

“Morally Bad” Beliefs: A Taxonomy

—draft—

Pamela Hieronymi
hieronymi@ucla.edu
April 22, 2026

My goals here are modest: I would like to lay out a few very different ways in which a belief, or a believing, may be morally wrong—or, to use an even vaguer locution from the literature, “morally bad”¹—and do a bit of ground-clearing by presenting some examples. Recent literature in the ethics of belief marks a distinction between “doxastic wrongdoing” and “moral encroachment,” but I feel we would be helped by some alternative structure, which may make the variety of types of moral or ethical concern more perspicuous.² I admit, though, that this task is fraught, both because the word “moral” is not philosophically well-behaved and because the terms “wrong” and “bad” are largely unhelpful. Each example I present will raise more questions than I will be able to answer. I will do the best I can. Pursuing these modest goals will unearth two surprises: first, it seems possible for a *belief* to “wrong” someone even though the *believer* does not wrong anyone.³ Second, it also seems possible for a believer to wrong someone by believing, despite being fully justified in believing. In

¹ Renée Jorgensen Bolinger’s helpful article “Varieties of Moral Encroachment” inspired me to use term “morally bad,” which is both helpfully and frustratingly vague. I will continue to put quotation marks around it as a reminder of just how vague it is. Renée Jorgensen Bolinger, “Varieties of Moral Encroachment,” *Philosophical Perspectives* 34, no. 1 (2020), <https://doi.org/https://doi.org/10.1111/phpe.12124>, <https://onlinelibrary.wiley.com/doi/abs/10.1111/phpe.12124>.

² In Rima Basu’s recent presentation, “the doctrine of moral encroachment is the view that moral considerations bear on the justification of belief.” Rima Basu, “A Tale of Two Doctrines: Moral Encroachment and Doxastic Wronging,” in *Applied Epistemology*, ed. Jennifer Lackey (Oxford University Press, 2021), 99. This raises the difficulty of identifying what is and is not a “moral consideration.” Sometimes “moral encroachment” is understood as the encroachment of “moral norms” onto “epistemic norms.” This raises the question of how to distinguish “epistemic” from “moral.”

Meanwhile, “doxastic wrongs” are wrongs “committed by beliefs” and are wrong “in virtue of the content of what is believed.” Basu, “A Tale of Two Doctrines: Moral Encroachment and Doxastic Wronging,” 108. I find it difficult to understand a belief, rather than believer, as the one who “commits” a wrong (though, below, I will consider how a belief can perpetuate a societal wrong without the believer doing or believing wrongly). I also think we will be helped by separating the significance of the content of a belief from the significance of the person’s reasons for believing.

³ I put “wrong” in quotation marks because, strictly speaking, I believe it must be agents (including believers), either individual or collective, that do wrong. (In the case I will discuss, it is society, as a collective, that wrongs.)

both cases, the believer is fully justified, but in the second, it seems to me, they nonetheless wrong someone. This second case thus seems a previously unidentified form of bad moral luck.⁴

THE QUALITY OF WILL: VIRTUES AND VICES

Let me start with cases in which it is natural to attach terms ascribing moral virtue or vice to someone's believing and, therein, to hold them morally responsible, in one sense, for their belief. We say, naturally enough, that someone is being uncharitable, generous, hasty, admirable, biased, prejudicial, spiteful, and the like, in believing what they believe. Perhaps I believe the worst of my rival, in a way that is petty.

When we apply these virtue terms, we are saying something about the persons' reasons for believing, about the reasons on which the belief is "based." In an extension of Scanlon's term "operative reason," I will call a person's reasons for believing that person's "operative justification."⁵ When I am petty in my believing, my operative justification gives too much weight to small matters.⁶ When someone believes hastily, they have drawn a conclusion when they ought to have taken more time to investigate or to think—their operative justification shows insufficient patience or care. A person's operative justification reveals features of a person's mind, what is sometimes called their

⁴ The classic case of moral luck can be found in Bernard Williams, "Moral luck," in *Moral luck* (Cambridge: Cambridge University Press, 1981). The cases discussed by Nagel in Thomas Nagel, "Moral luck," in *Mortal Questions* (Cambridge: Cambridge University Press, 1979). include some amount of negligence. The literature is large. Below I will note Julie Tannenbaum, "Mere moral failure," *Canadian Journal of Philosophy* 45, no. 1 (2015)..

⁵ See T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), ch 1. Note that "Operative reasons" need not be *good* reasons and, likewise, an operative justification may not in fact justify. Operative reasons are the considerations the person employs in answering a question, the answering of which explains their activity. By employing those considerations as their reasons, the person *treats* them as good reasons, but the person may be mistaken in doing so. I detail the role a person's reasons play in explaining (rather than justifying) their action or attitude in Pamela Hieronymi, "Reasons for Action," *Proceedings of the Aristotelian Society* 111 (2011)..

⁶ Perhaps I believe my neighbor should not be trusted with a certain task, because of some previous minor error at which I took offense. Drawing this conclusion for this reason is what earns the belief the characterization as petty.

“quality of will” or what I elsewhere called their “moral personality.”⁷ A belief can be “morally bad” by being based on a morally problematic operative justification, one that reveals some vice.⁸

Notice, further, if you see I am being petty, hasty, or uncharitable in believing what I believe, you may then come to distrust me on this account. If I have been generous in what I believe about you, giving you the benefit of the doubt, you might be grateful. If I refuse to believe what you tell me about some matter on which you enjoy expertise, you may be offended.⁹ Responding to my beliefs with these “reactive attitudes” of distrust, gratitude, or offense is one way—an admittedly minimal way—in which you hold me responsible for my believing.¹⁰ If you did not so respond, you might then be discounting the quality of my will, as you might a child or someone suffering from some serious mental impairment. By responding in the ordinary way, you treat me as a responsible adult, someone whose quality of will matters to other adults in the usual way. Elsewhere I have called this “responsibility-as-mattering*.”¹¹ Our beliefs can be morally bad, or morally good, by revealing a good or bad quality of will, one that matters to others in this way.

Thus far we have simply applied, to believing, observations familiar from the case of intentions and intentional action: A person’s action is petty when they have given too much weight to small matters; a person acts generously when they discount costs to themselves; they are disrespectful when they fail to take others’ interests appropriately into account. The vice, here, traces to the

⁷ The reasons one employs will cohere, imperfectly, with the rest of one’s mind. Thus, one’s operative justification will reveal, imperfectly, facts about one. I discuss this at length in Pamela Hieronymi, “Responsibility for Believing,” *Synthese* 161, no. 3 (April 2008), <https://doi.org/10.1007/s11229-006-9089-x>, <http://www.springerlink.com/content/tu33736539117270/>.

⁸ “Vice,” as I am using it here, need not be grand or serious, nor need it be an enduring trait of character. It is possible for someone who is generally generous to form a petty belief in a special context and thereby display vice.

⁹ In fact, you need not enjoy expertise, if it is an ordinary matter.

¹⁰ See Peter F. Strawson, “Freedom and Resentment,” *Proceedings of the British Academy* xlviii (1962), which is reprinted in Pamela Hieronymi, *Freedom, Resentment, and the Metaphysics of Morals* (Princeton: Princeton University Press, 2020).

¹¹ The * stands for “to others in the way that the quality of adults’ wills typically matter to others.” It is meant to remind us that those who’s wills do not matter* may still matter in other ways. Pamela Hieronymi, “Minds Matter,” in *Gifford Lectures* (University of Glasgow, 2024).

operative justification of the intention. (The intention and action share the same operative justification. When all goes well, an intention is executed in action.) A person's operative justification for their intentional action reveals their quality of will, to which we might respond with distrust, gratitude, or offense. When we so respond, we treat the person as a responsible adult, someone whose quality of will matters to us in the usual way.

It is worth noticing a complication that will become important later: The reasons in the operative justification of an action might include considerations that do not “directly” count in favor of or against acting—that do not, themselves, show something good or bad about so acting. To illustrate: It may be that I would not have decided to attend the concert unless you had cancelled our dinner plans (that is, I would not have cancelled on you in order to attend the concert, but, given that you had to cancel, I decided to go). The fact that you cancelled our plans is among my reasons for deciding to attend, and it plays a role in explaining my attending. It will thus appear in the operative justification of my decision to attend the concert. But your canceling our dinner date does not, itself, show anything good about attending the concert. It does not, in *that* way (“directly”), count in favor of attending. Rather, it removes what would have been a barrier to attending. (It undercuts an obstacle.¹²) Even so, by appearing in my operative justification, it reveals something of the quality of my will (in this case, that I care about you or that I take my dinner dates seriously). We will return to this.

Philosophers are sometimes skeptical about whether we can be responsible for believing in the way we are responsible for acting, because our beliefs are not voluntary in the way that our actions are voluntary: we do not enjoy the same kind of discretion over our beliefs that we enjoy over our intentional actions. I have argued elsewhere that such skepticism is misplaced.¹³ It is true that we do not enjoy discretion over our beliefs, but that is because we do not enjoy discretion over what we

¹² For a helpfully formal way of representing distinct relations between reasons see John F. Horty, *Reasons as Defaults* (New York: Oxford University Press, 2012)..

¹³ Hieronymi, "Responsibility for Believing."

find reason-giving or how we weigh reasons—which is to say, we do not enjoy discretion over our operative justifications—and this is just as true when we intend as when we believe. In fact, though we enjoy discretion over our *actions*, we enjoy no more discretion over our *intentions* than we do over our beliefs. But we are surely responsible for our intentions.¹⁴ I argue elsewhere that the voluntariness or discretion characteristic of action is not required for responsibility-as-mattering*. More, responsibility-as-mattering* is the most fundamental form of responsibility.¹⁵

PERMISSIBILITY, OBLIGATIONS, AND BELIEF MANAGEMENT

Shifting topics. In addition to caring about the quality of one another's will—about how we figure into each other's world or how we regard one another—we also have more concrete concerns about the overt actions of others and how they impact our material well being.¹⁶ In fact, we live in a complex framework of expectations and obligations concerning how we affect and impact others' interests and well-being (a framework that includes “deontological” concerns about fairness and justice). These two sets of concerns—about the quality of others' wills and about their overt actions and how they impact us—are related, but the relation is interestingly complex.¹⁷

Without exploring that complexity, we can, for now, simply note that believing can sometimes create foreseeable harms that you ought to take steps to avoid. Failing to take those steps may be

¹⁴ If we are responsible for our actions, we must also be responsible for our intentions. See Hieronymi, "Responsibility for Believing."; Pamela Hieronymi, "Reflection and Responsibility," *Philosophy and Public Affairs* 42, no. 1 (2014)..

¹⁵ Hieronymi, "Responsibility for Believing."; Hieronymi, "Minds Matter."

¹⁶ This way of understanding moral concerns—as arranged around these to distinguishable foci—appears in T. M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Belknap Press, 2008).. I have found it extremely useful.

¹⁷ Most obviously, our overt intentional actions manifest the quality of our will. More, the standards for a good quality of will are set, in part, by our need to live together cooperatively while sharing a world of limited material resources. P. F. Strawson seems to suggest that it is by responding to the quality of others' wills with reactive attitudes such as distrust or offense that we construct, so to speak, the framework of expectations and demands. For this interpretation of Strawson, see Hieronymi, *Freedom, Resentment, and the Metaphysics of Morals*; Pamela Hieronymi, "Strawson's Ethical Naturalism: a Defence," (Leiden, The Netherlands: Brill, 2025).. T. M. Scanlon provides a different account of the relation in Scanlon, *Moral Dimensions*.. I find it extremely useful to note both the difference between and the interaction of these two concerns.

morally wrong and may result in a belief that seems to be, itself, morally bad. Considered in this way, a belief is no different than anything else—an allergic reaction or exhaustion, say—that would foreseeably create harm that you could take steps to avoid.

To illustrate, suppose you are part of a committee awarding a prize and fairness requires that the applicants remain anonymous.¹⁸ You then find yourself in a conversation in which it is obvious that the identity of an applicant is about to be revealed. You ought to excuse yourself from the conversation (or else change its course to avoid learning the identity of an applicant). But suppose you do not. Perhaps you think you are impervious to bias and will be better able to exercise your judgment if you know the identity. Or perhaps you are just curious. In any case, you come to know the identity of the applicant. You now have a belief you ought not to have, one that creates some unfairness. I am happy to call that a “morally bad” belief.

For another example, we have certain expectations about privacy, and violating those expectations is wrong. If, during my party, you walk into my bedroom and read my diary, you have wronged me. More, you now have beliefs you have no business having—about my ambitions, or my feelings about my children—and those beliefs violate my privacy. I am happy to call those “morally bad” beliefs.

In these cases, it seems plain that, even though the belief is “morally bad,” both the moral obligation and the moral vice attach to the action, so to speak, not the belief. In these cases, in which one has what we might call “obligations of belief-management,” the obligation gives you reason to act, but it does not give you reason to believe. Your obligation to avoid learning the identity of the candidate does not give you reason to doubt what you are told in the conversation. The speaker may be entirely reliable. Likewise, your obligation to respect my privacy does not give you reason to disbelieve what you have read in my diary. You have very good evidence for your

¹⁸ This claim about fairness is itself the consequence of complicated ethical reasoning that we are here taking for granted.

illicit beliefs about my feelings about my children, and your obligations do not bear on whether I am, in fact, disappointed in my children. The obligation, it seems, attaches to the action.

Relatedly, when you knowingly violate your obligations, in these cases, the vice appears in the operative justification of the action, not the operative justification of the belief. If you intentionally read my diary or decide to learn the identity of the applicant, you will have been disrespectful, nosey, or unscrupulous. These vices are manifest in your reasons for reading the diary or staying in the conversation. In contrast, the reasons on which your belief is based—your evidence for your belief—need not show these vices. In fact, the operative justification of your belief may not show any vice, or it may show different vices, such as gullibility. The vice attaches to the action.

Here is a rough comparison: If you allow yourself to become intoxicated in certain situations, the decision that led to the intoxication may show a lack of conscientiousness or misplaced self-confidence. But if, while drunk, you insult your friend's partner, that insult may reveal other vices, such as jealousy or meanness. Of course, the possibility of showing those latter vices may be among the reasons that recommended against drinking. Still, the vice manifest in your decision to continue drinking is not yet the jealousy or meanness manifest in the later comment.¹⁹ The earlier decision and the later comment are distinct activities, subject to distinct expectations and distinct evaluations. In a similar way, it seems, your decision to read my diary is distinct from the beliefs you form when you read it, and those distinct activities are subject to distinct expectations and can manifest distinct vices.

EVIDENCE, “PRAGMATIC ENCROACHMENT,” AND REASONS FOR BELIEVING

In these cases of belief-management, it is tempting to think that the moral obligation and vice “attach to” the action, and not to the belief, because the reasons for avoiding the belief—the obligations of privacy or fairness—are not evidence for or against the truth of what is believed, and, therefore, they must be the “wrong kind of reason” for believing, and, therefore, could only be

¹⁹ Disregarding the risk of being mean it is not yet being mean. It is being foolish or disrespectful.

reasons for self-management. Because the reasons for avoiding the belief are only reasons for self-management, the obligation and vice concern only the activity of self-management.

But this is not quite right: some considerations that are admittedly not *evidence* for or against *P* can nonetheless be among one's reasons for believing *P*.²⁰ For example, if I trust you and therefore take you at your word when you tell me that *P*, I am not treating your say-so as evidence that *P*.²¹ I would be treating your say-so as evidence if, instead, I fail to trust you but nonetheless think that, on this occasion, you are likely a reliable indicator of the truth (perhaps because of your access to information and the incentives at work). In that case, I would not be taking you at your word.²² In contrast, when I trust you, I do not take your say-so as evidence. Yet your say-so is certainly my reason for believing—it is the consideration I employ in answering for myself the question of whether *P*. Thus, the mere fact that a consideration is not taken to be evidence does not yet show that it is the wrong kind of reason for believing. (And so the mere fact that it is not taken to be evidence does not yet show that it could only be a reason for belief-management.)

Recall, from above, that your reasons for acting (the operative justification of your action) may include considerations that do not directly count in favor of acting, by showing the action good or bad, but that instead undercut reasons against acting. A further class of non-evidential reasons for

²⁰ Thus they can be the “right kind” of reason for believing *P*. The fact that some consideration, *C*, is not evidence that *P* does not yet establish that it is the wrong kind of reason for believing *P*. On the distinction between the “right” and “wrong kind” of reason of belief, see Pamela Hieronymi, “The Wrong Kind of Reason,” *The Journal of Philosophy* 102, no. 9 (September 2005), <http://www.jstor.org/stable/3655632>.

²¹ I owe both the point that testimony is not naturally thought of as a kind of evidence and that trusting someone is incompatible with taking their assertion as evidence to Richard Moran. See Richard Moran, “Getting Told and Being Believed,” *Philosopher's Imprint* 5, no. 5 (August 2005); Richard Moran, *The Exchange of Words: Speech, Testimony, and Intersubjectivity* (Oxford University Press, 2018), <https://doi.org/10.1093/oso/9780190873325.001.0001>, <https://doi.org/10.1093/oso/9780190873325.001.0001>. I puzzle about this in Pamela Hieronymi, “The Reasons of Trust,” *Australasian Journal of Philosophy* 86, no. 2 (June 2008), <http://www.informaworld.com/smpp/content%7Econtent=a792047025%7Edb=all%7Eorder=page>.

What is “evidence,” strictly speaking? It seems that evidence must stand in some sort of explanatory relation to the fact that *P*. Testimony, when not taken as evidence, does not. But this is an obviously inadequate answer.

²² Compare the attitude of a jury members toward a witness. Jurors are not to take the witness at their word, but rather to assess the likely veracity of their testimony.

believing seem analogous to this—they are considerations that are not evidence for *P* but instead bear on the role that other considerations play. One category of such reasons sometimes goes under the head “pragmatic encroachment” (and I will appropriate the term “encroachment,” using it more narrowly than it is sometimes used).²³

The classic presentation of “pragmatic encroachment” compares two cases.²⁴ In the first, you have some evidence that the bank will be open later this afternoon (perhaps it was open in the afternoon last month, or perhaps the hours are posted on their not-always-reliable website), and the cost of being wrong is low—perhaps, if it is closed, you will need to reroute to get to the ATM. You believe, on the basis of your evidence, that it will be open. In the second case, you have the same evidence, but in this case being wrong about whether the bank is open would cost you dearly—you need a cashier’s check to close the purchase of a home, say. In the second case, you will require more evidence before you believe.

In light of such cases, it seems that the foreseen badness of an incorrect belief can be a reason for requiring more evidence before you believe, while the relative unimportance of being incorrect can be a reason for being satisfied with the evidence at hand. If you believe the bank will be open because the website says so and I remind you that the website is often incorrect, you may point out that, if you are wrong, you can just use the ATM. It seems you are giving me your reasons for

²³ I will use “encroachment” for cases in which the foreseen risks or costs of being mistaken (or being correct) become one’s reasons for requiring (or not requiring) more evidence before believing. Some of the literature on “encroachment” is about knowledge. I am interested in a question about belief. Some of the literature about encroachment is about justification, about what it takes for a belief to be justified. I am interested in agency, in understanding the class of reasons on which one might *base* one’s belief, in contrast to reasons for which one might merely *manage* it. The question of agency will delineate one form of justification (as explained at Hieronymi, “The Wrong Kind of Reason,” 443.).

(I am here setting aside discussion of “higher-order” evidence, though it is obviously relevant.)

²⁴ The foundational paper is Keith DeRose, “Contextualism and Knowledge Attribution,” *Philosophy and Phenomenological Research* 52, no. 4 (1992).

continuing to believe that the bank is open.²⁵ The low cost of being incorrect is your reason for taking the evidence at hand to be sufficient. In contrast, if you need to get the cashier's check to close the purchase, you will require more evidence. If you go on to acquire that evidence and believe the bank will be open, the foreseen costs of being wrong will appear (not only in the operative justification of the action you took to acquire the extra evidence, but also, it seems to me) in the operative justification of your belief, as your reason for believing (only) on the basis of the expanded set of evidence. While the costs of being wrong are not evidence regarding the bank's operating hours, they nonetheless seem to bear (indirectly, we might say) on the question of whether the bank is open, in something like the way that your cancelling of our dinner plans does not directly count in favor of attending the concert but nonetheless bears on whether to attend.

It may seem puzzling or surprising that the costs of error could play this role—that they could be among one's reasons for believing—given that you do not take them to be evidence. I agree, but this case seems to me less puzzling than the case of believing a person you trust. It seems less puzzling because of the following feature of believing for reasons: when you believe the bank will be open, you not only employ certain considerations to answer the question of whether the bank will be open (that it was open last month, or that those hours are posted), but you also, in employing those reasons to reach that conclusion, *treat* those considerations as *sufficient* to answer that question.²⁶ That is to say, when answering the question of whether *P*, you inevitably, in so doing, make a determination about *how much* evidence is enough evidence to conclude that *P*. We might say

²⁵ One might think that you are, here, only *accepting* that the bank is open, not believing it. But, when you say “*if I am wrong*, I can just go to the ATM,” it seems you are reaffirming your belief. The alternative would seem to identify “belief” with certainty, which seems to me too strict a usage.

²⁶ “Taking” here is not believing. You do not form a higher-order representation of the considerations as sufficient. Rather, in coming to the conclusion, you treat them as sufficient. Likewise, “employing” a reason is not an attitude about the reason, but rather using the reason. Thus I understand “taking” differently than Boghossian. See, e.g., Paul Boghossian, “Inference, Agency, and Responsibility,” in *Reasoning: New Essays on Theoretical and Practical Thinking*, ed. Magdalena Balcerak Jackson and Brendan Balcerak Jackson (Oxford: Oxford University Press, 2019)..

that, in employing some set of considerations C to come to the conclusion that P , you *therein* set a threshold of evidence required to show that P .

Cases of encroachment make clear that setting that threshold is not only something you inevitably do (implicitly, we might say) whenever you answer a question, but also something for which you might have your own reasons. You might require more evidence, because being wrong would be costly, or you might not, because being wrong would not be that costly. Further, we might regard your reasons for setting the threshold where you do as good or poor, and thereby conclude that you are believing hastily, or being gullible, or being overly cautious.²⁷

Given that setting the threshold is the sort of thing one might do for reasons, we might be tempted to think that “threshold-setting” is either a voluntary action (the action of “setting”) or forming an additional belief (the belief that C is sufficient evidence to conclude that P). Crucially, it is neither.

If “setting the threshold” were a voluntary action, you would be able to do it for any reason you take to count sufficiently in favor of doing *it*—just as an able-bodied person can raise their right hand for any reason they take to count sufficiently in favor of doing so.²⁸ But it seems plain that we do not enjoy this kind of discretion in determining how much evidence is enough evidence—rather, we encounter limits. The limits are set by our own sense of what is, after all, sufficient evidence. For an easy example, suppose you read that Bond dies at the end of the movie. You might take your future enjoyment of the movie to count sufficiently in favor of requiring more evidence, yet, you might still take the evidence you now have to be sufficient. That is to say, you take yourself to have sufficient evidence that Bond dies, despite also thinking that you have sufficient reason to require more evidence, if only requiring more evidence were something you could do voluntarily. You are not then able to raise the threshold simply in order to enjoy the movie. For another kind of case,

²⁷ You might not have reasons for setting the threshold where you do. It may be that you set the threshold simply by employing the reasons at hand to answer the question.

²⁸ This is just how I would define “voluntary.” See Pamela Hieronymi, “Believing at Will,” *Canadian Journal of Philosophy, Supplementary Volume 35* (2009).

suppose your good friend is a serious tennis player about to compete in her most important event. You are a tennis aficionado, with strong and well-informed opinions. Having carefully considered the situation, you believe your friend will lose her event. However, you foresee that she would be hurt and disappointed to learn that you believe this, and you believe that friends have obligations to believe the best about their friends.²⁹ You have therefore taken considerably more care in arriving at your opinion than you would have were she not your friend. Even so, given the evidence, you might find yourself unable to raise the threshold of required evidence and thereby avoid the conclusion that she is likely to lose—though, if raising your right hand would have the same effect (if it would prevent you from concluding she is likely to lose), you could certainly do it.

The possibility of such cases shows that setting the threshold is not a voluntary action; it is not something we can do “at will.” We do not enjoy the same kind of discretion with respect to threshold-setting that we enjoy with respect to ordinary voluntary action. Rather, we encounter the same kind of limits that we seem to encounter when we believe something we find inconvenient to believe.

Given that setting the threshold is not a voluntary action, we might be tempted to think that setting the threshold is forming an additional *belief*—a belief about how much evidence is required to show that *P*. But, if we posit an additional belief, we would thereby start a regress worthy of the tortoise (this additional belief would, again involving setting a threshold of evidence for it, and, if

²⁹ For discussion of the epistemology of friendship see Simon Keller, "Friendship and Belief," *Philosophical Papers* 33, no. 3 (2004/11/01 2004), <https://doi.org/10.1080/05568640409485146>, <https://doi.org/10.1080/05568640409485146>; Sarah Stroud, "Epistemic Partiality in Friendship," *Ethics* 116, no. 3 (2006), <https://doi.org/10.1086/500337>, <http://www.jstor.org/stable/10.1086/500337>. and subsequent literature.

setting the threshold is forming a belief then setting the threshold for the additional belief would require a third belief, which would also require a fourth, etc.).³⁰

So, what shall we say about threshold setting? For now I suggest we simply say that setting the threshold is a matter of employing reasons or seeing reasons as sufficient—either you find the evidence at hand sufficient or you see the risks of being wrong as reason to require more evidence. To avoid the regress, I suggest we simply say that, *in* answering the question of whether *P* for reasons, one also, *therein*, treats those reasons as sufficient to conclude that *P*, and therein commits to a threshold of sufficiency, and, moreover, one sometimes has reasons for so treating them. But treating considerations as sufficient, in answering a question, is not forming a belief about their sufficiency. It is, simply, so treating them and, therein, committing to a threshold.³¹

With this much understanding of “threshold setting” and this understanding of “encroachment,” we can return to the question of why the “moral” obligation and “moral” vice attach to the action, rather than to the belief, in the cases of belief management (the prize and diary cases). The obligation attaches to the action insofar as it provides one with reasons for acting, and the vice attaches to the action insofar as it describes the operative justification of the action. But we cannot say that obligations of fairness or privacy are not reasons for believing simply because they are not evidence.

³⁰ Lewis Carroll, “What the Tortoise said to Achilles,” *Mind* 4, no. 14 (1895), <http://www.jstor.org/stable/2248015>. Notice that we would also run the risk of regress if we thought that *employing* some consideration, *C*, to answer positively the question of whether *P* requires a *belief* that *C* bears positively on that question. But we can and should deny that employing a consideration to answer a question is forming a belief about the consideration, the question, and the answer. Employing a consideration is, simply, employing it.

What is *that*? Using the consideration or basing the belief on the consideration. I am not getting very far in giving a positive characterization, but we can deny it is a belief without yet being able to give a more positive characterization. See note XX above.

³¹ This opens a can of worms. (What, one might wonder, is a belief if not a commitment to of this sort?) However, once released and organized, the worms may be helpful for a number of other puzzles, including practical knowledge (a commitment to how the world will be seems implicit in intention) and emotions.

Indeed, we can now note that, depending on the circumstances, someone *might* take the fact that it would be unfair to know the identity of the candidates as reason to require more evidence before being fully confident, on the basis of the evidence at hand, that Jones is the candidate. Someone might find concerns of fairness and anonymity to be reason to raise the threshold. However, it seems that they could not avoid believing that they have *some* evidence that Jones is the candidate, even if they think that fairness is compromised by even that belief. Likewise, it would be hard to think that obligations of privacy are reasons to require more evidence about my ambitions than was provided by reading my diary.

In general, the question of whether a given consideration bears on the question of whether *P* is what we might call a “substantive” question, to be answered in the details of the subject matter at hand. Whether the position of the stars bears on the likelihood of success in my job interview is not a question to be answered by thinking about the nature of reasons or reasoning, but rather about the relation between celestial bodies and employment decisions. I suspect that, likewise, whether a given consideration bears or is taken to bear on the question of whether *P* “indirectly,” as we have put it, by bearing on whether other considerations are sufficient to conclude that *P*, is also a

“substantive” question, to be answered in the details of the case.³² Still, we can say that *insofar* as the obligations of fairness or privacy provide reasons for action, rather than belief, the obligation and vice attach to the action, not the belief.

The fact that considerations other than evidence might be the right kind of reasons for believing seems to open up the possibility, not only for “pragmatic encroachment” but also for what has been called “moral encroachment.”

“MORAL ENCROACHMENT”

We have already granted that beliefs can be uncharitable, petty, or spiteful, due to their operative justification, and that beliefs can thereby be “morally bad.” But the fact that we often require more evidence when the costs of being wrong are high provides a more specific way in which a belief might be morally criticizable. In certain cases, other people would bear the costs of our error, and, in such cases, one ought to do unto others as one would do for oneself: When error would be costly for someone else, one ought to require more evidence. Clifford provides one illustration with his

³² In order to support the claim that we cannot intend “at will” I often provide a case in which your partner’s feelings are hurt because you have no intention of marrying them. You would be happy to avoid their hurt feelings by *intending* to marry them, so long as you do not have to actually marry them. Here, you cannot intend even though you take yourself to have sufficient reason to intend, because you do not take yourself to have sufficient reason to act. This case is like a case in which you cannot believe *P* (that everything will work out, say) even though you take yourself to have sufficient reason to believe (to get a good night’s sleep), because you do not take the reason to bear on whether *P*.

However, as I have elsewhere noted, your partner’s unhappiness is not, strictly speaking, a “reason of the wrong kind” for intending. Compare: a friend is unhappy that you have no intention to attend their party (both of you understand that it is your *intention* that matters—if you were prevented from attending, the friend’s feelings would not be hurt). You might decide that their unhappiness is reason enough to attend the party, thereby form the intention, and thereby resolve their unhappiness. The fact that an intention to *x* would please your friend can be reason to *x*. Likewise, the fact that an intention to marry would please your party can be some reason to marry—though you do not take it to be sufficient reason to marry. Your partner’s desires bear on the question of whether to marry, and so are not reasons of the wrong kind. (In contrast, the Toxin Puzzle provides a reason of the wrong kind for intending: the possibility of earning the prize does not bear on the question of whether to drink, because the time for earning is past before the time for drinking arrives.) See note 13 in Pamela Hieronymi, “I’ll bet you think this blame is about you,” in *Oxford Studies in Agency and Responsibility: Essays on Themes from the Work of Gary Watson*, ed. Justin Coates and Neal Tognazzini (Oxford: Oxford University Press, 2019).

Finally, the marriage case, in which you require more reason to act, is different than another case I have considered (raised by Mark Schroeder) in which you have reason to get a decision made that is *not* a reason for acting. See Mark Schroeder, “The Ubiquity of State-Given Reasons,” *Ethics* 122, no. 3 (2012), <http://www.jstor.org/stable/10.1086/664753>; Pamela Hieronymi, “The Use of Reasons in Thought (and the use of earmarks in arguments),” *Ethics* 124, no. 1 (October 2013): 122–5.

seminal example of the ship owner who convinces himself to overlook the possibility that the ship is not seaworthy,³³ but we can provide more familiar cases. What might be sufficient evidence for believing something innocuous about your colleague ought not to be sufficient evidence for believing something damning about them. If you believe something damning on the basis of evidence that would be sufficient for believing something innocuous, your operative justification might reveal that you are too quick to believe the worst. You have been uncharitable in your believing.

It seems to me that many of the cases of “moral encroachment” discussed in the literature can be understood as taking roughly this form: they are often cases in which the costs, to others, of your believing *P* are reasons for requiring more evidence before believing that *P*.³⁴ You ought to raise the threshold. Failing to do so is a “moral” failing, in that you have not taken others’ interests into account. The operative justification of your belief will reveal this. If your belief has poor consequences—if either the ship or your colleague’s reputation sinks—these consequences will be due, in part, to your failure to require more evidence.³⁵

However, as noted above, threshold raising is not a voluntary action; we encounter limits. At some point you might find the evidence sufficient, despite the harms your belief will foreseeably cause, and you might then find yourself unable to avoid believing, despite those foreseen harms. What to say at that point? Are you now guilty of a “morally bad” belief, in light of the harms? Are all beliefs that foreseeably cause others harm “morally bad” or vicious? Or, are all beliefs that foreseeably cause “moral” harm, whatever that may be, thereby “morally bad”?

³³ W. K. Clifford, “The Ethics of Belief,” in *The Ethics of Belief and other essays* (New York: Prometheus Books, 1999).

³⁴ I am now shifting from the costs of error to the costs of believing more generally.

³⁵ Perhaps ironically, Clifford’s insistence that “it is wrong, always, everywhere, and for anyone, to believe anything upon insufficient evidence” (5), which is taken to be the rallying cry of “evidentialism,” appeals explicitly to sufficiency, and his cases readily suggest that the threshold of sufficiency can be affected by the costs of error.

These are crude and complicated questions.³⁶ I will attempt to clear some brush.

BUMMER BELIEFS

First, not every case in which a belief would foreseeably impose costs on someone else is a case in which the belief is “morally bad.” In fact, not even every case in which a belief would foreseeably impose *undeserved* costs on someone else is a morally bad belief. Some costly beliefs, and some undeserved costs, are, we might say, just a bummer.

Return to your tennis-playing friend. Having carefully considered the situation, you believe your friend will lose her competition. You do not reveal this belief to her, but, through events beyond your control, she comes to know your opinion. Her feelings are, understandably, hurt. More, there is no sense in which she deserves to have her feelings hurt in this way—this is an undeserved hurt. Yet, I claim, it is possible that there is no moral wrong, or even moral badness, to be found in this case.

One might think that, while perhaps there is no *moral* wrong, there is at least a failure of friendship, here, and that failures of friendship are importantly analogous to moral wrongs—proto-moral wrongs, or something of the sort. And, one might think, if we shift the case into the moral register, the wrong will appear.

However, I think this case need not involve even any failure of friendship. We can, again, grant that, because she is a friend, you owe it to her to think about her match differently than a bookie would.³⁷ We can grant that you owe it to your friend, as her friend, not only to be especially careful in assessing the evidence and hesitant to conclude that she will lose, but also to give her the benefit of every doubt. We might even call these “epistemic duties of friendship.” Even so, it seems to me

³⁶ One thing to say: the fact that we encounter what I have here called “limits” when reasoning—the fact that what we find reason-giving is not voluntary—is no bar to it being morally significant or something for which you are responsible. See, again, Hieronymi, “Responsibility for Believing.”; Hieronymi, “Minds Matter.”

³⁷ See note XX above. I owe the evocation of the bookie to Berislav Marusic’s important challenge in Berislav Marušić, *Evidence and Agency: Norms of Belief for Promising and Resolving* (Oxford: Oxford University Press, 2015).

that, in a given case, you might satisfy all epistemic duties of friendship and nonetheless come to the conclusion that she will lose.³⁸ And in such a case, her feelings might still, reasonably, be hurt. Yet, despite her hurt feelings, you will have not wronged her—either morally or simply as a friend.

Certain situations are, we might say, just a bummer.

Notice that, in this bummer situation, you will want to offer your friend a kind of apology; you will likely say something that starts with, “I’m sorry.” (You might also insist that you are rooting for her and emphasize how delighted you would be for her to prove you wrong.) More, offering some kind of apology will be appropriate. If you did not offer an apology, your friend might justifiably wonder if you understand how she feels. But in offering this apology, you need not be acknowledging any moral wrongdoing or even a failure of friendship. You may, rather, be acknowledging her hurt feelings and your role in them.³⁹

A prominent case in the literature describes a spouse who takes the smell of alcohol on her partner’s clothing to be evidence that he has “fallen off the wagon.”⁴⁰ The case is given as a central example of a “doxastic wrong”—a case in which a belief morally wrongs someone. That this is a moral wrong is supported by the fact that the partner will feel hurt, an apology seems appropriate, and the apology is appropriately “directed” at the hurt partner. But these features are insufficient to secure the claim that the belief is a wrong. The case, as yet, seems to me under-described. On many

³⁸ In fact, it seems to me that you could have the same belief in the likelihood of her losing as the bookie. It seems to me that friendship requires that you come to the belief more carefully, but not that you reach a different conclusion. (In contrast, the challenge of Marušić’s cases is that they seem to require a different conclusion.)

Note: if, instead, you thought that you should conclude that she will lose only if it is *very* likely, or that you should have different opinion of the odds than a bookie, then you might think that the friend’s hurt feelings become inapt, because her odds are so low. That is, you might think your friend’s hurt feeling should calibrate to your duties. I do not think this is so. (In fact, I think that thinking it so will lead to considerable strife in one’s interpersonal relations, as you discredit hurt feelings that fail to calibrate with what you take to be your duties.)

³⁹ I am not yet saying an apology is “owed.” Perhaps one *owes* an apology only for wrongdoings. But an apology seems appropriate. One might compare what Williams says about the lorry driver.

⁴⁰ Rima Basu and Mark Schroeder, “Doxastic Wronging,” in *Pragmatic Encroachment in Epistemology*, ed. Brian Kim and Matthew McGrath (New York: Routledge, 2018).

ways of telling the story, it seems to me likely to be, not a case of “doxastic wronging,” but just a bummer, like the belief that the friend will lose.⁴¹

Thus I conclude (arguing by appeal to the tennis example) that not every case of hurt feelings in which an apology is appropriate is a case of moral wrongdoing or even a morally bad belief. More generally, not every case in which a belief foreseeably generates undeserved costs for others is a morally bad belief.

SOCIETAL INJUSTICE: THE CLASSROOM CASE

Returning to our other crude and complicated question: is every case in which a belief foreseeably causes a “moral harm” a “morally bad” belief? The idea of a “moral harm” is difficult, so let us clear a bit more brush by focusing on something more concrete: a case in which a belief will foreseeably perpetuate an injustice.

Let us start with a familiar case. Suppose that you are teaching a college class which includes certain students who are underprepared for the course due to societal injustices. Resources have been allocated to primary and secondary schools in ways that systematically disadvantage certain groups relative to others, along lines that create systematic, morally arbitrary, structural disadvantages.

Let us make some important stipulations. First, let us stipulate that fairness requires that you apply a single, uniform standard to the entire class. Your grade for the assignment is to reflect the quality of the work submitted, measured against that impersonal standard. To ensure fairness, you

⁴¹ Admittedly, I am just trading intuitions. The tennis case seems to me clearly not a “moral” wrong, and the spouse case, under certain tellings, seems relevantly like the tennis case—just a bummer. Explaining when and why it would clearly be a case of moral wrongdoing would be an interesting exercise. In the case as introduced, the partner “has a long history of falling off the wagon.” Does the believing spouse, in addition to believing, harbor shaming and blaming attitudes? Is she angry or disappointed? If she is disappointed, does the disappointment lead her to frustration or giving up hope, or was she expecting some disappointment to come her way, given the difficulties of sobriety? What if she takes herself to be partly responsible for having failed to accompany her partner to the event? I suspect that intuitions about wrongdoing may be driven by assumptions about these other attitudes, rather than belief. Thanks to Ben Jia for the thought that the spouse may feel some responsibility for the imagined lapse.

grade student work anonymously. Let us further stipulate that you take all due care in reaching your evaluation of each submission.

Suppose, then, that one of your disadvantaged students earns a C+ on their assignment. In such a case, your belief may be both accurate and also a vehicle of societal injustice. Unless you refuse to reveal it, the belief will perpetuate the injustices that explain the substandard quality of the submission. Yet we have stipulated that the belief is fully justified—or, if you prefer, fully “epistemically” justified. You have taken all due care in reaching it (the threshold was raised sufficiently high) and it is well supported by the considerations that bear on it.

Is your belief “morally wrong”? That again seems too crude a question to ask.

Your belief is a vehicle that will perpetuate societal injustice. In that sense, it is certainly “morally bad” and I am even happy to say it is, in *that* sense (as a vehicle of injustice) morally wrong.

Yet it also seems to me that you, as the instructor, have not done *anything* wrong, either in acting or in believing. That is to say, this seems to me a case in which the injustice in your situation does not, so to speak, transfer to or impinge upon your own activities as the instructor. One might want to say that, though the belief is morally bad (or even, in a sense, morally wrong), you are not believing wrongly (or, you do nothing wrong in believing). The belief, as part of the larger state of affairs, perpetuates injustice, but the belief, as your activity, shows no vice (either in the evidence for it nor in the threshold set by it) nor is it something you ought to have managed yourself so as to avoid.

One might resist in three ways. First, one might claim that the belief is wrong because it represents the student in a “belittling” way. I would reply that there is nothing “belittling” in thinking that a student wrote a C+ paper. There is no shame in writing a C+ paper, either because you are underprepared or even because you are such that you will not excel at writing college essays. This objection problematically moralizes academic ability and achievement. (I will return, below, to genuinely belittling or morally problematic representations.)

Second, one might resist by claiming that, insofar as the instructor's fully justified belief perpetuates the injustice, and insofar as the belief is the instructor's own (the instructor's own activity, I would say), the instructor is, so to speak, "taken up" into the injustice.⁴² They are, one might say, blamelessly implicated in injustice—where "blameless" simply marks the fact we have marked, namely, that they are fully justified (both epistemically and morally) in believing. They are, then, "believing wrongly" simply insofar as the belief contributes to a larger structure of injustice. Though I would prefer to use the words "believing wrongly" to mark activities done wrongly, and so would prefer not to use the terms in this way, I do not otherwise disagree.

But one might want to resist in another way. One might think that the instructor can and should avoid perpetuating the injustice by, e.g., not recording the grade that reflects their judgment.⁴³ One might further think that, so long as they do not express their belief, they will have done no wrong, but if they do express it, then they will not only be acting wrongly but also believing wrongly.

First, it is odd to suggest that a belief can become wrong simply because it is expressed or acted upon. If one thinks that it is wrong to reveal the belief, it is more natural, I think, to say that it is the action of revealing that is wrong, while the belief is not.⁴⁴ But, further, one cannot ensure that the belief fails to perpetuate injustice simply by failing to reveal it. In hiding their belief, the instructor is not *avoiding* the injustice, but rather attempting to redirect it, so to speak, onto themselves, as the one to devise a way to avoid perpetuating the injustice for the student. They may also be redirecting the injustice onto certain other students, who may be disadvantaged in less obvious ways but still unjust

⁴² I have said something similar about the case in which one is "taken up into" the project of improving one's child's self-esteem by responding to their achievements with pride. One's pride serves the social function of improving their self-esteem, despite the fact that improving their self-esteem is the wrong kind of reason for pride. Hieronymi, "I'll Bet," notes 26, 36. I argue that beliefs are to be understood as activities in Pamela Hieronymi, "No Inertia in Consciousness," in *Analytic Existentialism*, ed. Berislav Marusic and Mark Schroeder (Oxford: Oxford University Press, 2024).

⁴³ Or perhaps by designing the course so that there is some way of offsetting the grades that are due to under preparation—though that will not eliminate the hit to self-confidence in the underprepared student.

⁴⁴ Cp. above: the obligation and the vice attach to the activity done for the reasons relevant to the obligation or vice.

ways. It thus seems doubtful to me that the instructor can avoid perpetuating the unjust situation by believing. Even so, it seems that the instructor may not be doing anything morally wrong—either in believing or acting.

If we accept this, the upshot is surprising: It is possible for a belief to be morally bad, and even to “wrong” someone, even if the believer has not done anything morally wrong.⁴⁵

OFFENSIVE CONTENT AND MORAL LUCK

I have so far avoided considering racist, sexist, homophobic, or other prejudicial beliefs, in part because these are complex forms of “moral badness,” and I find the cases discussed in the literature challenging but largely under-described.⁴⁶ I hope that what I have said above gives some ways of

⁴⁵ The quotation marks again marking the fact that wrong, here, is done the level of society. A state of affairs is wrong because of an individual or collective action or lack thereof. (Cp. fn XX.)

Julie Tannenbaum has explored another surprising possibility: She argues that purely “epistemic” errors, that is errors that display no moral vices (a simply mathematical error, say), can sometimes result in moral wrongdoing by explaining a failure to succeed in something one has a duty to succeed in doing (rescuing someone, say). Tannenbaum, “**Mere moral failure.**”

⁴⁶ One thing I do feel confident about: it cannot be wrong to instantiate a universal generalization. If I know that every Black resident of a given town has an 80% chance of having an open arrest warrant, I do no wrong in drawing the conclusion that *this* Black resident of that town has an 80% chance of having an open arrest warrant. Indeed, if I am to understand the injustices faced by this particular Black resident, it seems important for me to believe this of them.

Whether I should be particularly careful about believing that this person *has* an open arrest warrant (whether I should be particularly careful in moving from a belief about probability to an all-out belief) will depend, in ways we have seen, on the costs, to this person and others, of this belief. Shifting to another example from the literature, believing a Black member of a social club is a staff person seems to me costly to them (and others), and so this is a belief for which one should require extremely strong evidence—one should raise the threshold maximally. (Though, to be clear, the belief that someone is a staff person need not, itself, be demeaning or belittling—when and why it is so is worth considering.) It is not clear to me whether believing that someone who lives under an openly corrupt policing regime has an open arrest warrant is costly to them. I suspect it depends on the details; I would like to hear from those that do. (Compare: Given the facts revealed by the MeToo movement, it seems that believing a woman has been sexually harassed is not costly to that woman. Eliminating that stigma—eliminating the suggestion that the woman is at fault for her own harassment—seems to be one of the benefits of the widespread belief about the pervasiveness of the problem. Of course, that benefit depends on other factors that may not be present in the policing case.)

To more securely locate racism in the open-arrest-warrant case, it seems to me that we could either look more deeply into my reasons (my operative justification) for believing—into how I understand the fact about open arrest warrants (do I believe it is a product of prejudicial policing and racial injustice or do I think people unlike me are dangerous?) or else look more broadly, to facts outside of me—facts about how patterns of arrest, and even beliefs about them, effect people’s lives in systematically unjust ways. Like the belief that the submission merits a C+, a belief about open arrest warrants might perpetuate injustice. Alternatively, like the belief that most women have been harassed, the belief might, at some point, help to mitigate the injustice by exposing it.

It would enrich the literature to consider the complexity of these cases in more detail. I do not feel I have the knowledge base to do so.

breaking down the complexity, as follows: It is one thing to consider the quality of someone's will—that is, the values revealed in the operative justification of some action or attitude. It is another to consider the costs or impacts of the fact that they believe. The two can part company. More, the foreseeable costs or impacts of a belief—including the hurt feelings that may result from the appreciation of one's operative justification—can be either reason for belief management or reason to raise the threshold of evidence required for believing. Some beliefs will perpetuate larger societal injustices due to the role they play in certain contexts—such as the classroom. Such beliefs can be morally bad—unjust—without being something one has an obligation to avoid and without that injustice appearing in the operative justification of the belief. In that case, the belief can be “morally bad” even though the believer is not believing badly.

However, in addition to all this, a belief is a representational state, with representational content, and that content is meaningful, and certain meanings are, themselves, morally problematic cultural items, so to speak. Consider, for example, a slur. A belief that contains a slur is, for that reason, a “morally bad” belief. It representing someone, or some group, in a belittling way, simply in virtue of the slur—independently of any further badness that might be found in either the reasons on which the belief is based or the effects or impact the belief might have.

Of course, in the typical case in which a belief contains a slur, lots of other “moral badness” is to be found—in the operative justification and quality of will and in the effects that the belief will have. But the slur itself contributes. It is possible for someone who is new to a culture and language to use a slur entirely innocently, not yet fully understanding the word.⁴⁷ It is also possible for them to have a belief that uses the slur innocently. (For the plausibility of the case, it is important that the slur occurs in a *belief*—a committal, assertive state, not simply a passing thought. And it is important that the slur is *used*, not just mentioned, in the belief.) Even if they never

⁴⁷ Cp. the discussion of “deference” started by Tyler Burge, “Individualism and the Mental,” *Midwest Studies In Philosophy* 4, no. 1 (1979/09/01 1979), <https://doi.org/https://doi.org/10.1111/j.1475-4975.1979.tb00374.x>, <https://doi.org/10.1111/j.1475-4975.1979.tb00374.x>.

express that belief, and so even if that belief does not have any further effects, it seems to me that they might rightly feel that, simply by believing, they have (inadvertently and blamelessly) participated in a wrong, and so done something that wronged those maligned by the slur.

Interestingly, though, an apology would typically be inappropriate in this case. Apologizing for the unexpressed belief that innocently employs a slur would typically simply create another occasion for the slur to do its harm, so to speak. Still, I can imagine special circumstances in which the belief comes to light and in which an apology for believing (not just for the revelation of the belief) would be apt—though, as with other cases of moral bad luck, it seems that the apology would also be appropriately deflected.⁴⁸

The case of slurs, or of other morally problematic content, deserves much more attention.⁴⁹ I raise it here simply to note a bit of structure that might be helpfully distinguished: in addition to the reasons for which a person believes (the quality of will or operative justification), and in addition to the consequences or impacts of the belief (about which one might have obligations or duties of belief-management or in light of which one might raise the threshold of evidence), a belief has representational content, some of which might be, in itself, “morally bad.”

CONCLUSION

In conclusion, the category “morally bad” is both helpfully and unhelpfully vague. I hope here to have provided a bit of taxonomy for thinking a bit more carefully about “morally bad” beliefs.

First, beliefs can be “morally bad” because the believer’s reasons for believing—the believer’s operative justification, the considerations on which the belief is based—display some moral vice, a

⁴⁸ Williams points out that the lorry driver will apologize, we will say that there is no need for him to do so, but we would be appalled if he took us at our word. DOUBLE CHECK. Williams, “Moral luck.”

⁴⁹ Some words become slurs only when used as such (“Sit down, woman!”).

Some of the beliefs discussed in the literature are tropes or stereotypes (“women are emotional creatures with poor judgment”). Such beliefs are often used to rationalize (indeed, to attempt to “morally” justify) oppressive ways of thinking about and living with others. In the central cases, it seems that these beliefs are not the source, so to speak, but rather are rationalizations, in the pejorative sense, of a deeper and broader problem.

poor “quality of will.” The person’s reasons for believing show that they are believing uncharitably, hastily, peevishly, or that they are biased, prejudiced, or whatnot.

Second, beliefs can also be “morally bad” because of their effects, costs, or impacts on others. Thus, a belief can be morally bad in the same way that a non-rational state—such as an allergic reaction or exhaustion—can be morally bad: it can have foreseeably unfair, offensive, deceptive, or otherwise “morally bad” effects, and therefore be something that the believer should take steps to manage or avoid.⁵⁰ You ought not to learn the identity of the candidate or read my diary.

These two ways of being “morally bad” interact in various interesting ways. Rather than sort through the complexity, I have just tried to do some initial ground clearing.

First, the foreseeable costs of a belief can be among the person’s reasons for requiring more evidence before believing. Failing to “raise the threshold” in light of those costs can result in a belief that is hasty or rash. Sometimes this becomes a kind of “moral encroachment,” resulting in a belief that is uncharitable, prejudicial, or likely to cause harm.

However, there are limits to such encroachment. Sometimes, at that limit, a belief that is foreseeably hurtful or harmful is nonetheless a belief that shows no vice, disrespect, or even violation of the expectations of a relationship. These are simply bumner beliefs born of bumner situations.

Other beliefs that have foreseeable consequences that are not just a bumner, but are unfair or unjust. But, again just clearing ground, I argued that sometimes a belief, such as the belief that the underprepared student’s paper merits a C+, can perpetuate injustice independently, so to speak, of the believer. The belief can perpetuate injustice even though the believer has no obligation to avoid the belief and may be fully justified in holding it. The believer may not be doing anything wrong, even in believing.

⁵⁰ Note: merely appealing to “effects, costs, or impacts” does not render this category a “consequentialist” one. Many duties of fairness depend on facts about effects, costs, or impacts.

Finally, in addition to revealing an operative justification or quality of will, and in addition to having consequences and impacts, beliefs are representational states that often have culturally meaningful representational contents, some of which are demeaning. It seems to me that a believer may wrong someone simply in virtue of believing something that is, in virtue of such content, demeaning. They may do so even though the believer's operative justification shows no flaws and the believer has not failed in any obligation of belief management. These are not just “bummer” beliefs, insofar as the person demeaned is wronged by the believer's believing. These beliefs are, it seems, a kind of bad moral luck.

In sum, then, beliefs can reveal a poor quality of will. They can have or create unjust or otherwise morally problematic consequences or burdens. These two features interact. When a belief creates burdens, those burdens sometimes generate obligations of belief management, and the beliefs typically also face stiffer standards of justification—one should raise the threshold of evidence. Even so, it is possible that even the unjust burdens created by a belief are not due to any injustice committed by the believer. Yet a third feature of belief makes it possible for a believer to innocently wrong someone else—the belief, as a representational state, may represent someone in a belittling or demeaning way, even if the believer does not understand this.

Or so it seems to me.⁵¹

⁵¹ Thanks are owed to many. As the embarrassing amount of self-citation shows, this work builds on previous work and thanks owed there are again owed here. I am especially grateful to Rima Basu and Mark Schroeder, as many of the thoughts here began in conversation with them long ago. I also owe thanks to my graduate seminar at UCLA in Winter of 2026.

BIBLIOGRAPHY

- Basu, Rima. "A Tale of Two Doctrines: Moral Encroachment and Doxastic Wronging." In *Applied Epistemology*, edited by Jennifer Lackey, 0: Oxford University Press, 2021.
- Basu, Rima, and Mark Schroeder. "Doxastic Wronging." In *Pragmatic Encroachment in Epistemology*, edited by Brian Kim and Matthew McGrath, 181–205. New York: Routledge, 2018.
- Boghossian, Paul. "Inference, Agency, and Responsibility." In *Reasoning: New Essays on Theoretical and Practical Thinking*, edited by Magdalena Balcerak Jackson and Brendan Balcerak Jackson, 101–24. Oxford: Oxford University Press, 2019.
- Bolinger, Renée Jorgensen. "Varieties of Moral Encroachment." *Philosophical Perspectives* 34, no. 1 (2020): 5–26. <https://doi.org/https://doi.org/10.1111/phpe.12124>. <https://onlinelibrary.wiley.com/doi/abs/10.1111/phpe.12124>.
- Burge, Tyler. "Individualism and the Mental." *Midwest Studies In Philosophy* 4, no. 1 (1979/09/01 1979): 73–121. <https://doi.org/https://doi.org/10.1111/j.1475-4975.1979.tb00374.x>. <https://doi.org/10.1111/j.1475-4975.1979.tb00374.x>.
- Carroll, Lewis. "What the Tortoise Said to Achilles." *Mind* 4, no. 14 (1895): 278–80. <http://www.jstor.org/stable/2248015>.
- Clifford, W. K. "The Ethics of Belief." In *The Ethics of Belief and Other Essays*, 70–96. New York: Prometheus Books, 1999.
- DeRose, Keith. "Contextualism and Knowledge Attribution." *Philosophy and Phenomenological Research* 52, no. 4 (1992): 913–29.
- Hieronimi, Pamela. "Believing at Will." *Canadian Journal of Philosophy, Supplementary Volume* 35 (2009): 149–87.
- . *Freedom, Resentment, and the Metaphysics of Morals*. Princeton: Princeton University Press, 2020.
- . "I'll Bet You Think This Blame Is About You." In *Oxford Studies in Agency and Responsibility: Essays on Themes from the Work of Gary Watson*, edited by Justin Coates and Neal Tognazzini, 60–87. Oxford: Oxford University Press, 2019.
- . "Minds Matter." In *Gifford Lectures*. University of Glasgow, 2024.
- . "No Inertia in Consciousness." In *Analytic Existentialism*, edited by Berislav Marusic and Mark Schroeder, 82–114. Oxford: Oxford University Press, 2024.
- . "Reasons for Action." *Proceedings of the Aristotelian Society* 111 (2011): 407–27.
- . "The Reasons of 'Trust.'" *Australasian Journal of Philosophy* 86, no. 2 (June 2008): 213–36. <http://www.informaworld.com/smpp/content%7Econtent=a792047025%7Edb=all%7Eorder=page>.
- . "Reflection and Responsibility." *Philosophy and Public Affairs* 42, no. 1 (2014): 3–41.
- . "Responsibility for Believing." *Synthese* 161, no. 3 (April 2008): 357–73. <https://doi.org/10.1007/s11229-006-9089-x>. <http://www.springerlink.com/content/tu33736539117270/>.
- . "Strawson's Ethical Naturalism: A Defence." 170–96. Leiden, The Netherlands: Brill, 2025.
- . "The Use of Reasons in Thought (and the Use of Earmarks in Arguments)." *Ethics* 124, no. 1 (October 2013): 114–27.

- . "The Wrong Kind of Reason." *The Journal of Philosophy* 102, no. 9 (September 2005): 1–21. <http://www.jstor.org/stable/3655632>.
- Horty, John F. *Reasons as Defaults*. New York: Oxford University Press, 2012.
- Keller, Simon. "Friendship and Belief." *Philosophical Papers* 33, no. 3 (2004/11/01 2004): 329–51. <https://doi.org/10.1080/05568640409485146>. <https://doi.org/10.1080/05568640409485146>.
- Marušić, Berislav. *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford: Oxford University Press, 2015.
- Moran, Richard. *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oxford University Press, 2018. <https://doi.org/10.1093/oso/9780190873325.001.0001>. <https://doi.org/10.1093/oso/9780190873325.001.0001>.
- . "Getting Told and Being Believed." *Philosopher's Imprint* 5, no. 5 (August 2005): 1–28.
- Nagel, Thomas. "Moral Luck." In *Mortal Questions*, 24–38. Cambridge: Cambridge University Press, 1979.
- Scanlon, T. M. *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press, 2008.
- . *What We Owe to Each Other*. Cambridge, MA: Harvard University Press, 1998.
- Schroeder, Mark. "The Ubiquity of State-Given Reasons." *Ethics* 122, no. 3 (2012): 457–88. <http://www.jstor.org/stable/10.1086/664753>.
- Strawson, Peter F. "Freedom and Resentment." *Proceedings of the British Academy* xlviii (1962): 1–25.
- Stroud, Sarah. "Epistemic Partiality in Friendship." *Ethics* 116, no. 3 (2006): 498–524. <https://doi.org/10.1086/500337>. <http://www.jstor.org/stable/10.1086/500337>.
- Tannenbaum, Julie. "Mere Moral Failure." *Canadian Journal of Philosophy* 45, no. 1 (2015): 58–84.
- Williams, Bernard. "Moral Luck." In *Moral Luck*, 20–39. Cambridge: Cambridge University Press, 1981.