

Rima Basu and Mark Schroeder
University of Southern California
May 21, 2018

Doxastic Wronging

In the Book of Common Prayer's Rite II version of the Eucharist, the congregation confesses, "we have sinned against you in thought, word, and deed".¹ According to this confession we wrong God not just by what we do and what we say, but also by what we think. The idea that we can wrong someone not just by what we do, but by what think or what we believe, is a natural one. It is the kind of wrong we feel when those we love believe the worst about us. And it is one of the salient wrongs of racism and sexism. Yet it is puzzling to many philosophers how we could wrong one another by virtue of what we believe about them. This paper defends the idea that we can morally wrong one another by what we believe about them from two such puzzles. The first puzzle concerns whether we have the right sort of control over our beliefs for them to be subject to moral evaluation. And the second concerns whether moral wrongs would come into conflict with the distinctively epistemic standards that govern belief. Our answer to both puzzles is that the distinctively epistemic standards governing belief are not independent of moral considerations. This account of moral encroachment explains how epistemic norms governing belief are sensitive to the moral requirements governing belief.

I.1 Doxastic Wronging, What

Our interest in this paper is in the idea of doxastic wrongdoing. A doxastic wrongdoing happens if one person wrongs another in virtue of what she believes about him. There are three parts of this definition that we want to emphasize. First, doxastic wrongs are directed. When you wrong someone, you don't merely do wrong, you do wrong to them. Second, doxastic wrongs are committed by beliefs. So in particular, the wrong in a doxastic wrongdoing does not lie in what you do, either prior to, or subsequent to, forming a belief, but rather in the belief itself. And third, doxastic wrongs are wrongs in virtue of

¹ Book of Common Prayer.

what is believed. So a belief that is a doxastic wronging does not wrong merely in virtue of its consequences; the wronging lies in the belief, rather than in, or at least, over and above, its effects.

It is not obvious to most philosophers that there are any cases of doxastic wronging; indeed whether this is even possible is, we take it, deeply controversial within philosophy, and many philosophers find it to be puzzling. Yet this puzzlement seems to come from theoretical reflection, rather than from ordinary thought. There are many plausible intuitive examples of doxastic wrongs. As we have already noted, one common formulation of the Christian Eucharistic confession appeals to the idea that we can sin against God in thought, as well as in word and in deed. This language is clear that the wrong is a directed one, and the explicit contrast with word and deed makes clear that the wronging lies in the thought, rather than in what we do. As with all doctrinal matters, of course, other interpretations are possible, and will be sought by those who cleave to the view that doxastic wrongs are too philosophically puzzling to be possible. Perhaps the confession does not imply that our beliefs ever wrong God, but only other matters of thought – what we spend time thinking about, for example, or our doubts, rather than our beliefs. But in the absence of some particular reason to think that beliefs cannot wrong, the confession is naturally interpreted as at least suggesting the existence of doxastic wrongs.²

Another place to look for doxastic wrongs lies in the feelings we are prone to have when our loved ones believe the worst of us. Just to make the thought vivid, 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. If you are like us, then you will be prone to feel wounded by this. Yes, you have a long history of falling off of the wagon, and yes, there is some evidence that this time is another. You can see how it could be reasonable for someone to draw this conclusion. But it still hurts – not least because in your eyes, tonight was an achievement to stay on the wagon despite adverse circumstances.

The feeling of being wounded is arguably a sign of a directed wrong. If she owes an apology, it is to you, not to anyone else – again the sign of a directed wrong. And it seems to be her belief that you would like an apology for. For example, it would feel insincere and unsatisfying if she apologized for the upstream act of not investigating more carefully before forming this belief, but continued to believe

² For further discussion, see Schroeder [unpublished].

it anyway, or if she apologized for the downstream act of revealing her belief to you by the look in her eyes, but not for the belief itself. Again, in the absence of some particular reason to think that beliefs cannot wrong, this looks to us like a strong prima facie case of a doxastic wrong.

A third, more pervasive place to look for doxastic wrongs lies in the wrongs of racism, sexism, homophobia, and other forms of prejudice. It is important to be clear that there are many harmful consequences of racism (just to focus on one example). But it is also natural to take it that one of the distinctive wrongs committed by a racist lies in what she believes about another human being.³ The racist is paradigmatically disposed to be influenced by her perceptions of race in the beliefs that she forms about another person – more easily persuaded that someone is dangerous, for example, if they are perceived as black. Racist beliefs are naturally taken not just to be morally problematic, but specifically to wrong their subjects. Moreover, an apology for crossing the street to avoid you by someone who still believes that you are dangerous would strike us as insincere, as would an apology for not checking more carefully before forming this belief, by someone who still holds the belief. So that is again at least some prima facie evidence that at least an important part of the wrong lies in the belief, rather than acts leading up to or subsequent to it. Once more, in the absence of any reason to think that doxastic wrongs are impossible, this looks to us like strong prima facie evidence of a doxastic wronging.

Much more work would be required in order to convince the dogmatic, but it is our belief that prima facie evidence like that assembled here should be taken at face value. All of these cases in which some wrong is committed, and neither apologies for the upstream acts which lead to the belief nor apologies for the downstream acts that are caused by that belief are intuitively satisfactory. This suggests that the real wrong lies somewhere else – neither upstream nor downstream, but in the belief itself, or in its formation. For these reasons, we hold that we do sometimes wrong one another in virtue of what we believe about them, and that an adequate understanding of the morality of interpersonal relationships must take this into account. But in this paper our purpose is not to offer a positive argument for this conclusion. Rather, it is to make room for this idea, by confronting what we take to be two of the most obvious reasons why the existence of doxastic wrongs might seem to be philosophically puzzling.

I.2 Two Puzzles

As we have been noting, the idea that there are doxastic wrongs is a natural one. Indeed it is, we believe, the position of ordinary common sense. But it has been held by many theorists to be philosophically

³This interpretation is argued for explicitly and directly in Basu [ms].

problematic. Indeed, this idea is so far from being an acceptable one in most mainstream circles in analytic philosophy that intelligent interlocutors often correct us, explaining that we can't really mean what we say we mean when we describe this thesis. There are important puzzles about how we could wrong one another in virtue of what we believe about them, and it is our goal in this paper to come to grips with two such puzzles: the problem of control and the problem of coordination.

The problem of control is the most familiar such problem. There are a number of ways of trying to press it, but they are all modeled on a familiar challenge to so-called deontological theories of epistemic justification. According to traditional deontological theories, epistemic justification is a matter of what we ought, or are permitted, to believe, or a matter of what is prescribed by certain epistemic rules. Critics object that beliefs cannot be subject to deontological categories like ought, permission, or rules, on the grounds that only behavior over which we have voluntary control can be subject to such categories, and we lack such voluntary control over our beliefs. The problem of control tries to generalize this sort of objection to the idea that there are doxastic wrongs. The idea is that we lack sufficient control over our beliefs to be accountable for them in the way that we would need to be, in order for our beliefs themselves to constitute ways of wronging someone. In contrast, actions are paradigmatically the sort of thing over which we have control, and so, the proponent of this problem alleges, it is much more plausible that we wrong others in virtue of the actions that we take leading up to the formation of some belief, or in the actions that we take on its basis. Hence, proponents of this argument conclude, the apparent wrongs associated with racism, sexism, and so on must all take place either before or after the belief is formed, and cannot be identified with the belief itself. We will respond to the problem of control in sections 2.1 and 2.2.

The problem of coordination, on the other hand, arises because what makes something wrong someone is influenced by moral considerations. But it is generally held that moral considerations are not the sort of consideration that rightly bears on the epistemic matter of what one should believe. And in general, according to popular evidentialist accounts of epistemic rationality, what it is epistemically rational to believe – that is, what is rational in every way that is required for knowledge – depends wholly on whether it is supported by adequate evidence.⁴ If moral considerations affect whether something is wrong without affecting the evidence, then nothing will coordinate moral and epistemic standards governing belief. In particular, there may be cases in which a belief wrongs someone even though it is rational in every way required for knowledge – in fact, even if it is knowledge.

⁴ Conee and Feldman [2004].

But this is puzzling, because in general, we do not wrong someone by doing something for which we have the right sort of adequate justification. If Daniel fails to keep his promise to meet you for lunch because he has been called to his father's deathbed, it may be wise and gracious for him to apologize to you, but he does not owe you an apology. He may harm you by frustrating your expectations for a lunch partner, but that harm is not a wronging. That his father's death was suddenly imminent is an adequate justification for him to break his promise. Similarly, if a belief is rational in every way required for knowledge – or even if it is knowledge – it seems like that ought to be an adequate justification for holding this belief, even if it hurts your feelings. It may be wise or gracious to apologize to each other for our beliefs, on this view, but we surely wouldn't owe apologies. And if this is right, then beliefs do not wrong, if they are epistemically rational. Hence, they cannot wrong at all. We will respond to the problem of coordination in sections 3.1 and 3.2.

Finally, we will conclude in section 4 by gesturing towards an explanation of how it could be true that the epistemic norms governing belief must be sensitive to the moral requirements governing belief. Our strategy will be familiar from the more general literature on pragmatic encroachment in epistemology. Pragmatic encroachment is the thesis that whether someone knows, or whether it is epistemically rational for her to believe, can depend on intuitively practical features of her situation, and not just on the evidence or other truth-related factors.

Because one of the main routes through which the thesis of pragmatic encroachment entered mainstream epistemology was through detailed discussion of intuitions about examples by Jeremy Fantl and Matthew McGrath [2002] and by Jason Stanley [2005], it is often taken to stand or fall on the basis of such intuitive judgments about cases. But Stanley does not claim that the intuitions that support pragmatic encroachment have a role akin to the role of observational data for scientific data. Rather, according to Stanley, these intuitions are intended to reveal in an intuitively compelling way a powerful connection between belief and action that epistemologists had been ignoring.⁵ Similarly, we suggest that cases of doxastic wronging are not so much just one more data point. Instead, they are important because they dramatize the role that our beliefs about one another play in constituting our interpersonal relationships – a powerful connection between belief and human sociality.

2.1 The Problem of Control

The first hurdle for establishing the plausibility of doxastic wronging is the intuitively appealing thought that we lack control over our beliefs – that our beliefs are insufficiently voluntary for

⁵ Stanley [2005, 12].

appropriate ascriptions of praise and blame, and correspondingly for rightness or wrongness, to apply. This problem of control has traditionally been used to establish that beliefs are not subject to distinctively deontological epistemic obligations. That is, there is no sense to be made of the claim that we ought believe or are permitted to believe anything as such terms suggest that we have control over what we believe. Our beliefs, it is argued, are at the mercy of the evidence.⁶ That is, there is no interesting sense in which it is up to us what we believe, we lack the necessary voluntary control. Call this the *original* problem of control.

In this section we will show how similar arguments can be extended to argue that beliefs cannot wrong. This version of the problem has seldom been laid out explicitly, in part this is because the thesis that beliefs can wrong has not been given the serious attention it deserves. It is natural to think, however, that the problem of control is at the heart of much of the skepticism towards doxastic wrongs. Call this the *extended* problem of control.

As we will show in this section, there are many places in which the original problem of control goes wrong. As a result, it shouldn't be surprising if those difficulties also extend to the extended problem of control. However, in section 2.2, we will see that this answer is too quick. In fact, more needs to be said to solve the extended problem of control than to solve the original problem of control. Nonetheless, we go on to say it by drawing on one of the main insights of this paper: that the epistemic standards governing belief are not independent of moral considerations.

The intuitive case for the challenge posed by the problem of control starts by reflecting on instances in which despite feeling wronged, it would be a mistake to say that we are wronged. For example, when your house is destroyed by a tornado, you may very well feel wronged by the tornado. Further, if the tornado destroys only your house, leaving all the other houses on your block untouched, it might even feel like a directed wrong. You might express your frustration at the tornado by turning to the sky and asking why it had to be your house or shaking your fists at the sky demanding an apology for what was done to you. We do not think, however, in such cases that you are wronged by the tornado. Similarly, despite how natural it is to feel wronged by your partner when they believe you've been drinking perhaps we are just mistaken.

Tornados have no control over whether they form nor what path they take. If there is a thunderstorm with an unstable mix of warm moist air, cool dry air, a change of wind direction, and increase in wind speed, there will be a tornado. Further, once the tornado has formed, there is no

⁶ We borrow this apt expression from Hatcher [2017], with whom we join in rejecting the assumption. As we understand this metaphor, the sense in which beliefs are supposed to be at the mercy of the evidence is that what you believe depends simply on what evidence you have been given, and is not the result of choice or agency.

interesting sense in which it is up-to-the-tornado whether to turn south or to turn west. Proponents of the problem of control argue that beliefs are, in a way, just like tornados. Just as tornados are at the mercy of instability in the atmosphere and wind shear, our beliefs are similarly at the mercy of the evidence. For example, if you look outside and see drops of water falling from the sky and making the sidewalk wet, you will believe that it is raining outside. On this picture, beliefs are not exercises of agency at all, voluntary or otherwise – they are merely caused in a predictable way by the evidence with which we are presented. Further, just as it is a mistake to think that we are wronged by tornados and other things outside of our control, it is similarly a mistake to think that we are wronged by what others believe of us.

This idea that beliefs are at the mercy of the evidence in the same way tornados are at the mercy of the elements is a widespread sentiment. Consider, for example, the following selection of quotes.

In saying that his belief is based on particular evidence, we would mean not just that he has the belief and can defend it with the evidence, but that he has the belief just because he has the evidence. This says that if he ceased to believe the evidence, then, other things being equal, he would cease to have that belief.⁷

With respect to almost all normal perceptual, introspective, and memory propositions, it is absurd to think that one has any such control over whether one accepts, rejects, or withholds the proposition. When I look out my window and see rain falling, water dripping off the leaves of trees, and cars passing by, I no more have immediate control over whether I accept those propositions than I have basic control.⁸

Suppose I am at the grocery store deliberating whether to get lemon sorbet or rocky road ice cream. Suppose I go with lemon sorbet. That was my choice, and thus the decision is mine. It is “up to me.” By contrast, suppose I am looking out my window at the sky, and it is raining. I deliberate whether to believe it is raining. I cannot bring myself to believe that it is not raining. It is not up to me what I believe; rather it is up to the evidence. I can’t go against definitive evidence that stares me in the face. It is not “up to me” in this way. Belief is not “up to me,” but is rationally determined.⁹

Implicit in these quotes is the following thought: for deontological concepts such as obligation or duty to apply in the domain of epistemology, beliefs would have to be different from tornados in the following way. We must be able to choose what to believe, agents need to have control over their beliefs to be responsible for what they believe. We (typically) are not responsible for what is not on our control. Proponents of this original problem of control argue that given that we are no more in

⁷ Williams [1970] p. 141.

⁸ Alston [1988], p. 270.

⁹ Flowerree [2016], p. 6.

control of our beliefs than a tornado is in control of what path it takes, it is just as odd to say that an agent ought believe p as it is to say that a tornado ought take the southeasterly path.

To extend this argument to establish the conclusion that beliefs cannot wrong we can add the following premise that emerges from the tornado analogy: for moral concepts, e.g., wronging, to apply to what an agent believes, agents need to have control over their beliefs. That is, for it to be the case that I can wrong you by what I believe of you, the belief must be in my control. My beliefs, to continue to mix metaphors, cannot be like tornados. It must be up to me whether my belief can turn south or turn west if I am to wrong you given what I believe. In short, we lack sufficient control over our beliefs to be accountable for them in the way that we would need to be in order for our beliefs themselves to constitute ways of wronging someone.

To briefly summarize, our beliefs are psychologically and conceptually limited in a way that makes them more like tornados than something to which deontological or moral concepts apply. With this in hand, we can formulate the two versions of the problem of control we've been discussing as follows.

Original Problem: The Argument Against Deontological Epistemic Obligations

- (1) For deontological concepts such as obligation or duty to apply in the domain of epistemology, agents need to have control over their beliefs.
- (2) It is not the case that agents have control over their beliefs.
- (C) Therefore, deontological concepts do not apply in the domain of epistemology.

Extended Problem: The Argument Against Doxastic Wrongs

- (1) For beliefs to wrong, agents need to have control over their beliefs.
- (2) It is not the case that agents have control over their beliefs.
- (C) Therefore, beliefs cannot wrong.

In the original problem, the two premises – (1) and (2) – have both come under attack for various reasons. Our goal here is not to come down one way or another on which way, in this space of possibilities, is the right way of responding to both versions of this argument. Instead, our goal is to give a sense of the wide space available for objecting to both versions of this argument and convey the point that the argument goes wrong in many places. Further, given that there's more than one place to object to the original problem, it shouldn't be surprising if those objections also apply to the extended problem.

Let us begin with the premise that is most commonly attacked: (2). Shah [2002] argues that a closer examination the connection between deontological concepts and belief reveals an alternative conception of control that is sufficient for deontological concepts to apply in the domain of

epistemology. Namely, we have a capacity to be moved by an appreciation of reasons. As a result, Shah rejects (2) on the grounds that we exert control over our beliefs in the following way: through our appreciation of the evidence we are agents with regard to our beliefs and are capable of regulating our beliefs. Steup [2012], in a manner similar to Shah, argues that arguments of the form given above assume that voluntariness requires responsiveness to practical reasons, but our beliefs are responsive to epistemic reasons and that is the way in which we exercise control over our beliefs. Weatherston [2008] has persuasively argued that this argument relies on too narrow a conception of voluntary. Flowerree [2016] argues that we must be agents with respect to our beliefs if we are to avoid the following reductio that plagues accounts that argue we are not agents with respect to our beliefs: if we are not agents with respect to our beliefs, then we are not agents with respect to our intentions. And Benbaji [2016] argues something even stronger: the argument regarding our lack of control generalizes to show that we're not even agents with respect to our actions.

Resisting (2), however, is not the only strategy. Some philosophers reject the assumption underlying (1) in the original problem: that for deontological concepts to apply in the domain of epistemology control is needed. For example, Hieronymi [2006], [2008] argues that it is not troubling to say that we can no more intend at will than believe at will. Many of our intentions unreliably produce actions, and the unreliability is highly suggestive of a lack of control. For example, consider Janae's intention to drink scotch. Janae would like to one day be the kind of person that can order scotch at a bar and drink it as easily as wine or beer. She did not start out as someone that liked the taste of wine or beer, but over time she has brought herself to like the taste. There are, however, many items she has failed to bring herself to like. No matter how much she would like to be the kind of person that likes goat cheese, for example, she finds that she cannot stand the taste. It seems that Janae lacks any reliable control over her intentions. Nonetheless, it is not odd to grant that given Janae desires to like goat cheese and drink scotch there is some sense in which Janae ought to like goat cheese and ought to drink scotch, further it is some sort of failing when she does not succeed. Despite the fact that she lacks any reliable kind of control over her intentions, deontological concepts seem to apply. Hieronymi takes these sorts of cases to suggest that control is not central to our notions of responsibility, praise, and blame.¹⁰

In general, these strategies against the problem of control take the following form. They show that for deontological concepts to apply in the domain of epistemology either we have sufficient

¹⁰ McHugh [2012] similarly argues that voluntariness is not a central condition, but rejects Hieronymi's criterion—answerability—in favour of reasons-responsiveness and thus also rejects (1).

control for deontological concepts to apply or there is no control needed for deontological concepts to apply. These strategies all concern different ways of understanding “control” and “sufficiency” and the relationship between “sufficient control” and the applicability of deontological concepts to a domain. Now we must ask whether these arguments might extend to respond to the extended problem of control. Unfortunately, as we will now see, the extended problem of control is more forceful than the original problem. But, we will argue that it can still be answered successfully.

2.2 The Problem of Control, Reconsidered

Let us start by taking a closer look at one of the objections to the original problem of control in order to see why these objections do not automatically transfer to the extended problem of control. Hieronymi [2006] grants that we can exercise a manipulative form of control over our beliefs. For example, we can have extrinsic reasons to produce within ourselves a belief.

Suppose you can't fall asleep because you are worried about whether your friends arrived safely home through the storm. Wanting sleep, you have an extrinsic reason for the belief that they are safely home. This extrinsic reason gives you reason to produce in yourself a belief. The obvious thing to do, in this case, is to conduct a little investigation: call your friends. If you find them home, you will have brought yourself to the desired belief.¹¹

These extrinsic reasons you have, reasons for why the belief is good to have, e.g., you'll be able to fall asleep if you believe your friend arrived safely home through the storm, only support you taking the steps to induce a belief. These reasons, Hieronymi argues, do not bear on the question of whether the belief is true, and as a result cannot support any particular conclusion you might otherwise hope you'd arrive at.

According to Hieronymi's answer to the original problem of control, the reason why being at the mercy of the evidence is compatible with epistemic agency is that evidence is the right kind of reason – a constitutive reason – for forming beliefs. So in being led around by the evidence, in being at the mercy of evidence, we are responding to reasons that are paradigmatic of epistemic agency. This objection to the original problem of control does not automatically transfer to the extended problem of control. The latter problem concerns whether moral reasons could be the right kind of reason for or against forming beliefs.

Nevertheless, the structure of Hieronymi's response offers a promising route to an answer to the extended problem of control. What we must show is that evidence is not the only right kind of reason for or against forming beliefs. Once we can show that, we will have opened the door for moral reasons

¹¹ Hieronymi [2006, 54]

being among the reasons that are active in epistemic agency and have cleared a theoretical hurdle for establishing the possibility of doxastic wrongs.

It seems wrong to us to say that only evidence matters for epistemic agency. On the face of it a lot of things matter. We can find support for this suspicion in an example from Nelson [2010, 87].

Given the appearance of some distinctive dark, winged shapes, moving across my visual field, what should I believe? That visual evidence, joined with other factors, may license me to believe propositions such as:

1. There are things moving through the air in front of me.
2. There are birds flying in front of me.
3. There are jackdaws flying in front of me.
4. At least three jackdaws exist.

Of these, which proposition I do believe will depend on, among other things: how my perceptual abilities have developed (e.g., have I learned to discriminate different kinds of bird on the wing?); the background information I happen to have (e.g. do I know what a jackdaw is?); and my particular interests at that moment (e.g. what do I want to know or do now?).”¹²

When we ask which proposition we should believe, the answer depends partly on our epistemic situation – in Hieronymi’s terms, the constitutive reasons, i.e., the evidence – but also partly on our needs and interests. For example, if Charlie is interested in launching a model airplane, then Charlie should believe 1. But if Zenobia is afraid of birds and wishes to avoid them, then she should believe 2. If D’angelo is conducting a species survey, then depending on the specificity of the survey he should believe either 3 or 4. However, if Diana’s interest is to hail a cab, she need not believe either 1-4. The lesson here is that none of Charlie, Zenobia, D’angelo, or Diana are simply being pulled around by the evidence. Given the same evidence – distinctive dark, winged shapes, moving across their visual field – they are capable of either forming a belief or not. Evidence, then, is not the only thing matters for epistemic agency.

This argument from Nelson shows that we have no trouble not believing things for which we have evidential support. This argument also provides support for thinking that we do respond to a wide variety of considerations that are not just evidence. This picture, however, could be made consistent with the account offered by Hieronymi. Hieronymi offers a two-stage view that can offer a competing diagnosis of each of these cases. For example, although neither Charlie, Zenobia, D’angelo, nor Diana are pulled around by the evidence, how they decide what to believe is settled by what question they ask. When it comes to believing, there are two separate activities answerable to two separate sorts of reasons. First, there is the activity of asking a question, of beginning a course of inquiry. Second,

¹²Nelson [2010, 87].

there are the constitutive reasons, i.e., the evidence that is unearthed during the investigation that settles which particular belief we should hold. Hieronymi argues that Charlie, Zenobia, D'angelo and Diana all exercise managerial control over their beliefs *through* their direct control over the first activity – the question they ask – but with regard to the second activity, their belief, insofar as they are rational, is entirely a matter of the evidence. That is, in the second stage of inquiry they are at the mercy of the evidence.

We can now see how this two-stage account prevents Hieronymi's objection to the original problem of control argument from automatically transferring to the extended problem of control. For deontological epistemic obligations it is enough that we have managerial control over our beliefs, i.e., control in the first stage of deliberation. For beliefs to wrong, however, it is not enough that we have managerial control over our beliefs. Recall that we are not upset with our spouse because they began an investigation into whether we had been drinking. Although that is another reason we might seek an apology, if that is all that our spouse apologized for her apology would be incomplete. What we seek apology for is what our spouse believes. We seek apology for the second activity, the activity over which Hieronymi argues it is not up to us how we settle the investigation once we begin it. We seem to seek apology for the activity over which Hieronymi argues we lack control. Generally it is thought to be inappropriate, however, to feel wronged by what someone did if what they did was not in their control. This lack of control seems like it could be excusing. If this lack of control is excusing, then we must look elsewhere for an explanation of why we feel wronged by what our spouse believes of us. That is, we need an alternative explanation for what we have been arguing we need doxastic wrongs to explain. That is, we need an explanation for why we might feel wronged by what others believe of us that does not require the controversial assumption that agents have control over their beliefs as the latter seems needed to establish doxastic wrongs.

The two-stage view offers such a debunking argument against doxastic wrongs. This debunking argument starts with a plausible alternative to doxastic wrongs: the observation that in contrast to beliefs, our actions are paradigmatically the sort of thing over which we have control. Given the strong intuitive sway of the idea that we do not control our beliefs, in the way needed for doxastic wrongs, perhaps then it is much more plausible that we wrong others in virtue of the actions that we take leading up to the formation of some belief, or in the actions that we take on its basis. Perhaps, then, the apparent wrongs associated with racism, sexism, and so on must all take place either before or after the belief is formed, and cannot be identified with the belief itself.

There are many reasons to reject this line of reasoning,¹³ but as we have focused on in this paper, an apology that stops at the level of what your partner does or what they reveal about their beliefs feels incomplete. In the example we've been using to motivate doxastic wrongs – wronged that are directed, wrongs that are committed by beliefs, and wrongs that are wrongs in virtue of what is believed – the apology is not complete until your spouse apologizes for the belief. It is more commonplace than philosophers seem to suggest, given the theoretical pull of the problem of control, to exclaim, “You shouldn't have believed that of me.” Thus, in response to this debunking alternative, we suggest that this candidate alternative explanation fails to capture the incompleteness of our spouse's apology. Hence we can reject this alternative explanation and we have good grounds to look back at suggestions made by the two-stage view and figure out where it goes wrong.

To recap, according to the two-stage view, once we get to the second activity, once the investigation is underway, it is no longer up to us what we believe. That is, what we will believe at the conclusion of the investigation is entirely a matter of epistemic considerations, i.e., what the evidence suggests. To answer the two-stage view what we must do is show that we are not at the mercy of the evidence once we consider the evidence. That is, we must show that evidence is not the only kind of reason that's the right kind of reason for or against forming beliefs. As we will now go on to show, 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.

Consider the following example. Suppose that Neil is ambushed in the street with a bunch of evidence that Taylor Swift has a dog named 'Taylor Swift'. Managerial control over what he believes has been taken away from him, the question has been asked and the evidence has been presented. Neil, however, does not care at all about Taylor Swift, whether she has a dog or what that dog's name is. Does Neil have any control over what he believes now of Taylor Swift and her dog or is Neil required to believe that Taylor Swift has a dog named 'Taylor Swift'? If Neil just doesn't care, it seems odd to say that he is required to believe anything about the name of Taylor Swift's dog. The two-stage view, however, is committed to saying that Neil has no choice in what he believes once the evidence is presented to him. This strikes us as implausible and we think we're in good company in rejecting this consequence of the two-stage view.

Sherlock Holmes for example, as he is presented in Sir Arthur Conan Doyle's work *A Study in Scarlet*, remarks at one point: “there comes a time when for every addition of knowledge you forget something that you knew before. It is of the highest importance, therefore, not have useless facts

¹³For further discussion see Basu [2018], especially chapter I.

elbowing out the useful ones.”¹⁴ Similarly, the version of Sherlock Holmes that appears in the latest BBC adaption by Stephen Moffat also finds himself chastised by Dr. Watson for being ignorant of Copernican Theory and unaware that the Earth travels around the sun. In *A Study in Scarlet*, when Holmes is similarly chastised in this way he responds, “What the deuce is it to me? [...] You say that we go round the sun. If we went round the moon it would not make a pennyworth of difference to me or to my work.” Along with the previous examples from Nelson of how evidence underdetermines what we believe, these examples together suggest that we really are capable of responding to a wide variety of considerations that are not just evidence, even once questions have been opened for us. These examples push back against the thought that the problem of control is really a problem. Evidence is not the only thing that matters for epistemic agency, a lot of things matter, and presumably some of those things could be moral reasons. These examples illustrate that we have more control over what we believe than traditionally thought and 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.

Recent work in social psychology further reinforces the picture that we have been sketching concerning belief, apology, interpersonal commitments, and the way in which our interests shape our beliefs. Kruglanski’s [2014] work shows that when we choose to settle belief is intimately tied to what he calls our epistemic motivations. The epistemic motivations are goals that we possess with respect to knowledge.¹⁵ He classifies these motivations along two dimensions: closure seeking versus avoidance, and specificity versus nonspecificity. As a result, there are four epistemic motivations explored in Kruglanski’s experimental work:

- (1) Need for nonspecific closure: you desire a firm answer to a question, but any answer will do
- (2) Need to avoid a nonspecific closure: you desire to keep your options open, as a result you are engaging in judgmental noncommitment
- (3) Need for specific closure: you are looking for a particular answer to a particular question and it is not the case that you would be happy with any answer to the question you are seeking an answer to, your motivations are directionally biased
- (4) Need to avoid specific closure: you wish to avoid believing a particular belief

The two-stage view can make sense at most of cases involving the need for nonspecific closure and need to avoid nonspecific closure. This is because in the two-stage view, belief formation decomposes into

¹⁴ Sir Arthur Conan Doyle, *A Study in Scarlet*.

¹⁵ See Nagel [2008] for more on applying Kruglanski’s work to issues in epistemology, e.g., our understanding of stakes and knowledge ascriptions.

two separate stages, and all practical considerations matter for is whether to go in for forming a belief at all. They do not matter for which belief you get.

In contrast, there is no way to choose at which stage the need for and need to avoid specific closure could operate. Take some simple cases of the need to avoid specific closure. Tobacco companies have expended considerable money on research investigating whether smoking causes lung cancer. They've avoided coming to the conclusion that it does not by refusing to consider the question at all, but by considering it asymmetrically – applying higher evidential standards to the answer that inconveniences them. The need to avoid specific closure works in just this way at an individual level – allowing even someone who is engaged in considering a question to apply differing standards to each answer.¹⁶

The way that Kruglanski and his colleagues understand this picture, one's motivation toward closure lies on a spectrum that can be manipulated by a variety of factors, e.g., your need for nonspecific closure can be heightened when you are working towards a deadline. Alternatively you might be tired and given that seeking particular answers to particular questions is taxing, you may instead exhibit more of a need to avoid nonspecific closure. These studies suggest that our belief formation practices are more sensitive to non-evidential considerations and it is unclear whether the two-stage view's limited conception of epistemic agency can accommodate these various epistemic motivations.

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 and we've cleared the first theoretical hurdle for establishing the suggestive idea that moral considerations could matter directly for the epistemic rationality of belief.¹⁷

3.1 The Problem of Coordination, First Pass

So far, we have given our response to the problem of control. The other obstacle to doxastic wrongs arises from the possibility that this leads to a lack of coordination between moral and epistemic standards governing beliefs. The problem is simple: if beliefs can wrong, then there will be non-coordinating epistemic and moral standards governing belief, because while the epistemic standards

¹⁶ Compare Schroeder [2012], [2013].

¹⁷ We have not exhausted the ways in which this could be established. For example, it is consistent with the idea that we are responsive to epistemic reasons, and the epistemic reasons are not exhausted by the evidence. See, for example, Schroeder [2012].

governing belief depend only on the evidence, whether a belief wrongs depends instead on moral considerations. And this could easily lead to troubling choices between being rational and being good.

Many authors who are sympathetic to the possibility of doxastic wrongs have described the moral standards governing belief as in conflict with the demands of epistemic rationality. For example, Tamar Gendler [2011, 57] claims that avoiding the doxastic wrongs of racist beliefs requires “explicit irrationality through base-rate neglect”. And according to Simon Keller [2004, 329], “a good friendship sometimes requires epistemic irresponsibility”.¹⁸ At a first pass, both of these appear to be the claim that there are genuine conflicts of epistemic and moral duties regarding beliefs – cases in which one ought morally to be one way with respect to one’s beliefs, but ought epistemically to be a contrary way.

The most natural interpretation of the problem of coordination is that a lack of coordination between moral and epistemic norms governing beliefs is bad, precisely because it would result in such deontic conflicts between epistemic and moral requirements on belief. We reject this interpretation of the problem of coordination, but it is worth taking the time to see why. This first pass articulation of the problem looks like this:

Non-coordination as Conflicts:

- PI) If there are doxastic wrongs, then there are deontic conflicts between requirements of morality and rational requirements on belief.
- P2) If there are deontic conflicts between requirements of morality and rational requirements on belief, then that creates some problem for rational requirements on belief.
- P3) There is no problem for rational requirements on belief.
- C) There are no doxastic wrongs.

We’ll leave open, here, whether there is a way of making good on an appropriate answer to what would be so bad about deontic conflicts between morality and epistemic rationality that is not bad about deontic conflicts between, for example, prudence and courage. The main reason why this does not give us a general problem for the idea that there are deontic wrongs is that PI is false – there could easily be deontic wrongs without any deontic conflicts between moral and epistemic requirements on belief.

If there is to be a deontic conflict between the requirements of morality and the rational requirements on belief, then there must be something, A, which some person X, ought morally to do but ought rationally not to do. So since we are concerned with beliefs, one of the following must be the case, for some proposition P and agent X:

¹⁸ Compare Stroud [2006].

Case 1: X is morally required to believe that P, and X is rationally required not to believe that P.
Case 2: X is morally required not to believe that P, and X is rationally required to believe that P.

The problem for PI that there can be deontic wrongs without either of these cases ever being possible. The cases of doxastic wrongs that we used to motivate our initial interest in this paper are cases in which one person wrongs another person by what she *believes* about him – not cases in which someone wrongs another person by what she does *not* believe about him.¹⁹ If there are doxastic wrongs of belief, but no doxastic wrongs of non-belief, then there are *ipso facto* no case 1 deontic conflicts.

Similarly, in order for there to be case 2 conflicts, there must be some things that we are epistemically required to believe. But it is not obvious that there are enough things that we are epistemically rationally required to believe, in order to create such conflicts. As we have noted earlier in section 2.2, Mark Nelson [2010] has argued, for example, that there are no positive epistemic duties at all – nothing that anyone is ever epistemically rationally epistemically required to believe at all, only things that it is rationally epistemically irrational to believe. And even if Nelson is wrong and there are some things that we are rationally epistemically required to believe, it is quite plausible that these are going to be insufficiently common, in order to generate conflicts of the required sort. For example, if our only positive epistemic duties are to believe logical truths, then all it takes for there to be doxastic wrongs without deontic conflicts with the requirements of epistemic rationality, is that believing a logical truth never wrongs anyone. But none of our leading examples of doxastic wrongs are cases of believing logical truths, and it is in any case not plausible that you can wrong someone by believing a logical truth. Similar points go if the beliefs that are rationally required include non-logical truths as well, so long as they are still relatively limited.

If any of the views in the last paragraph are correct, then at most few beliefs are epistemically required. You are doing poorly, as a believer, if you do not believe these things. And many are epistemically forbidden – you are doing poorly, as a believer, if you believe these things. Finally, some are morally forbidden – you are doing poorly, as a moral agent, if you believe these things. But the goals of doing well as a believer and doing well as a moral agent are in no way in tension, because as the following diagram illustrates, there are many ways of believing all of the things that you are rationally required to believe while believing none of the things that you are rationally forbidden from believing and believing none of the things that you are morally forbidden from believing.

¹⁹We don't mean to deny, here, that one can wrong someone by what you do not believe about them. The authors are divided over this question. We just mean to be arguing that there is a problem of conflicts even if we do deny this.




Fig. 1.

Consequently, we don't believe that a lack of coordination between epistemic and moral standards governing belief can be problematic simply in virtue of leading to deontic conflicts between the demands of epistemic rationality and the demands of morality. But we *do* believe that there is something more general to be troubled by, in the neighborhood. There is a more forceful way of saying why a lack of coordination between moral and epistemic standards governing belief would be bad.

3.2 The Problem of Coordination, Second Pass

The more forceful way of putting the problem of coordination does not depend on the assumption that there are any proper deontic conflicts between moral and rational epistemic requirements on belief. It targets the bare assumption that morality forbids things that epistemic rationality does not – the shaded area in figure 2:




Fig 2.

We believe that nothing can occupy this shaded area. Suppose for purposes of *reductio* that it does. If morality forbids beliefs that epistemic rationality does not, then there will be cases of belief that are rational in every way required for knowledge, but are nevertheless morally wrong. But we believe that there are no such cases. Indeed, it is puzzling how it *could* be wrong – and in particular, how it could wrong some individual – to believe something about them that satisfies every rational standard required for knowledge.

In order to draw out what is puzzling about this possibility, we will again focus on the characteristic role of apology, in relationship to directed wrongs. Apologies, it is important to

acknowledge, are a currency that helps to grease the mechanics of interpersonal relationships under a wide variety of circumstances, allowing reconciliation in the face of conflict. As such, it is often wise to apologize in many cases in which an apology is not, strictly speaking, owed, and successful interpersonal relationships may involve many such apologies. But we take it that apologies are owed, and therefore most at home, in the case of directed wrongs. If you wrong someone, you owe them an apology, and if you owe someone an apology – as opposed to if it is merely wise to give them an apology – then you must have wronged them in some way.

The close relationship between apology and directed wrongdoing is supported by consideration of the relationship between apology and excuse, as well as the relationship between apology and acting in a way that is all-things-considered wrong. Some actions wrong someone but are not wrong all-things-considered; this can happen because we encounter tradeoffs. For example, even if it is morally permissible to turn the trolley onto the side track, that does not show that turning the trolley does not wrong the single person on the side track; it may wrong her but nevertheless be justified all things considered. And strikingly, apology does still seem owed in this case.

Similarly, an action can both wrong someone and be wrong all-things-considered, but be excused. And we owe apologies for excused wrongs, as well. If you slip and spill red wine on someone's brand new white carpet for which they have been saving for months, you owe them an apology even if you were bumped by someone else. The fact that apologies are still owed in cases in which an act is permissible all-things-considered and in which though wrong, it is excused, supports our contention that owing an apology closely tracks having committed a directed wrong.

The problem for doxastic duties that extend beyond what is required by epistemic rationality is that apology does not seem owed, by someone whose belief meets every epistemic standard short of truth required for knowledge. To appreciate this point, it helps to imagine, as we have done before, the case in which you come home smelling of alcohol though you have not fallen off of the wagon. We contended before, in section I.I, that if your spouse apologizes only for what they did upstream from forming your belief or downstream after forming their belief, this does not feel like enough – there is something else left, unapologized-for, which is still a wrong.

Imagine, in this context, your spouse saying, 'well, I'm sorry for believing that you fell off the wagon, even though my belief was epistemically impeccable, short of being true'. This does not seem to us to be much of an apology. And the reason why it does not seem to be much of an apology, it seems to us, is that satisfying every epistemic standard appropriate to belief short of truth looks like the right kind of thing to defeat any presumption of being a wrong. It is therefore a bit like saying, 'I'm sorry for

taking back the book that you borrowed from me, even though I had every right to do it', which is again, we believe, not much of an apology.

If this is right, then there cannot be beliefs that are both rationally epistemically permissible and also constitute doxastic wrongs. That is, doxastic wrongs are all epistemically impermissible. As a result, our picture of the deontic space of beliefs must look more like the following:

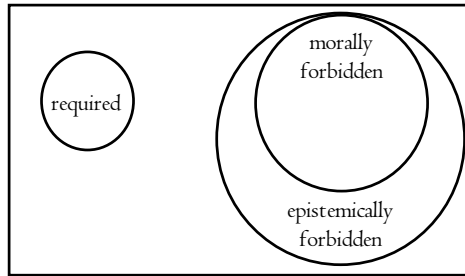


Fig 3.

This is a completely consistent picture. But unfortunately, given one more plausible background assumption that is commonly accepted in epistemology, this picture is not consistent with many of the best examples of apparent doxastic wrongs.

The problem is that if moral and epistemic considerations were completely independent of one another, then epistemic and moral impermissibility would not be expected to coordinate in the way required by figure 3; instead, the space of possibilities would be as depicted in the following diagram:²⁰




Fig 4

²⁰ In the diagram, we use the vertical dimension to represent evidential support for beliefs, with higher support toward the top of the diagram, and the horizontal dimension to represent the moral dimension of beliefs, with more morally problematic beliefs toward the right of the diagram. There are many different ways to make each of these dimensions more precise, and we do not intend our diagram to decide between them; merely to make a very rough picture of what is going on a little bit more visually tractable. Intuitively, the horizontal line that distinguishes epistemically rational from epistemically irrational beliefs is the threshold for what makes evidence 'sufficient' to justify belief.

If epistemic permissibility and moral permissibility are in fact coordinated, it follows that moral and epistemic considerations cannot be completely independent. So at least one of the following must be the case: either the moral requirements governing belief must be restricted in order to respect the epistemic requirements governing beliefs, or else the epistemic requirements governing belief must be expanded in order to respect the moral requirements governing beliefs.

Unfortunately, the first of these possibilities—restricting the moral requirements governing belief to respect the epistemic requirements governing belief—is not consistent with the full range of plausible and pre-theoretically compelling cases of doxastic wrongs. For example, Tamar Gendler [2011] offers the case of John Hope Franklin, an African-American member of an elite social club in Washington DC who is mistaken for staff by a fellow club member. In this case, the evidence that is possessed by the club member who forms the belief that John Hope Franklin is staff is quite good, in probabilistic, purely truth-directed terms, in comparison to many cases of ordinary beliefs that we form on an ordinary basis, because the club’s members are nearly all white, and its staff is nearly all black. This problem is that this belief appears to be high on the evidence scale. So in order for this belief to constitute a doxastic wrong, the standard of sufficiency for evidence must be very high across the board. But then it would follow that very many ordinary beliefs about the weather, the day’s news, and so on would all be epistemically irrational, because they would not meet the correspondingly high bar for evidence:

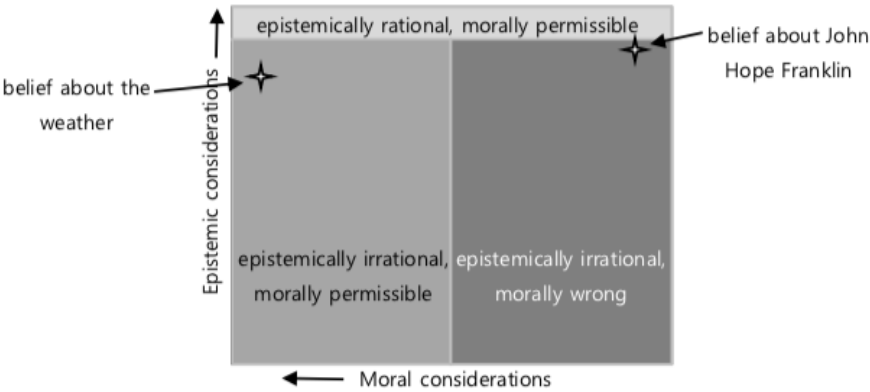


Fig 5

To avoid this consequence, we need a story of how the standard of sufficiency can be raised without also being raised across the board. This is what we will do in section 4.I.

In this section we have argued for two conclusions: first, that if there are doxastic wrongs, then there must be some coordination between the moral and epistemic norms governing belief, in order to rule out the possibility of morally wrong beliefs that meet every standard required for knowledge. And second, that the coordination between moral and epistemic standards governing belief cannot come

only by restricting morally wrong beliefs to cases of beliefs that are *independently* epistemically irrational. From these two conclusions we may infer that if there are doxastic wrongs, then the epistemic requirements governing belief must be sensitive to the moral requirements governing belief.

4.1 Moral Encroachment on Epistemic Rationality

In the last section we showed how the fact that epistemic rationality makes apology inappropriate generates a puzzle for the existence of doxastic wrongs – at least for the existence of doxastic wrongs in the kinds of plausible cases in which we believe that they are interesting and important, as opposed to cases in which they merely represent an additional layer of what is wrong with a belief that is independently recognized as epistemically irrational. If there are such doxastic wrongs, then in order to avoid the puzzle about inappropriate apology, the epistemic norms governing belief must be sensitive to the moral requirements governing belief.

This conclusion flies in the face of the epistemic orthodoxy of most of the last two and a half millennia. Distinctively epistemic rationality, it is commonly assumed – that is, the strongest kind of rationality that is required for knowledge – depends only on truth-related factors such as evidence, reliability, and subjunctive connections to the truth. In particular, according to core evidentialism, a view that we take very seriously indeed, what it takes for a belief to be epistemically rational is nothing other than that it is supported by sufficient evidence. So as long as the bar for sufficiency of evidence does not depend on moral factors, it follows that the epistemic requirements governing belief cannot be sensitive to moral considerations. So it is no wonder that evidentialists have not believed that we can wrong one another in virtue of what we believe about one another.

But we infer, in contrast, that the bar for sufficiency of evidence *does* depend on moral factors. As the moral considerations against belief increase, so does the evidence that is required in order to epistemically justify that belief. Our picture, in short, looks something like the following:




Fig 6

This picture is not inconsistent with core evidentialism. On our view, any belief is epistemically rational just in case it is supported by sufficient evidence. But the bar for sufficiency on evidence is sensitive to moral considerations. This reconciles the sets of morally and epistemically forbidden beliefs, so that there are no beliefs that wrong someone despite being epistemically impeccable. On this view, we cannot fully understand knowledge or purely epistemic rationality unless we come to grips with when and why beliefs can constitute wrongs.

This completes our defense of the possibility of doxastic wrongs from the hardest challenge behind the Problem of Coordination, and our argument for the deep connection between the possibility of doxastic wrongs and the need to make room for moral encroachment on purely epistemic rationality. But we'll consider one important objection, in closing. On our view, epistemically rational beliefs never constitute moral wrongs. But one might try to construct counterexamples to this principle. For example, on a natural view, the right to privacy is the right that others not know certain things about you, at least without permission. If you have a moral right that we not know certain things about you at least without permission, and you have not given us permission, and we know those things, then maybe it is precisely the fact that our belief has a very *positive* epistemic status – in particular, that it is knowledge, and hence epistemically rational – that makes it wrong you.

We agree that it follows from our account that the right to privacy cannot consist in or entail a right that others not know things about you without your permission. But this, we think, is an acceptable result. As natural as it might seem to describe the right to privacy in terms of a right against knowledge, in fact we believe the right to privacy is much more expansive. If it were simply a right against knowledge, it could be protected by Gettierizing anyone who believes something deeply personal about you. For example, if there are lots of fake nude celebrity photos on the internet, it would not

interfere with someone's privacy to publish real stolen nude photos of them, because given the abundance of fake photos, no one could come to know anything about you on that basis. This seems deeply wrong to us. Even with invasions of privacy, the wrong is one of belief, not of rational belief, or of knowledge.

4.2 A Word in Defense of Moral Encroachment

If you take the answers we've provided to the problem of control and the problem of epistemic conflicts seriously, as we have argued that we ought, then the existence of doxastic wrongs gives us an argument for a kind of pragmatic encroachment on epistemic rationality, and correspondingly, we take it, although we will not pursue this further here, on knowledge.²¹ Some, of course, will take this to be merely another way of sharpening the problem of epistemic conflicts for the existence of doxastic wrongs. Some may also take this to merely be another way of sharpening the problem of control. That is only to be expected; one philosopher's modus ponens is another's modus tollens. But we believe that cases like ours – cases involving doxastic wrongs – are both the best cases for pragmatic encroachment on the rationality of belief, and a return to pre-philosophical common sense. Sometimes philosophy may serve as a corrective to common sense, of course, but philosophy is hard, and can also lead us astray.

As we have been noting, the idea that there are doxastic wrongs is a natural one. But has been held by many theorists to be philosophically problematic. Jason Stanley has noted that the intuitions supporting traditional pragmatic encroachment reveal a powerful connection between belief and action that epistemologists had been ignoring. Similarly, though this is not a thought that we have had the opportunity to develop here, we suggest that cases of doxastic wrongdoing are important because they dramatize the role that our beliefs about one another play in constituting our interpersonal relationships – a powerful connection between belief and human sociality. If this defense of ordinary thought stands at odds with theoretical commitments, all the worse for our theoretical commitments.

²¹ Others who have come to the same conclusion that there is moral encroachment on epistemic rationality or knowledge include Pace [2011], Fritz [2017], and Moss [2018].

References

- Alston, William [1988]. 'The Deontological Conception of Epistemic Justification.' *Philosophical Perspectives* 2: 257-299.
- Basu, Rima. [2018]. *Beliefs that Wrong*. PhD Thesis, University of Southern California.
- _____. [ms]. 'The Wrongs of Racist Beliefs.' Unpublished manuscript.
- Benbaji, Hagit. [2016]. 'What can we not do at will and why.' *Philosophical Studies* 173(7): 1941-1961.
- Book of Common Prayer. <https://www.bcponline.org/>.
- Conee, Earl, and Richard Feldman [2004]. *Evidentialism*. Oxford: Oxford University Press.
- Flower, Amy [2016]. 'Agency of Belief and Intention.' *Synthese* 194 (8): 2763-2784.
- Fritz, James [2017]. 'Pragmatic Encroachment and Moral Encroachment.' *Pacific Philosophical Quarterly* 98(s1): 643-661.
- Gendler, Tamar [2011]. 'On the Epistemic Costs of Implicit Bias.' *Philosophical Studies* 156(1): 33-63.
- Hatcher, Michael [2017]. *A Deontological Explanation of Accessibilism*. PhD Dissertation, USC, <https://search.proquest.com/openview/17ecbcf034eb521b866ea243da550826/1.pdf?pq-origsite=gscholar&cbl=18750&diss=y>.
- Hieronimi, Pamela [2006]. 'Controlling Attitudes.' *Pacific Philosophical Quarterly* 87(1): 45-74.
- _____. [2008]. 'Responsibility for Believing.' *Synthese* 161(3): 357-373.
- Keller, Simon [2004]. 'Friendship and Belief.' *Philosophical Papers* 33(3): 329-351.
- Kruglanski, Arie W. [2014]. *The Psychology of Closed Mindedness*. Psychology Press.
- McHugh, Conor [2012]. 'Epistemic Deontology and Voluntariness.' *Erkenntnis* 71(1): 65-94.
- Moss, Sarah [2018]. *Probabilistic Knowledge*. Oxford: Oxford University Press.
- Nagel, Jennifer. [2008]. 'Knowledge ascriptions and the psychological consequences of changing stakes.' *Australasian Journal of Philosophy* 86(2): 279-294.
- Nelson, Mark [2010]. 'We Have No Positive Epistemic Duties.' *Mind* 119(1): 83-102.
- Pace, Michael [2011]. 'The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will to Believe'.' *Noûs* 45(2): 239-268.
- Schroeder, Mark [unpublished]. 'Sins of Thought.' Unpublished paper, presented at the Third Annual Theistic Ethics Workshop, College of William and Mary, October 2017.
- _____. [2012]. 'Stakes, withholding, and pragmatic encroachment on knowledge.' *Philosophical Studies*, 160(2): 265-285.
- _____. [2013]. 'State-Given Reasons: Prevalent, if not Ubiquitous.' *Ethics* 124(1): 128-140.

- Shah, Nishi [2002]. 'Clearing Space for Doxastic Voluntarism.' *The Monist* 85(3): 436-445.
- Stanley, Jason [2005]. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Steup, Matthias [2012]. 'Belief Control and Intentionality.' *Synthese* 188(2): 145-163.
- Stroud, Sarah [2006]. 'Epistemic Partiality in Friendship.' *Ethics* 116(3): 498-524.
- Weatherson, Brian [2008]. 'Deontology and Descartes' Demon.' *Journal of Philosophy* 105(9): 540-569.
- Williams, Bernard [1970]. 'Deciding to Believe.' In *Language, Belief, and Metaphysics*, ed. Howard E. Kiefer and Milton K. Munitz, 95-111. Albany: SUNY Press, 1970.