcohol problem for many years, but have been sober for eight months. Tonight you attend a departmental reception for a visiting colloquium speaker, and are proud of withstanding the temptation to have a drink. But when you get home, your spouse smells the wine that the colloquium speaker spilled on your sleeve while gesticulating to make a point, and you can see from her eyes that that she thinks you have fallen off of the wagon.

In their analysis, this is an instance of doxastic wronging: the belief itself—

your immediate thought that your partner has fallen off the wagon—constitutes a wrong against them.

In the end, I disagree. Cases like this do, indeed, involve some wronging, but that wronging is not the result of the belief—at least, not the belief alone. And, there are real harms that can be rooted in the mental life of the believer. But those harms are not the result of the belief—at least, not the belief alone. In the case of the alcoholic partner, the occurrence of the belief, *they fell off the wagon!* is not the wrong. Rather, I'll argue, insofar as a wronging has occurred in this case, it can be accounted for in terms of wrongful patterns of attention. The wrongly not-yet-accused is wronged when their partner chooses to ruminate upon the proposition, perhaps setting their sights on its confirmation or lingering on it throughout the evening. Given their relationship, this devotion of negative attention may be wrongful. But, should the believer bat this thought down, no wrong has been committed. This analysis not only captures the intuition that something is awry in this case, but also leaves room to account for particularities that might undermine the sense of there having been wrongdoing, such as the believer's past trauma, the partner's recent strange behavior, and so on.

I'll begin by reviewing the notions of epistemic morality, epistemic wronging, and doxastic wronging, with particular attention to the causal and normative structure of doxastic wronging. With that in mind, I turn to the Problem of Control (§2) and examine *Basu and Schroeder's* (2018) response to it (§2.1). Then, §3 presents a new dilemma that arises when we take responses in this vein seriously. In hopes of finding an answer to this dilemma, §4 takes a closer look at the wrongs of doxastic wronging. I suggest we focus on the role of attention in §5, where I offer an account of attention and show that it meets the criteria we need for epistemic wronging: it is a part of our epistemic practice that is agential and, therefore, morally evaluable. In light of this, §6 argues that we should understand cases like **Party** as instances of *attentional epistemic wronging*. This broader concept allows us to accommodate many of the intuitions surrounding doxastic wronging, while also avoiding the dilemma of §3.

1 Doxastic Wronging

Epistemic morality deals with epistemic wrongs, which are wrongs carried out in the course of our epistemic practice.¹ Here, I mean ‘epistemic practice’ in a fairly broad sense: epistemic practices are practices governed by epistemic normativity which, at minimum, is concerned with attaining true beliefs and avoiding false ones.² Clearly, doxastic wronging is a kind of epistemic wronging involving beliefs, but the concept is much more nuanced than this. *Basu and Schroeder (2018)* offer the following initial characterization:

Doxastic Wronging. A wrong committed by one person against another in virtue of what the first person believes about the second. (Paraphrased, p. 181)

They go on to highlight three key features of this idea:

1. Directedness: Doxastic wrongs wrong a particular person.
2. Causal Structure: An instance of doxastic wronging is caused by the belief itself, not its predecessors or consequences, and
3. Normative Structure: The wrongness of an instance of doxastic wronging is owing to the belief itself, not its predecessors or consequences.

¹Note that this is a normatively broad definition: there is nothing necessarily epistemic about the nature of the wrong itself because source of normativity involved is not specified in this definition. Many of the wrongs discussed in this paper are *morally* wrong, rather than epistemically wrong. That is, they are not violations of epistemic norms. Rather, they are violations of moral norms that involve epistemic practice. See *Saint-Croix (2022a)* for further discussion of these distinctions. A normatively narrow definition of epistemic wronging would constrain the source of normativity involved to epistemic normativity; whether such wrongings are possible is a matter of contention, especially within the moral encroachment literature (*Gardiner, 2018; Jorgensen, 2020*). This paper takes no stance on whether they are possible.

²Delineating the boundaries of epistemic normativity is a much-discussed challenge beyond the scope of this paper. I take it, at a minimum, epistemic normativity concerns attaining true beliefs and avoiding false ones, though even this is contested. Nevertheless, this proposed delineation suggests that epistemic normativity governs not only what to do with the evidence one has, but how to go about gathering evidence, which inquiries to open, when and how to close them, and so on. For discussion, see *Kornblith (1993)*.

The causal and normative aspects of Basu and Schroeder's description of doxastic wrongdoing are of particular interest because they stand in stark contrast with a more straightforward way that beliefs can wrong. To see this, consider the following example:³

Seaworthy. The SS Minnow is not seaworthy. The owner of the Minnow, however, has customers to serve and neither the time nor the money to repair the Minnow. So, he carries out a series of dubious inspections, chats with overly-supportive fellow shipowners, and so on. Soon after adopting these habits, he comes to believe that the ship is, in fact, seaworthy. He then lends the ship to Skipper Jonas Grumby, who charts it for a three-hour tour. The boat, its crew, and its passengers are never seen again.

In this example, the harm—the loss of the Minnow, its crew, and its passengers—is caused by the consequences of the belief: because he believed the ship to be seaworthy, he was willing to lend it to Grumby. Had the shipowner simply retired the boat that day, no harm would have been done. So, the causal structure of this wrongdoing requires action on the part of the believer. But this is not the case in instances of doxastic wrongdoing. Rather, it is the belief itself, not “what you do, either prior to, or subsequent to, forming a belief,” that brings about a doxastic wrongdoing (*Basu and Schroeder, 2018*, p. 181). As Basu clarifies elsewhere, “forming a belief” is meant in the completed sense, and should not be understood as including anything ‘upstream’ of the belief state itself, such as the epistemic practices, cognitive limitations, dispositions, motivations, or affective states that might contribute to the process of forming the belief (*Basu, 2019a*, §3). Similarly, turning downstream, it is not the actions you take, if any, as a result of having the belief that are relevant to the wrong of doxastic wrongdoing. It is simply the belief.

Before moving onto the normative structure of doxastic wrongdoing, it is worth highlighting a distinction raised in the previous discussion: doxastic

³This case is modified from *Clifford's (1877)* shipowner case.

harm versus doxastic wrong. By way of analogy, suppose my car rear-ends another vehicle while I'm driving. Whether I'm blameworthy for the impact depends on whether I'm responsible for the harm caused. And, whether I'm responsible depends on how the harm came to be. If I was paying attention to my phone or trying to change the radio station, my negligence renders me responsible and I am, therefore, blameworthy. I have wronged the other driver. On the other hand, if my brakes fail despite my diligence in auto maintenance and I do my best to avoid the other car, I am not responsible in a way that transmits moral blame. I have not wronged the other driver. (The insurance companies and state laws may disagree, but this has more to do with the exigencies of real-world adjudication than it does with moral analysis.) In this case, perhaps a poor mechanic or bad manufacturing is to blame. Or perhaps it's simply bad luck. In this case, there is harm done to the individual whose car was hit, but they have not been wronged. So, let's distinguish doxastic harms from doxastic wrongs:

Doxastic Harm. A harm done to an individual in virtue of what the first person believes about the second.

As defined, it is plausible that doxastic wrongs involve doxastic harms.⁴ But, the reverse need not be so: doxastic harms may not always engender doxastic wrongs. This is a conceptual distinction that comports with standard ethical reasoning—not all harms are wrongs, and when a harm has occurred, you have wronged someone only insofar as you are responsible for the harm done to them, supposing that harm isn't somehow warranted. For example, a referee's rejection of my paper may harm me, but that does not make the

⁴This is the standard assumption in the literature: *Basu (2019c)* explicitly argues that the wronging does involve harm and *Schroeder's (2018, §4)* discussion of diagnosing wrongs suggests that he, too, assumes there must be some harm involved. One might nevertheless think that doxastic wrongs are a species of harmless wrong. For example, on a contractualist picture like the one offered in *Scanlon's (1998) What We Owe to Each Other*, doxastic wrongs would be construed as violations of expectations that, arguably, do not constitute harms. To my knowledge, however, no such account is defended in the literature on epistemic morality. See *Begby (2018)* for some additional discussion of this question in the context of doxastic morality, and *Healey (2017)* for a general discussion of the relationship between interests, harms, and wrongs.

rejection wrongful. Moreover, if the referee slips into a coma ten minutes after accepting the assignment to referee my paper, I might be harmed by their failure to complete the job, but they've clearly not wronged me in this failure. This latter example illustrates the *control condition* on responsibility, which we will return to in §2.

For the moment, we return to doxastic wrongdoing to consider the other key feature: its normative structure. Basu and Schroeder emphasize that “a belief that is a doxastic wrongdoing does not wrong merely in virtue of its consequences; the wrongdoing lies in the belief, rather than in, or at least, over and above, its effects,” (2018, p. 181). It is neither the fact that the subject of your belief would feel wounded if they knew about your belief, nor your reaction of disgust or pleasure, nor the role the belief plays in your mental life. These are all downstream consequences of the belief. So, in **Party**, the thought is that the belief *my partner fell off the wagon* is harmful on its own, independent of any feeling of woundedness the partner might have, were they to find out about the belief. While Basu and Schroeder state that “the feeling of being wounded is arguably a sign of a directed wrong” (2018, p. 182) and often use the counterfactual sense of woundedness as evidence that a wrongdoing has occurred, this negative consequence is nevertheless unnecessary.⁵ This is markedly distinct from **Seaworthy**. In this case, the belief itself is benign; all relevant harms arise from consequences of the actions taken on the basis of the belief.

This illustrates a principle implied in *Basu and Schroeder's* (2018) and *Basu's* (2019a) conception of doxastic wrongdoing:

Evaluative Resilience. With respect to doxastic wrongdoing, the wrongs and harms, if any, of a particular belief will be resilient in the face of changes to (a) the believer's upstream and downstream epistemic practices, dispositions, capacities, or motivations, and (b) the normatively salient consequences, epistemic and moral, of the belief.

⁵And, it's insufficient: a hypothetical feeling of woundedness may be misplaced.

The list in (a) is drawn from *Basu's* (2019a, p. 2507) illustration of the antecedent influences on belief formation from which doxastic wrongings are distinct, and which do not contribute to the wrongs of doxastic wrongdoing. It is important to determine whether this evaluative resilience holds when considering whether a particular belief is an instance of doxastic wrongdoing. This is because, if those upstream or downstream influences are responsible for the apparent wrongfulness of the belief, and we count the belief as an additional wrong beyond those influences, then we seem to be merely re-describing the same wrongdoing instead of identifying two separate wrongs. If the wrongdoing is only identifiable in consideration of these influences, or if their presence or absence significantly alters our intuitions with respect to whether some wrongdoing has taken place, this suggests that those influences are part of the causal or normative structure of the wrongdoing, meaning that those wrongs are not instances of doxastic wrongdoing. The harm, if there is any, and the wrong must be located in the believing of the proposition alone, exclusive of wrongs that can be causally or normatively grounded in these antecedent and subsequent influences or actions.

Before moving on, two clarifications are in order. First, note that evaluative resilience does not imply that all features of the broader context of a belief are irrelevant to determining whether it is an instance of doxastic wrongdoing. For example, the relationships involved, the moral facts, and so on might all play a role in this evaluation.⁶ Additionally, it is worth highlighting that evaluative resilience does not imply that the enumerated features do not give rise to additional wrongs. Those wrongs may be part

⁶Generally, whether an act is a wrongdoing depends also on whether it is justified. In this case, there are two sources of normativity in question: epistemic and moral, which seem to have distinct standards of justification. However, since epistemic practice includes justificatory practice, I will follow *Basu* (2019a) in assuming that the justification of the belief is part of the epistemic practice ruled out by evaluative resilience. Doxastic wrongs, *Basu* writes, can be committed by beliefs that “reflect reality and seem to be rationally justified,” (p. 3). Thus, the justification in question is moral justification. Note, however, that this discussion is intimately tied to moral encroachment and the so-called “coordination problem”, which concerns whether epistemic and moral normativity can come into conflict. I will not engage with these questions here; see *Basu and Schroeder* (2018, S3.1), *Enoch and Spectre* (forthcoming).

of the broader category of epistemic wrongs, but they are not instances of doxastic wrongdoing. And, those additional wrongs may co-occur with any instance of doxastic wrongdoing. The condition of evaluative resilience merely states that changes in these influences should not change whether there has been an instance of doxastic wrongdoing.

In the next two sections I will argue that the standard intuitions around paradigmatic cases of doxastic wrongdoing, such as **Party**, are not compatible with evaluative resilience. To account for these intuitions, we are compelled to consider the believer's actions surrounding her belief. But, as we will see, this incompatibility points the way toward more solid grounding for the wrongdoing in question.

2 The Problem of Control

While there are other objections to doxastic wrongdoing,⁷ our focus will be the looming *problem of control*. The core of this objection is the claim that there is no bridge between the doxastic and the normative because an essential component of that bridge—control over one's beliefs—is missing. As a result, doxastic wrongdoing must be impossible.

Let's begin with the normative side of the objection. Here, the thought is that, in order for moral notions like wrongdoing, blameworthiness, and responsibility to apply to an agent's action, that agent must be able to exercise meaningful control over the action in question. For example, when the doctor taps your patellar tendon with a reflex hammer, you don't have control over your body's response: your leg kicks briefly, then goes back to rest. So, if the doctor stands directly in front of you when administering this test, you are not morally responsible for the resulting kick, regardless of any harm it might cause. This is the control condition.

On the epistemic side, we set aside the question of whether beliefs can harm. Instead, the trouble arises with epistemic agency: It is a commonplace

⁷See *Enoch and Spectre* (forthcoming); *Begby* (2018)

in epistemology to argue that we are at the mercy of our evidence.⁸ When we encounter new evidence, we involuntarily update on it. The evidence ripples through our priors and we find ourselves with new posteriors. *BonJour* (1985) illustrates the point this way:

Consider then the following example of (putative) observational knowledge: As I sit at my desk (or so I believe), I come to have the belief, among very many others, that there is a red book on the desk. In fact, of course, the content of the belief is a good deal more precise and specific than the formulation just given would suggest: I do not believe simply that there is a red book on the desk, but rather that there is a book of a certain approximate size, of an approximately rectangular shape, which is a certain fairly specific shade of red, and so on. But what matters for the moment is that I do not infer that there is a red book on the desk, nor does the belief result from any other sort of deliberative or ratiocinative process, whether explicit or implicit. Rather it simply occurs to me, “strikes me,” in a manner which is both involuntary and quite coercive; such a belief is, I will say, cognitively spontaneous. (p. 177)

Cognitively spontaneous beliefs like these leave little room for control or agency. Once we have the evidence, what happens with it is, according to this argument, beyond our control.

Putting these epistemic and moral observations together, we arrive at the problem of control: we cannot be morally responsible for any harm caused by our beliefs because we lack control over our beliefs. As Basu and Schroeder themselves put it, “it is just as odd to say that an agent ought to believe p as it is to say that a tornado ought to take the southeasterly path,” (2018, p. 187). In finer course, the objection goes like this:⁹

1. In order for an agent’s action to wrong someone, that action must be

⁸See, for example, *Alston* (1988); *Williams* (1970); *BonJour* (1985).

⁹*Basu and Schroeder* (2018) offer a slightly different version of this argument on p. 186.

under that agent's control.

2. Therefore, in order for a belief to wrong, an agent's beliefs must be under their control.
 3. But, agents do not have control over their beliefs.
- C. So, beliefs cannot wrong.

Thus, cases in which one feels as though they have been wronged by a belief are *all* false positives, akin to being angry at a passerby's accidental elbowing you on a crowded street. If the problem of control is correct, the partner in **Party** may have been wronged in some other way, but there is no doxastic wronging afoot.

2.1 Responding to the Problem

Basu and Schroeder's (2018) response to the problem of control targets premise three. They argue that we do, in fact, have some direct control over how we evaluate the evidence for our beliefs. To illustrate this, they consider the role of so-called *epistemic motivations*, which are "goals that we possess with respect to knowledge," (p. 193). Drawing from Arie Kruglanski's work on the psychology of closed mindedness,¹⁰ they enumerate four varieties of epistemic motivation:

1. Need for a specific closure: Preference for a particular answer to some question.
2. Need to avoid a specific closure: Preference to avoid a particular answer to some question.
3. Need for a nonspecific closure: Preference for a firm answer, as opposed to ambiguity.
4. Need to avoid a nonspecific closure: Preference for "keeping options open" and avoiding a firm answer.¹¹

¹⁰Summarized in *Kruglanski* (2013).

¹¹Paraphrased, *Kruglanski* (2013, p. 5-7).

So, for example, suppose that you have a strong preference to answer the question “is global warming real?” with ‘no’. This would constitute a need for specific closure (1) on Kruglanski’s view. Owing to this motivation, you might limit your inquiry, only performing searches for things like “global warming hoax” or “scientists lie about global warming”, or restricting your reading to news outlets that display an anti-climate change bias. In these and similar ways, epistemic motivations can alter the ways that we search for evidence, the kinds of inquiries we open and close, and so on. Thus, they can be an integral part of our belief-formation processes.

Basu and Schroeder argue that such motivations operate not only upstream and downstream, but also in the believing: they allow us to control how we interpret the evidence. Specifically, they argue that motivation (2), the need to avoid a specific closure, allows individuals to control their beliefs by applying different evidential standards to particular answers. So, in the global warming case above, if you had a type-(2) motivation to avoid coming to the conclusion that, yes, global warming is real, you might apply very high evidential standards to confirmatory evidence and relatively low standards to other evidence.¹² They conclude that,

When it comes to answering a question about the content of a belief, it is often the case that how we answer that question is a matter of our own psychology, and thus up to us. Evidence is not all that matters for epistemic agency, a lot of things matter and presumably those things could include moral reasons. Hence, there is no problem of control for doxastic wrongs. . . . (p. 194)

In this context, I take it that “up to us” means “under our control”. Given this, the argument against the problem of control is as follows: What is a matter of one’s own psychology is under their control. Since one’s goals and desires are matters of their own psychology, they are under their control. And, one’s response to evidence can be influenced by their goals and desires.

¹²It is worth noting that this does not seem to be an instance of direct control over an agent’s evaluation of their evidence, so much as it is control over their broader pattern of inquiry.

Therefore, one has control over their beliefs through their goals and desires. With this, they take themselves to have offered a full response to the problem of control.

With this picture of control, we now have a clearer picture of how doxastic wrongdoing works: Having acquired the evidence, an agent can ask a question, such as whether p , and, by employing their desires and motivations, determine how they interpret the evidence so as to arrive at or avoid a particular conclusion, in accord with those motivations. So, in a case like **Party**, the agent's belief that p is wrongful, regardless of the rest of the agent's epistemic practice, upstream and downstream, because the agent could have employed this motivation-based faculty of control and did not.

3 The Dilemma

This does not solve the problem, however, because not all beliefs are created equal.¹³ Paradigmatic cases of doxastic wrongdoing often involve beliefs that are not (or need not be) the ends of this kind of motivation-supervised, goal-responsive inquiry. Rather, these beliefs are brought about by sparse evidence and momentary reactions. They are cognitively spontaneous beliefs, at least potentially unsupervised by desires and motivations.

Let's consider two variations on **Party**: **Spontaneous** and **Suspicious**. In **Spontaneous**, we fill in the details of **Party** as follows: the believer's response to her evidence is unconscious, unreflective, and automatic. As soon as the wafting scent hits her olfactory receptors, she finds herself having connected the dots and believing that her partner has fallen off the wagon.

¹³ While this paper focuses on the dilemma presented in §3, the problems with *Basu and Schroeder's* (2018) response to the problem of control go beyond the present discussion. For example, it is far from clear that something being a matter of our own psychology implies that it is "up to us" in any way sufficient for the relevant form of agential control. However, exploring these worries, especially as they relate to Basu and Schroeder's discussion of *Hieronymi* (2006), is well beyond the scope of this paper. See *Saint-Croix* (2022b) for further discussion. This paper argues that the evidence presented in *Basu and Schroeder's* (2018) only provides support for managerial and evaluative control, in *Hieronymi's* (2006) terms, and not the direct control necessary to respond to the argument that Basu and Schroeder take as their target.

By contrast, in **Suspicious**, her response is not automatic. The scent of wine hangs in the air, and, a moment later, the believer thinks to ask herself, “Wait, did my partner fall off the wagon?” She chooses to pursue this question and, upon then evaluating her evidence, finds herself with the belief that her partner has, indeed, fallen off the wagon. In both cases, the agent evaluates her evidence identically: she finds the wafting scent, in combination with the rest of her evidence, sufficient support for the proposition *my partner fell off the wagon*, and thereby believes it.

If you share the intuition that *something* has gone wrong in **Party** even before that belief causes the agent to behave in some way, and irrespective of whether the partner is aware of the belief, this pair of cases poses a problem for doxastic wronging as a diagnosis. Either these cases differ with respect to whether they involve an instance of doxastic wronging, in which case evaluative resilience is false, or they do not, in which case the problem of control resurfaces.

To begin, suppose that **Spontaneous** and **Suspicious** are both instances of doxastic wronging. Initially, then, it appears that there is no violation of evaluative resilience: the evaluation of the cases go hand-in-hand, upstream difference notwithstanding. But, a new problem of control arises. Recall that Basu and Schroeder’s response to the original problem of control discussed in §2 concedes that the control condition must be met. Their response argues that we can exercise control in the evaluation of evidence, choosing the answer at which we are motivated to arrive. Thus, in **Party**, it is the failure to exercise this control that is at issue. The wrong here is the same kind of wrong that I commit if I fail to check my blind spot before merging onto the freeway—this failure is a kind of negligence. However, we must take care with the notion of control in this context.

While we may be able to choose to ask and answer particular questions for ourselves, it is not at all clear that we can *prevent* ourselves from connecting evidence to an evaluation thereof. If we are to be blameworthy for failing to prevent ourselves from answering a question in a particular way, we must have been able to do so in the first place. And, in order to be able to do so,

we must be aware of and able to intervene in the process.

But, a great deal of psychological evidence suggests that we do not have this kind of awareness or capacity to intervene and control. To see this, refrain from thinking about a white bear for the next five minutes. If you are like the subjects in *Wegner et al.'s* (1987) famous study, you will have a difficult time with this task. *Wegner et al.* (1987) found that, “participants indicated thinking about a white bear (through a bell ring, mention, or both) more than once per minute even when directly instructed to try not to think of a white bear,” (p. 7). Furthermore, they found that subjects who had previously been asked to suppress this thought “showed significantly more tokens of thought about the bear than did subjects who were asked to think about a white bear from the outset,” (p. 5). In other words, we are largely unsuccessful at suppressing unwanted thoughts. This is an important observation because unwanted thoughts are all too common. Tracking the responses of participants across 13 countries, *Radomsky et al.* (2014) found that 93.6% of participants reported experiencing unwanted, intrusive thoughts. Not only do we fail in suppressing unwanted thoughts, but we also almost all have them. So, while we might be able to observe our responses to our evidence and shift our behavior around that response, our minds do a great deal without our input, consent, awareness, or preference, and there is no reason to think that evaluating evidence is uniquely free from such meandering—not all beliefs are created equal.

At this stage, one might object: Even if much of our evidence evaluation is outside of our direct control, we can surely recognize patterns in our ways of responding to evidence and, perhaps over time, change those patterns. For example, some have argued that implicit bias can be combated by recognition and conscious effort to diminish the bias. Strategies for carrying this out include deliberate intergroup contact (*Dasgupta and Rivera, 2008*), exposing oneself to counter-examples to the bias (*Dasgupta and Greenwald, 2001*), and evaluative conditioning (*De Houwer, 2011*). Insofar as these strategies are effective, then we clearly have at least some control over how we respond to

our evidence and, thereby, over the resultant beliefs.¹⁴ And this is entirely correct. However, this does not provide the kind of control required by Basu and Schroeder's account of doxastic wrongdoing because the causal structure of a wrongdoing brought about by failing to exercise this kind of control involves upstream management of belief-forming dispositions. But, doxastic wrongdoing requires that the causal structure of the wrongdoing be located in the belief itself, not the any upstream epistemic practices of the agent. So, while this kind of control may be possible, and may even be morally required, it is not the kind of control necessary to circumvent the problem of control.

All of this suggests that if we stick to evaluative resilience, so that our hypothesis is that both **Spontaneous** and **Suspicious** wrong, Basu and Schroeder's commitment to the control condition is violated in cases of spontaneous belief formation.

This might lead one to the other horn of the dilemma. Suppose, as seems intuitive, that there is no wrongdoing in **Spontaneous**, so that the two cases differ with respect to whether doxastic wrongdoing has occurred. According to evaluative resilience, however, given any two instances in which an agent believes P both beliefs must either both doxastically wrong or fail to wrong together, whenever all differences between the two cases concern either the agent's upstream epistemic practices, dispositions, capacities, or motivations or the downstream consequences of the belief. Here, the difference is that the believer in **Suspicious** chooses to pursue the question of whether her partner fell off the wagon. But, this is not a difference in the belief itself or the broader context. This is a difference in her upstream behavior, be it a nasty disposition, a motivation for epistemic closure, or simply a matter of her having a particularly thorough epistemic practice of trying to answer every question that occurs to her. Granted, it may not be far upstream, but it is upstream nonetheless. So, if the immediate belief in **Spontaneous** does not wrong while the belief in **Suspicious** does, then we have a situation that violates evaluative resilience: in both cases, the same agent believes the same

¹⁴This is quite close to what *Fischer and Ravizza* (1998) refer to as *guidance control*.

proposition, but one is an instance of doxastic wronging while the other is not. Evaluative resilience must be false and we must look elsewhere for the cause of the wronging.

To this, Basu and Schroeder might respond that the upstream difference is still incidental; it is in the believing that the difference occurs. But, if the upstream difference merely makes it possible for the agent to intervene in the evaluation of evidence in **Suspicious** a way that she was not able to do in **Spontaneous**, so that the issue is a matter of negligence, then we still have a failure that is upstream from the belief itself. Failing to intervene on the belief-forming process is distinct from believing. Similarly, if the issue is that the evaluation of evidence goes differently, so that the believer either brings herself to believe in **Suspicious** (perhaps by choosing to apply lax standards to the evaluation of her evidence, which just so happen to be the same lax standards that occur in **Spontaneous**), this, too, is distinct from believing. It has predictable consequences for one's beliefs, but bringing oneself to believe is distinct from believing.¹⁵ It is upstream from the belief itself.

So, it seems we must either give up on evaluative resilience or violate the control condition. It is important to note that this argument does not show that doxastic wronging is impossible. It merely shows that *Basu and Schroeder's* (2018) response to the problem of control does not account for purportedly paradigmatic cases of doxastic wronging like **Party**. A proponent of doxastic wronging might save it by somehow jettisoning cases like **Spontaneous** while finding a way to keep cases like **Suspicious**, thereby maintaining evaluative resilience. Or, they might offer a more comprehensive response to the problem of control. For that has been said here, **Suspicious** may still be a case of doxastic wronging.¹⁶

With this dilemma in mind, however, it is worth thinking more carefully about what actually goes wrong in instances of doxastic wronging—doing

¹⁵Hieronymi, a key interlocutor for Basu and Schroeder's argument, points this out as well (2006, p. 61-62).

¹⁶But, see *Saint-Croix* (2022b) for further discussion. *cf* fn. 13.

so will, I believe, shed some light on the situation and suggest an alternative way forward.

4 The Wrongs of Doxastic Wronging

In “What we (epistemically) owe to each other”, [Basu \(2019c\)](#) argues that there is a harm involved in these cases: *relational harm*. She draws this idea out by considering the following case:

Racist Hermit. Suppose a racist hermit in the woods discovers trash containing an alumni newsletter from Sanjeev’s university, which includes Sanjeev’s photo. The hermit immediately concludes that the pictured person—Sanjeev—smells of curry. Suppose also that Sanjeev happens to have recently made curry, so in this instance the hermit’s belief is true—Sanjeev does smell of curry. Has the hermit wronged Sanjeev? (p. 919)

In Basu’s analysis, the wronging in this case arises from a harm: the hermit fails to relate to Sanjeev in the way that he ought to, and this harms Sanjeev. In particular, the hermit is failing to take the *participant stance* toward Sanjeev. Basu draws on Strawson to illuminate this idea, writing

Strawson argues that we ought take an involved stance towards others and recognize that, just as our own self-understanding relies on the attitudes that others take towards us, so too does their understanding rely on the attitudes we take towards them. Extending this idea, we can presumably hold that the participant stance creates demands upon the beliefs we hold of others. ([Basu, 2019c](#), p. 923-4)

Even if it is right that this Strawsonian involvement extends to our epistemic attitudes regarding others, it nevertheless seems that this judgment of harm depends on much more than the believing of a particular proposition. In the setting of the case, we learn that the hermit is a racist. This informs our

interpretation of the details: At the very least, we assume that the hermit is using a cultural stereotype to infer that Sanjeev smells of curry. We might also assume that he takes this stereotype to be a negative one and that this is not his first foray into racism. We might imagine him as a child cruelly bullying Indian classmates about their lunches, or mocking their accents. And, we might assume that he takes this to be an acceptable way to reason about Sanjeev—that upon reflection, he will neither regret nor reconsider his formation of the belief. All of these deeply troubling possibilities inform our understanding of how he relates to Sanjeev, of the stance he takes toward him.

As a result, it is difficult to separate out the moral weight of the hermit's upstream and downstream behaviors in analyzing this case. But, we can compare it with other cases in which the hermit forms the same belief. Were the hermit to have read an autobiography of Sanjeev in which Sanjeev himself explains that he always smells of curry (unlikely though that certainly is), the hermit's forming that belief would be a matter of accepting Sanjeev's own testimony.¹⁷ This would not plausibly be an instance of relating improperly to Sanjeev. But, the difference here is in the evidential basis for the belief and the manner of its being formed, both of which are distinct from the belief itself. We cannot see the hermit's stance from his belief alone. This suggests that the causal and normative structure of any wronging in these cases is not isolated to the belief itself, contrary to *Basu and Schroeder's* (2018) conception of doxastic wronging.¹⁸

Moreover, returning to **Party**, relational harm does not seem like a plausible diagnosis of the wronging at hand. Spontaneously or unwittingly

¹⁷Even here, however, it is difficult not to draw the upstream dispositions into the moral analysis: one easily imagines the racist hermit smugly self-satisfied at this development.

¹⁸Beyond this, it is difficult to see how the belief amounts to a *harm* to Sanjeev. It may be poor moral conditioning for the hermit and diminish his capacity to relate to others, but Sanjeev seems to be equally well-off regardless of the hermit. Sanjeev may prefer that everyone in the world believe only pleasant things about him, but even so, it is difficult to see how the hermit's failure to cooperate with this preference harms Sanjeev. The mere fact that this is something Sanjeev wants does not seem to be sufficient to generate a moral claim on the mental states of others.

forming a belief about someone does not constitute taking any *substantive* stance toward them. To see this, consider two further variations on **Party: Rumination** and **Regret**. In the original description of **Party**, the story stops before the believer has a moment to react to her immediate response to the evidence. We do not learn whether we are in the **Regret** case, in which she goes on to regret her reaction and dismiss the momentary thought, or the **Rumination** case, in which she chews on her suspicions, focusing her attention on the smell of her partner's breath and the quality of her step as they share small talk about the party. Yet, it seems like this matters deeply to the question of how she relates to her partner—it defines this instance of relating to her. This is evident from considering a version of **Rumination** without the belief: had the believer not come to the belief that her partner had fallen off the wagon, but nevertheless continued to ruminate on the question, this would be just as much a matter of poorly relating to her partner. So, relational harm is an unlikely candidate for the harm of doxastic wronging, if we are to maintain the account of doxastic wronging given in *Basu and Schroeder (2018)*.

Another account comes from *Schroeder (2018)*, who argues in “When Beliefs Wrong” that beliefs wrong when they falsely diminish their subject. By this, Schroeder means that the propositional content of the belief is false and that believing it of the subject underestimates their *agential contribution*. An agential contribution can be less because “because it is a *worse* contribution” or “because it is *less* of a contribution” (*Schroeder, 2018*, p. 124). On this account, then, believing of a coworker that she was only hired because she is a woman is an instance of doxastic wronging because it diminishes the subject's agency in the hiring process, such as her hard work in school, her careful crafting of application materials, her quality interviewing skills, and so on.

Yet, in some cases, it seems that overestimation of someone's agential contribution might be wrongful. For example, consider a hiring manager who accurately assesses the resume of this woman applicant, but subsequently *overestimates* the agential contributions of a male applicant with an otherwise

similar application. In this case, it seems like the overestimation of the male applicant is wrongful, yet, because it involves no underestimation of the degree or quality of either applicant's agency, there is no doxastic wronging involved.

The fundamental attribution error, a common cognitive bias, gives rise to another odd consequence of this view. This bias is characterized by a tendency to attribute events or behaviors to individuals' stable character traits rather than their circumstances. So, for example, we might attribute someone's tripping up the stairs to their clumsiness rather than the unevenness of the stairs. On Schroeder's analysis of doxastic wronging, this bias turns out to be a kind of shield against doxastic wronging: by systematically overestimating the agential contributions of others, we avoid the risk of wronging them.

Most importantly, this view leaves out plausible instances of doxastic wronging without a satisfying explanation. To see this, let's begin with a case familiar from the literature on both doxastic wronging and moral encroachment:¹⁹

Cosmos Club. In the summer of 1995, historian John Hope Franklin was awarded the Presidential Medal of Freedom, the nation's highest civilian honor. On the night before the award ceremony, Franklin hosted a dinner for a small group of friends at the Cosmos Club, a Washington DC social organization of which he was a member. During a stroll through the club, a woman presented him with her coat check and demanded that he retrieve her coat.

In the moral encroachment literature, this example is meant to draw out the intuition that the believer is not (epistemically) justified in her belief. Here, we're interested in the very different intuition that this belief is an instance of doxastic wronging: the believer wrongs Franklin, not by acting

¹⁹Versions of this case appear in *Gendler (2011)*, *Basu and Schroeder (2018)*, *Basu (2019b)*, and many others. This version is lightly modified from *Gendler (2011, p. 35)*

on her belief, but rather simply by having that belief. On Schoeder’s account, this case *does* constitute an instance of doxastic wronging because it falsely diminishes Franklin. But, now consider a variant of the case in which the woman happens to address the Black man standing next to Franklin, who *is* a staff member. In this case, she makes exact same set of inferences on the same information, but this time is lucky and correct. But this is an unwelcome result: if the first version of the case is odious, this version ought to be as well, accuracy notwithstanding.²⁰

Moreover, even if we accept that false diminishment is the harm of doxastic wronging, the same problem that arose for Basu’s account arises here: merely having the belief is insubstantial—it is maintaining the belief, choosing not to re-evaluate it, allowing it to be part of your understanding of the subject that diminishes the person. Merely being struck by a belief or a thought, then quickly reassessing and dismissing it does not substantively diminish the subject.

With these worries in mind, I take the question of what, exactly, goes wrong in cases of doxastic wronging remains somewhat unresolved—at least, if belief is at the core of the wronging involved.

5 Attention

The opening lines of “Doxastic Wronging” take a brief ecclesiastical excursion into the Book of Common Prayer, quoting the ritual immediately preceding the Eucharist: “we have sinned against you in thought, word, and deed,” (*Basu and Schroeder, 2018*, p. 181). This is meant to demonstrate that, in contrast with philosophers’ epistemic scrupulousness, doxastic wronging is a commonplace, intuitive idea among the folk. Yet, the notion of a sin of

²⁰*Schroeder (2018, §4)* responds to this kind of worry by appealing to a moral encroachment thesis, and arguing that someone in a case like this does not wrong the staff member by believing what they do, but does subjectively wrong them by believing on insufficient evidence owing to moral encroachment. One of the benefits of a view like the one to be presented in §6 is that we can account for cases like these without appealing to a controversial thesis like moral encroachment. See *Basu (2019b)* for introduction to the concept of moral encroachment and *Gardiner (2018)* for a critical discussion.

thought, at least as it is taught in the Catholic Church, is markedly distinct from the strictly constrained concept of doxastic wrongdoing. For example, in the *Summa Theologiae*, Thomas Aquinas makes clear that it is not the occurrence of a thought, but the entertainment of that thought (the *morose delectation*) that constitutes a sin.²¹ Finding oneself with a thought that would be problematic were one to dwell upon it, it is the choice to do so that makes it sinful. Sins of thought are sins of willful thinking, not momentary thoughts.

Nevertheless, I think Basu and Schroeder are right that something goes wrong these cases, and that it should be understood as a matter of epistemic morality. Furthermore, I agree that it has to do with the epistemic motivations they identify. However, these epistemic motivations are not themselves epistemic attitudes. They are *desires* with epistemic objects. For example, epistemic motivations falling under the category *the need for specific closure* are desires to be in a state in which you bear the belief relation to a particular proposition. As such, it is difficult to offer an analysis of these motivations as any kind of epistemic misstep in and of themselves—desires are not subject to epistemic evaluation. For example, your desire to believe that your spouse is faithful when they are not is epistemically impeccable so long as it is epistemically inert—it is merely a desire.

But, these goals desires can influence our epistemic practice by influencing our patterns of attention. In doing so, they may inspire epistemically dubious behavior, such as willfully ignoring evidence or failing to give equal scrutiny to supporting and undermining evidence. But the desire is merely giving reasons. Choosing to act epistemically on those reasons is a different matter. Such choices also resonate with Aquinas' *morose delectation*—allowing oneself to pay attention to a spontaneous or intrusive thought rather than allowing it to pass or directing one's attention elsewhere makes the difference between a sin of thought and the misfortunes of an active mind.

The epistemology of attention, however, is under-explored.²² This over-

²¹See, in particular, ST II-I, q. 74, a. 4-6 .

²²Some exceptions include *Smith and Archer (2020)*, *Gardiner (2022)*, *Siegel (2017)*, *Munton*

sight is unfortunate because it is here that we may have a clear form of epistemic agency, as I will argue below. Moreover, this is a conception of epistemic agency that will better trace the contours of our intuitions about **Party** and other purported cases of doxastic wronging. If this is right, it points to a way to ameliorate any worries about whether there is some work for the concept of epistemic morality to do here—though it means giving up on evaluative resilience.

5.1 Goal-Oriented Structuralism

While attention is sorely neglected within epistemology, it is also only recently that philosophers of mind, cognitive science, and related fields have begun to focus on attention in earnest. Perhaps both lacunae owe themselves to the fact that, as William James wrote, “everyone knows what attention is,” (1890). Unsurprisingly, though, this claim is false. Where attention is widely discussed—within psychology, neuroscience, and behavioral economics—conceptions of attention vary dramatically. Sebastian Watzl observes that, even as many of these papers and presentations begin with the remainder of James’ pithy description of attention²³—“[Attention] is the taking possession by the mind, in clear and vivid form, of one out of what seem several simultaneously possible objects or trains of thought. Focalization, concentration, of consciousness are its essence. . .” (James, 1890, p. 403-4)—they soon diverge. A common strand among them, however, is the idea that agents have at least some measure of direct control over their attention.²⁴ This is not to say that conscious attention is always under the control of an agent. In fact, it is very often not. For example, while you might marshal your conscious attention toward practicing the violin, a car crash outside your window will inevitably draw your attention away from that task, at least momentarily. Nevertheless, while much controversy attends the question of whether belief

(2021), and *Fairweather and Montemayor* (2017).

²³Watzl (2011a)

²⁴See, e.g., *Spence* (2012), *Mackie et al.* (2013), and *Schmeichel* (2007).

is voluntary,²⁵ the claim that attention is voluntary is far less controversial.

For our purposes, I will adopt a subject-level conception of conscious attention based largely on Watzl's work.²⁶ But, it is worth noting that there are many alternatives to this account, some reductivist or even altogether eliminativist.²⁷ The Goal-Oriented Structuralism I'll suggest here is neither.

Our focus is on subject-level attention, which may be contrasted with understanding attention as a sub-subject (or subpersonal) phenomenon.²⁸ On this view, attention is a matter of an individual agent, rather than being identified with some subsystem of that individual. There are two key features of a subject-level account of attention: First, it captures the sense in which there is "something it is like" to attend—there is a phenomenal experience of attending that is difficult to capture in terms of sub-subject mechanisms. Second, a subject-level account of attention provides a unifying account of the nature of attention through this phenomenal perspective. This is a vexing problem for sub-subject level accounts, because there seems to be such a wide variety of mechanisms recruited in service of attending that some, such as *Allport* (1993), have argued that there is no such thing as attention.

Watzl's (2011b) implementation of this approach is *structuralism*:

Structuralism. Consciously attending to something consists in the conscious mental process of structuring one's stream of consciousness so that some parts of it are more central than others.

On this account, the key phenomenological experience is centering—whatever one is attending to central to one's consciousness, while everything else is relatively peripheral. To get a sense of this phenomenology, imagine listen-

²⁵See, e.g., *Williams* (1970), *Hieronymi* (2006), *Bennett* (1990).

²⁶*Watzl* (2011b, 2017)

²⁷Jesse Prinz' work (e.g. 2011) adopts a reductionist stance, identifying attention with underlying mechanisms, and *Anderson* (2011) defends an eliminativist position. There are many other varieties of theories of attention that are neither reductivist or eliminativist. See *Watzl* (2017, Ch. 1) and *Mole* (2021) for respectively more and less opinionated reviews of these alternatives.

²⁸I follow *Watzl* (2017) in using the term 'subject-level' rather than 'personal-level', to avoid the implication that personhood, rather than mere subjecthood, is required of attending things.

ing to a friend's voice in a crowded pub. Focusing on what they are saying requires letting the other conversations around you take on a characteristically peripheral feeling. This is the kind of structuring around which Watzl builds the account.

In more detail, *Watzl's* (2017, Ch. 4) explains that attending to something is a matter of regulating the relevant constituents of one's conscious experience by priority, be they mental states, processes, or any other part of one's occurrent psychology. The object of one's attention, on this view, is whatever is at the top of that structure at a particular time.²⁹

The question, then, becomes: on what basis do we form these priority structures? To a significant degree, our goals and interests that guide our regulation of these priority structures. Consider what happens when you sit down to write a paper that you're excited to write. You know that you need a collection of things in place before you can accomplish your goal: you'll need good lighting, a decent chair, and a cup of hot coffee. You push these things to the top of your priority structure. Once they are there, aspects of your stream of consciousness that are relevant to them become foregrounded. The location of the coffee bag, the numbers on the scale, the temperature of the kettle—all of these things are vividly in focus, whereas the rest (What's the cat doing? Does the floor need to be swept? Are the plants drooping?) moves to the periphery. From the periphery, they come into focus only to the extent that they are relevant to things higher in your priority structure.

Watzl's structuralism also incorporates what he calls "freedom", which is the claim that "attention does not consist in consciously selecting that thing for any particular purpose," (*Watzl*, 2011b, p. 154). Watzl argues that we should reject views like *Wu's* (2011) on which attention is merely "selection for action", because such selection for action is not sufficient for attention. And this is right: there are many cases in which we unconsciously act on perceptions and the like. For example, when ascending a staircase, you perceive the distance between one stair and the next, and you select this perception

²⁹I have said little here about what is for something to be at the top of a priority structure. For the sake of space, I refer the reader to (*Watzl*, 2017, Part II).

for the action of placing your foot on the next step. This selection for action is evident in the fact that slight variations in these distances frequently cause missteps. But, we are almost never conscious of doing this, and certainly not with the degree of precision we in fact employ. So, it is reasonable to conclude that selection for action is not sufficient for attention. However, this does not mean that purpose plays *no* role.

For example, purpose plays an important role discerning whether we are *successful* in attending to one thing or another. When speaking of subject-level attention, we mark a distinction between trying to attend to something and succeeding at attending to it.

Corner Coffee. It's Sunday morning and you've decided to try out a new coffee shop, Corner Coffee, hoping to get work done on your newest writing project. Unbeknownst to you, Corner Coffee is a favorite hot spot for the church-going crowd and, just as you're settling in with your iced coffee, they begin filling in the other tables in twos and threes. They chat, gossip, and laugh, distracting you from your writing.

In this case, the fact that you are structuring your stream of conscious experience—pushing aside this, focusing on that—does not give us a full picture of the relevant phenomenal experience. And, capturing the phenomenology is an important part of any subject-level account. But, for this we need to know that you are trying to focus on writing—it is only relative to this goal that the actual structuring of your consciousness is accurately intelligible.

We can also structure our consciousness by *not* paying attention to anything. Take a moment to stare idly into space. This, too, is structuring one's stream of consciousness, but it is not attending to anything—attention has an object. It is attention for a *reason*. So, structuring one's stream of consciousness alone is also not sufficient either. Rather, attention structuring one's consciousness toward some goal. Nevertheless, I take Watzl's broader view, that attention is a matter of structure, to be correct.

With this in mind, I propose the following amended version of Watzl's structuralist account:

Goal-Oriented Structuralism. Consciously attending to something consists in the conscious mental process of structuring one's stream of consciousness toward a particular purpose, so that some parts are more central than others.

There is a great deal more that might be said about the nature of attention, but this much should suffice for the purpose of exploring the extent to which attention is an epistemic practice.

5.2 Attention as an epistemic practice

Given this account of conscious, subject-level attention, it is clear that attention is an important driver of our epistemic lives. Consider the following examples:

- **Closing inquiry:** Conventionally, one might close an inquiry upon reaching a clear answer to it. But, this is not the only way to close an inquiry. Less directly, shifting one's attention may also close inquiries. Suppose, for example, that you're back at Corner Coffee, and this time you overhear an intriguing conversation several tables over. While waiting for a friend, you focus in, drawing the discussants' voices to the foreground of your mind. It's a strain, but you're able to follow along, and you try to figure out who exactly broke up with who. But, when your friend arrives and you shift your attention away from the other patrons to focus on your friend, you thereby abandon your inquiry. Even though you never found out who broke up with who, your restructuring your conscious attention in this way effectively closes the inquiry.
- **Opening inquiry:** When some topic is central to your conscious experience, related things in your stream of consciousness become likely objects of inquiry because they are less peripheral than they might

otherwise be. For example, as you try to figure out who broke up with who, you might realize that one of the participants in the conversation has a slight East Coast accent, and begin to wonder where they are from. This is an inquiry that would not have been opened were it not for your attention to the conversation.

- **Selecting evidence-gathering processes:** Sticking with this case, we might suppose that you *don't* abandon the inquiry when your friend arrives. Instead, you motion furtively in the direction of the pair, explain what's going on, and recruit your friend's keen ear as well. Now, keeping their conversation at the center of your attention has changed the kind of evidence-gathering process you're employing. Not only are you eavesdropping, but you're also gaining the testimony of a fellow eavesdropper.
- **Obtaining Evidence:** Insofar as our consciousness is structured by our attention, so too is the evidence we obtain. You're more likely to catch the details of the conversation if it, rather than your friend's new puppy, is at the center of your attention.
- **Success in inquiry:** Were the situation a few tables over merely at the periphery of your attention, you might accidentally catch the story. But, intentionally eavesdropping, thereby structuring your conscious experience in such a way as to foreground their conversation, clearly increases the likelihood that you'll succeed that this task.

With these examples in hand, it is clear that attention is quite central to our epistemic practices. And, insofar as this conscious, subject-level attention is under our control, it is a locus of epistemic agency. So, while attention does not provide *direct* control over our doxastic attitudes, it is a means by which we manage our inquiries and, thereby, influence their outcomes.³⁰

But, does all of this make attention the kind of thing that is subject to epistemic morality? In some cases, I take it that the answer is “no”: if you

³⁰In *Hieronymi's* (2006) terms, attention is a mechanism by which we enact both *managerial* and *evaluative* control over our beliefs, but it does not provide direct control.

are focusing on your breath with the goal of meditation, this is likely not a practice governed by epistemic norms, morally-infused or otherwise. However, insofar as your goal is alethic, this is more plausible. In particular, when the kinds of epistemic motivations Basu and Schroeder identify are the goals setting your priority structures, and you are using these goals to shape your conscious experience, this does seem to be an epistemic practice. Thus, goal-oriented structuralism allows us to divide practices of attending into epistemic and non-epistemic in virtue of the goals that ground them. Epistemic attention is attention structured by epistemic goals. So, for example, if I decide to open an inquiry into the pharaoh Tutankhamun, I thereby adopt the goal of learning about him. Insofar as I structure my conscious experience in service of this goal, I am engaging epistemic attention, since learning about Tutankhamun is an epistemic goal.

6 Rumination's Relation to Wronging: Attentive Epistemic Wronging

We return now to epistemic morality. This section draws on our discussion of attention to offer an account of what I'll call *attentive epistemic wrongdoing*. This concept, I argue, better accounts for the kinds of harms that Basu and Schroeder identify and provides a clear locus of agency.

To get there, let us return to **Party**. In this case, it is easy to see how rumination might make a difference to wrongdoing. The evidence the believer receives is strong, but not incontrovertible, despite the fact that the evidence initially convinces the believer that her partner has fallen off the wagon. In our **Regret** variant of the case, the believer might bat the belief down by focusing on the question of what else is relevant to whether her partner fell off the wagon. For example, she might recall that it's been several years since there was any issue or bring to mind a forgotten remark the partner made about how proud she is of herself. She might also ask herself whether there are other ways that wine smells might end up wafting in her direction, and,

considering these possibilities, let her confidence in the proposition fall shy of belief. In this case, it would seem strange for the partner to feel wronged or to claim that the believer has demonstrated some moral failing—the believer did what any of us would hope: as soon as she had the opportunity, she considered her partner fairly and with charity, and changed her mind.

By contrast, in the **Suspicious** and **Rumination** variants, the believer's attention is and remains focused on the negative possibilities. She seeks only confirmatory, incriminating evidence, prioritizing affirmation of her belief over everything else. She does not take the time to consider alternative explanations, and she is not focused on the conversation she's having with her partner, nor on appreciating the good evening she's had, etc—all of these things drift to the periphery. This disordered regard for the partner is a betrayal of their relationship (absent excusing conditions) because long-term romantic relationships generally bring with them mutual expectations of attention, such as listening to one another and appreciating each other rather than focusing on one another's faults. This betrayal is what makes the resulting conviction that the partner fell off the wagon hurtful. In both of these variations, attention makes the difference. The believer structures her attention around her partner and the proposition in ways that constitute morally significant disregard for the partner.

With this in mind, I want to suggest that at least some paradigmatic cases of doxastic wrongdoing are better understood as instances of (directed) attentional epistemic wrongdoing:³¹

Attentional Epistemic Wronging. A wrong committed by one person against another in virtue of the patterns of attention that brought them to their doxastic attitude.

A few points of clarification are in order. First, attentional epistemic wrongdoing does not involve evaluative resilience. Unlike Basu and Schroeder's doxastic wrongdoing, attentional epistemic wrongdoing requires that we look

³¹I leave open the question of whether there are non-directed attentive epistemic wrongs, though I suspect that the answer is 'yes'.

upstream or downstream of the belief itself to understand the causal and normative structure of the wrongdoing. An individual belief may be a part of or a consequence of a pattern of attention, but cannot itself constitute a pattern of attention. Second, because the control that renders us morally responsible in cases of attentive epistemic wrongdoing is control over attention, which comes in degrees, our responsibility for wrongdoing of this kind also comes in degrees. Finally, the failing in cases of attentive epistemic wrongdoing is a moral failing, not an epistemic one. While epistemic failings, such as poor reasoning or carelessness in evaluating one's evidence, may attend instances of attentive epistemic wrongdoing, they are neither necessary nor sufficient to their existence.

This account of the wrongdoing involved in cases like **Party** has some important advantages over doxastic wrongdoing. For example, it allows us to make sense of a common intuition about **Party**: the wrongfulness of the believer's conviction is not only a matter of control, but also of how reasonable it is for her to worry. If it was only last week that the partner last fell off the wagon, the partner's focus on whether it happened again this week is neither irrational nor wrongful. If it was ten years ago, on the other hand, it seems quite unfair. And there seems to be some continuous relation of reasonableness between these points. By accounting for the wrongdoing in terms of patterns of attention rather than the belief itself, it is clearer where and why factors like these can influence the analysis of particular instances.

Similarly, this account also provides a clearer distinction between cases in which it seems there has been some wrongdoing and those in which there has not. Consider a variant of **Party** on which the partner is violent and abusive when drunk. In such a case, it seems that the believer is well within her rights to be vigilant about her partner's behavior, regardless of whether the most recent relapse was five days ago or five years ago.

Another important feature of attentional epistemic wrongdoing is that it is not restricted to beliefs alone. Any doxastic attitude—belief, suspension of judgment, uncertainty, and so on—can be involved. This allows us to capture neglectful cases. Suppose, for example, that you believe of a student,

call her Genevieve, that *Genevieve is a decent student*. But, as it happens Genevieve is an excellent student and has been trying to get you to read her most recent essay for months. In this case, your belief is entirely rational, but nevertheless is neglectful—you ignore the evidence available to you to improve the quality of your epistemic attitudes. And, in this case, it's clear that the wrongfulness scales with the relationship: If Genevieve is an undergraduate at another university in a different field, it's difficult to identify any sense in which you've wronged her. However, if you are the chair of her dissertation committee, this is a serious wronging. But, again, notice that this wronging depends not only on the context and the content of the belief, but also the attentive history behind it: If you *did* read the paper and simply misunderstood it or otherwise failed to sufficiently appreciate it, this is not wronging her. Rather, the wronging that manifests in belief lies in the attentive choices that came before. The causal structure of the wronging requires this unexcused, neglectful behavior. What makes this wronging epistemic is not that it is a violation of epistemic normativity, but rather that the agency involved—control over epistemic attention—is epistemic.

In sum, then, in cases like **Party**, it is not the belief alone that causes any wronging. Rather, it is the upstream decision making—the patterns of attention—that result in that belief that wrong. By way of analogy, if I'm holding a rock and decide to throw it at someone, I have wronged them by throwing the rock at them (given that I reasonably believe that they would not want me to hit them with the rock and that I don't have good reason to harm them). The rock's hitting them does not constitute an additional wronging. After it leaves my hand, the final state of affairs—whether it hits them or not—is out of my control. Of course, the amends I'll need to make will differ depending on whether the rock hits them and how much damage it does, but my throwing it is what makes me responsible for whatever harms come of that action. I can be more or less careful about how I direct my attention, but what happens to my doxastic states when I encounter (or fail to encounter) new evidence as a result of my choice to direct my attention is beyond my control.

6.1 Obligations of Attention

It is worth dwelling on obligations of attention. Why think that particular patterns of attention, such as those observed in the **Suspicious** and **Rumination** variants, amount to wrongs?

Our relationships with others can engender obligations to them, some of which concern attention. Plausibly, for example, our children and pets have strong claim to our attention. As do spouses, partners, friends, and others. In the case of children and pets, the source of the moral significance of our attentional obligations is clear: these are relationships in which one party is dependent on the other for meeting its basic survival needs. So, all else being equal, the other party is obligated to attend to the dependent insofar as is required to correctly interpret and meet those needs. Attending is an enabling act for many other obligations as well, such as training and play, which are less dire but no less real.

In the case of social relationships, attentional obligations and expectations are often unspoken; they are simply how the culture we are a part of understands, e.g., long-term partnered relationships to operate.³² In some cases, such as student-teacher relationships or client-provider relationships, these obligations are reasonably clear: teachers, for example, have an obligation to pay attention to their students' assigned work. But, they may also require negotiation and careful navigation. In fine detail, the contours of attentional obligations are determined by the particular relationship, and it is the responsibility of the parties to that relationship to understand them. Nevertheless, because we rely on these relationships and the expectations they involve, this reliance can generate morally significant obligations. Furthermore, since attention, as construed above, is susceptible to agential control, failure to uphold these obligations may reasonably be construed as wrongdoing.

If my analysis of cases like **Party** is correct, we do lose evaluative resilience—the wrongdoing in this case is no longer a matter of the belief alone. But, in

³²This is not to suggest that such arrangements are ideal. Indeed, asymmetrical expectations and understandings of relationships are the source of much unintended pain, especially when it comes to failing to meet obligations of attention.

its stead, we gain a clearer explanation for how and why these epistemic practices are morally significant: They are morally significant because they involve violations of obligations of attention that arise from the relationship between the wronged party and the wronging party.

This is also why attentional epistemic wrongdoing makes better sense of the kinds of diagnoses of harm we saw in section 4 from *Schroeder (2018)* and *Basu (2019c)*. Both identify ways that we relate poorly to others in instances of doxastic wrongdoing. *Schroeder's (2018, p. 124)* view is that the wrong of doxastic wrongdoing is false diminishment: “beliefs wrong you only when they diminish you, [...] only when they bring you down.” Yet, it is quite difficult to see how I can in fact bring someone down when I am merely struck by a belief that I soon come to disavow—especially in a case like **Regret**, where it is my reflection on the quality of the individual and their actions that alters my belief. By contrast, the patterns of attention we cultivate around others *can* be diminishing. If the believer in **Party** shifts from treating her partner as a trusted companion to treating her as a duplicitous addict, this diminishes her. If, upon realizing that he is Black, the woman in **Cosmos Club** treats John Hope Franklin (or the man standing next to him) as having no more inferentially relevant qualities than his skin color, this diminishes him, regardless of whether her conclusion is correct. Both of these manifest primarily in ways we attend.

Basu (2019c) proposes that we wrong others when we do not engage with them in the participant stance—when we do not relate to them as we ought to. Again, forming a spontaneous belief is not substantially taking *any* stance toward someone, at least not in the morally significant, personally involved way necessary for the participant stance. Rather, it is how we manage our epistemic practices upon their being brought to our attention that determines the stance we take toward them. As the story begins, the woman in **Cosmos Club** is focused on retrieving her coat. That goal is structuring her attention. Upon seeing someone she takes to be instrumental to that goal—Franklin—he exists in her stream of consciousness only as a means to that end. She either does not consider or does not care about the harm she might do in

treating him as a staff member, and this failure to re-organize her attention when given the opportunity is a failure to engage with him in the participant stance.

6.2 What happened to belief?

Nevertheless, one might wonder what has happened to beliefs. In a case like **Party**, for example, one can imagine the subject of the belief complaining to her partner upon finding out about it by saying, “I can’t believe you would think that of me!” Or, “how could you believe I would do that?” Both reproaches seem to target the belief. So, if the source of the wrongdoing is actually a matter of decisions about attention, what explains this?

To return to the rock-throwing analogy, while the throw is wrongful, it’s the impact that hurts. In the case of beliefs, the focus of reproach (or guilt) is the belief itself because that is the outcome that the subject wishes to avoid. But, the outcome and its cause are distinct. The fact that our emotional focus is often on the belief itself does not indicate that the belief is a wrongdoing, but only that it is the foreseeable consequence of some choice the agent made.

Moreover, even if the belief alone does not constitute a pernicious pattern of attention (or diminishment, or taking a poor stance toward someone), it is a likely outcome thereof, making it a reasonable indicator of such patterns (especially for a partner who is worried that they might be accused of having fallen off the wagon). And, from the believer’s side, belief is a particularly difficult outcome to hide. In their description of the case, *Basu and Schroeder* (2018, p. 182) write that “you can see from her eyes that that she thinks you have fallen off of the wagon.” In this instance, it is not the belief itself, but the believer’s attention to the belief—the role it is playing in her psychology, such that it changes the way she looks at her partner—that causes the harm. Knowing how easy it will be for one’s partner to identify the belief, and how easily that knowledge might ruin their evening and lead them to feel betrayed, it is no wonder that the believer feels guilt, even if her belief is warranted.

7 Conclusions

One of the major motivations behind the epistemic morality literature is to understand senses in which our epistemic behavior might also be susceptible to moral criticism. Early forays into this literature, such as *Clifford (1877)*, assumed that the harm would come not from the belief itself, but from the actions inspired by the belief. But, more recent explorations have begun from the idea that believing a proposition about someone can be wrongful, apart from whether anyone acts on it. In this paper, I've tried to take that proposal seriously, and understand the conditions that need to be part of our understanding of doxastic wronging in order to make sense of it.

As I've argued, belief does not delineate the boundaries of doxastic wronging well. However, by expanding our focus to attentional epistemic wronging, so that we include both epistemic attitudes and the patterns of attention involved in forming them, we accomplish three things: First, because this added condition rules out cases like spontaneous belief-formation, we do a better job of tracing the contours of reasonable intuitions about wronging. Second, since conscious, subject-level attention can reasonably be understood as both agential and, when pointed toward epistemic goals, epistemic, we maintain the idea that these wrongs fall under the category of epistemic morality. Attention, in other words, is a plausible candidate for being governed by both moral and epistemic normativity. Finally, because patterns of attention concerning individuals and what's true of them are sustained engagements with those individuals, we provide better grounding for the kinds of explanations proponents of doxastic wronging offer for the harm of doxastic wronging.³³

³³Many thanks to audiences at Macalester College, the 13th Orange Beach Epistemology Workshop, the University of Michigan, and the 16th *Episteme* Conference. Special thanks, as well, Sarah Wright for extensive written comments on an early draft, and to Georgi Gardiner, Amy Flowerree, Will Fleischer, Rima Basu, Elise Woodard, Cameron McCulloch, and Laura Soter for helpful comments at many stages of writing.

References

- Allport, A. (1993), Attention and control. have we been asking the wrong questions? a critical review of twenty-five years, in *Attention and Performance XIV*, edited by D. E. Meyer and S. Kornblum, pp. 183–218, MIT Press.
- Alston, W. P. (1988), The deontological conception of epistemic justification, *Philosophical Perspectives*, 2, 257–299, doi:10.2307/2214077.
- Anderson, B. (2011), There is no such thing as attention, *Frontiers in Psychology*, 2, doi:10.3389/fpsyg.2011.00246.
- Basu, R. (2019a), The wrongs of racist beliefs, *Philosophical Studies*, 176(9), 2497–2515, doi:10.1007/s11098-018-1137-0.
- Basu, R. (2019b), Radical moral encroachment: The moral stakes of racist beliefs, *Philosophical Issues*, 29(1), 9–23, doi:10.1111/phis.12137.
- Basu, R. (2019c), What we epistemically owe to each other, *Philosophical Studies*, 176(4), 915–931, doi:10.1007/s11098-018-1219-z.
- Basu, R., and M. Schroeder (2018), Doxastic wronging, in *Pragmatic Encroachment in Epistemology*, edited by B. Kim and M. McGrath, pp. 181–205, Routledge.
- Begby, E. (2018), Doxastic morality, *Philosophical Topics*, 46(1), 155–172, doi:10.5840/philtopics20184619.
- Bennett, J. (1990), Why is belief involuntary?, *Analysis*, 50(2), 87–107, doi:10.2307/3328852.
- BonJour, L. (1985), *The Structure of Empirical Knowledge*, Harvard University Press.

- Clifford, W. K. (1877), The ethics of belief, in *The Ethics of Belief and Other Essays*, pp. 70–97, Amherst, New York: Prometheus Books.
- Dasgupta, N., and A. Greenwald (2001), On the malleability of automatic attitudes: Combating automatic prejudice with images of admired and disliked individuals, *Journal of personality and social psychology*, 81, 800–14, doi:10.1037//0022-3514.81.5.800.
- Dasgupta, N., and L. M. Rivera (2008), When social context matters: The influence of long-term contact and short-term exposure to admired out-group members on implicit attitudes and behavioral intentions, *Social Cognition*, 26(1), 112–123.
- De Houwer, J. (2011), Evaluative conditioning, *Associative learning and conditioning theory*, pp. 399–417.
- Enoch, D., and L. Spectre (forthcoming), There is no such thing as doxastic wrongdoing, *Philosophical Perspectives*.
- Fairweather, A., and C. Montemayor (2017), *Knowledge, Dexterity, and Attention: A Theory of Epistemic Agency*, New York: Cambridge University Press.
- Fischer, J. M., and M. Ravizza (1998), *Responsibility and Control: A Theory of Moral Responsibility*, Cambridge University Press.
- Gardiner, G. (2018), Evidentialism and moral encroachment, in *Believing in Accordance with the Evidence*, edited by M. Kevin, Springer Verlag.
- Gardiner, G. (2022), Attunement: On the cognitive virtues of attention, in *Social Virtue Epistemology*, edited by M. Alfano, J. D. Ridder, and C. Klein, Routledge.
- Gendler, T. S. (2011), On the epistemic costs of implicit bias, *Philosophical Studies*, 156(1), 33–63, doi:10.1007/s11098-011-9801-7.

- Healey, R. (2017), Interests, wrongs, and the injury hypothesis, *Journal of Ethics and Social Philosophy*, 12(1), 102–109, doi:10.26556/jesp.v12i1.215.
- Hieronymi, P. (2006), Controlling attitudes, *Pacific Philosophical Quarterly*, 87(1), 45–74, doi:10.1111/j.1468-0114.2006.00247.x.
- James, W. (1890), *The Principles of Psychology*, Dover Publications.
- Jorgensen, R. (2020), Varieties of moral encroachment, *Philosophical Perspectives*, 34(1), 5–26, doi:10.1111/phpe.12124.
- Kornblith, H. (1993), Epistemic normativity, *Synthese*, 94(3), 357–376, doi:10.1007/bf01064485.
- Kruglanski, A. (2013), *The Psychology of Closed Mindedness*, Essays in Social Psychology, Taylor & Francis.
- Mackie, M.-A., N. T. Van Dam, and J. Fan (2013), Cognitive control and attentional functions, *Brain and cognition*, 82(3), 301–312.
- Mole, C. (2021), Attention, in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta, Winter 2021 ed., Metaphysics Research Lab, Stanford University.
- Munton, J. (2021), Prejudice as the misattribution of salience, *Analytic Philosophy*.
- Prinz, J. (2011), Is attention necessary and sufficient for consciousness?, in *Attention: Philosophical and Psychological Essays*, edited by C. Mole, D. Smithies, and W. Wu, pp. 174–204, Oxford University Press.
- Radomsky, A. S., et al. (2014), Part 1—you can run but you can't hide: Intrusive thoughts on six continents, *Journal of Obsessive-Compulsive and Related Disorders*, 3, 269–279.
- Saint-Croix, C. (2022a), What is epistemic purism?, Manuscript.
- Saint-Croix, C. (2022b), Doxastic wrongdoing is out of control, Manuscript.

- Scanlon, T. (1998), *What We Owe to Each Other*, Belknap Press of Harvard University Press.
- Schmeichel, B. J. (2007), Attention control, memory updating, and emotion regulation temporarily reduce the capacity for executive control., *Journal of Experimental Psychology: General*, 136(2), 241.
- Schroeder, M. (2018), When beliefs wrong, *Philosophical Topics*, 46(1), 115–127, doi:10.5840/philtopics20184617.
- Siegel, S. (2017), *The Rationality of Perception*, Oxford University Press.
- Smith, L., and A. Archer (2020), Epistemic injustice and the attention economy, *Ethical Theory and Moral Practice*, 23(5), 777–795, doi:10.1007/s10677-020-10123-x.
- Spence, C. (2012), Attention, in *Encyclopedia of Human Behavior (Second Edition)*, edited by V. S. Ramachandran, second edition ed., p. 211–217, Academic Press, San Diego, doi:https://doi.org/10.1016/B978-0-12-375000-6.00037-9.
- Watzl, S. (2011a), The nature of attention, *Philosophy Compass*, 6(11), 842–853, doi:10.1111/j.1747-9991.2011.00433.x.
- Watzl, S. (2011b), Attention as structuring of the stream of consciousness, in *Attention: Philosophical and Psychological Essays*, edited by C. Mole, D. Smithies, and W. Wu, p. 145, Oxford University Press.
- Watzl, S. (2017), *Structuring Mind. The Nature of Attention and How It Shapes Consciousness*, Oxford, UK: Oxford University Press.
- Wegner, D. M., D. J. Schneider, S. R. Carter, and T. L. White (1987), Paradoxical effects of thought suppression., *Journal of Personality and Social Psychology*, 53(1), 5–13, doi:10.1037/0022-3514.53.1.5.
- Williams, B. (1970), Deciding to believe, in *Problems of the Self*, edited by B. Williams, pp. 136–51, Cambridge University Press.

Wu, W. (2011), Attention as selection for action, in *Attention: Philosophical and Psychological Essays*, edited by C. Mole, D. Smithies, and W. Wu, pp. 97–116, Oxford University Press.