

Moral Encroachment and Positive Profiling

(forthcoming in *Erkenntnis*. please cite published version)

Abstract Some claim that moral factors affect the epistemic status of our beliefs. Call this *the moral encroachment thesis*. It's been argued that the moral encroachment thesis can explain at least part of the wrongness of racial profiling. The thesis predicts that the high moral stakes in cases of racial profiling make it more difficult for these racist beliefs to be justified or to constitute knowledge. This paper considers a class of racial generalizations that seem to do just the opposite of this. The high moral stakes of the beliefs we infer from these generalizations make it *easier* rather than *harder* for these beliefs to be justified or to constitute knowledge. I argue that the existence of this class of cases—cases of “positive profiling”—give us reason to expand our account of moral encroachment in a way that brings it closer to the ideal of pragmatic encroachment that motivates it in the first place.

1 Introduction

Some claim that moral factors affect the epistemic status of our beliefs. Call this *the moral encroachment thesis*. It's been argued that the moral encroachment thesis can explain at least part of the wrongness of racial profiling. The thesis predicts that the high moral stakes in cases of racial profiling make it more difficult for these racist beliefs to be justified or to constitute knowledge. This paper considers a class of racial generalizations that seem to do just the opposite of this. The high moral stakes of the beliefs we infer from these generalizations make it *easier* rather than *harder* for these beliefs to be justified or to constitute knowledge. I argue that these cases provide some much needed insight into the structure of moral encroachment.

The first three sections of this paper briefly lay out some background. In §2, I describe and motivate the thesis that practical factors sometimes affect the epistemic status of our beliefs via pragmatic encroachment. In §3 and §4, I describe the parallel argument for moral encroachment and show how it leads to a puzzle, one that the literature already acknowledges. The moral encroachment thesis cannot *both* explain cases of racial profiling and *also* mimic the structure of pragmatic encroachment.

The remainder of this paper argues that we can resolve this puzzle by expanding our account of racial profiling in a way that we already have reason to do. In §5, I introduce a class of racial generalizations that give rise to what I call “positive profiling”. An example of this kind of generalization is the true statistic that *Black Americans are more likely to have their pain ignored and undertreated*. I argue that, on the basis of this generalization, we owe a duty of care to avoid the positive profiling error of failing to believe Black Americans when they claim to be in pain. Just as the harm of holding a belief about some individual, on the basis of their membership in a racial group, should sometimes lead us to raise the standards for justifiably holding this belief, the harm of *failing* to hold a belief about some individual, on the basis of their membership in a racial group, should sometimes lead us to *lower* the standards for justifiably holding this belief. We should be more positively inclined to believe that Black Americans are telling the truth when they claim to be in pain, on the basis of a statistic that tells us that their pain is so often ignored. In §6, I show that expanding the moral encroachment thesis so that it allows us to both lower and raise the standards for justified belief, in cases of racial profiling, brings it closer to the ideal of pragmatic encroachment that motivates it in the first place by defending a novel account of the harm of racial profiling.

2 From Classical Pragmatism to Pragmatic Encroachment

The observation that non-evidential factors seem to bear upon what we should believe isn’t a new idea. Consider the classic case of Pascal’s wager, which suggests that we have practical reason to believe in God. Given the possibility of theism, things will go better for us if we believe in God, or so the argument goes. And we should clearly perform those actions that make things go better for us.

The idea that our beliefs should be beholden to our practical interests, cast in this very crude way, is difficult for many to swallow. This view contrasts with *evidentialism*, according

to which what you should believe is a function only of your evidence.¹ More generally, this view contrasts with the idea that there is a clear boundary between the practical and the epistemic. For it suggests that all of our reasons, including our reasons for belief, are grounded in our practical interests.²

While Pascal's wager has largely fallen out of favor these days, there's a more palatable pragmatism in the neighborhood. Like Pascal's classical pragmatism, this different pragmatic view contrasts with evidentialism. It holds that not all of those factors that bear upon the justification of our beliefs must be evidential. But this more sophisticated view is able to preserve the distinction between the practical and the epistemic by maintaining that our practical interests don't directly provide reasons for our beliefs. Instead, our practical interests set the threshold for how much evidence is required for our beliefs to be justified or to constitute knowledge.

To see this idea in action, consider the BANK CASE from Stanley (2005, p.3-4), adapted from DeRose (1992), which is made up of the following two vignettes:

LOW STAKES. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it wasn't very important that their paychecks are deposited right away, Hannah says, 'I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning'.

HIGH STAKES. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, "I guess you're right. I don't know that the bank will be open tomorrow".

These cases differ only with respect to the practical stakes in question. In the second case, the consequences of being wrong about the bank being open are much worse than in

the first case. That the stakes are so high in the second case, but not in the first, suggests that we have justified belief in the first case, but not in the second.³ The higher risk of being wrong about the bank being open in the second case should raise the evidential threshold and make it harder for this belief to be justified:

HRHJ: The higher the risk of acting on p , the higher the evidential threshold and the harder it is to have a justified belief in p .

The combination of HIGH STAKES and LOW STAKES illustrates how what is at stake can change whether or not our beliefs are justified. In order to be justified in believing that the bank is open, when the consequences of being wrong are so dire, we need more evidence than we do when the consequences are relatively less severe. This idea supports the following general principle:

The Pragmatic Encroachment Thesis: The justification of a belief can depend upon its pragmatic features.

The pragmatic encroachment thesis implies that the evidence alone does not settle what we should believe. The practical stakes of holding a belief determine how much evidence is needed for this belief to be justified.

3 Moral Encroachment and Racial Profiling

Proponents of moral encroachment also hold that the evidence alone does not settle what we should believe. They endorse the following principle:

The Moral Encroachment Thesis: The justification of a belief can depend upon its moral features.⁴

As this principle suggests, it's tempting to think we can take any example of pragmatic encroachment and simply turn it into a case of moral encroachment by replacing the agent's practical interests with her moral interests. To borrow an example from Fritz (2017, p. 650), suppose there is a maniacal traffic officer who will kill five innocent people if you have the false belief that your car is legally parked. In such a case, it would be problematic for you not to check to make sure that your car is legally parked. Moreover, it would be problematic in the very same way that it would be problematic for you to believe that the bank is open on Saturday, on the basis of past experience alone, if there is an impending bill coming due.

The moral encroachment thesis is, then, made plausible by its similarity to the pragmatic encroachment thesis. But that's not all that makes it plausible. Many think that the moral encroachment thesis can be used to identify one sense in which racial profiling is epistemically problematic. Consider the following definition of racial profiling:

Racial Profiling: The practice of forming opinions about a person on the basis of statistics about members of their racial group.⁵

To begin to see how the moral encroachment thesis might be used to explain the badness of racial profiling, consider the following case, borrowed from Basu (2019a, p.10), who adapts it from Gendler (2011):

SOCIAL CLUB. Agnes and Esther are members of a swanky D.C. social club with a strict dress code of tuxedos for both male guests and staff members, and dresses for female guests and staff members. While preparing for their evening walk, the two women head toward the coat check to collect their coats. As they approach the coat check, they both look around for a staff member. As Agnes looks around she notices a well-dressed black man standing off to the side and tells Esther, "There's a staff member. We can give our coat check ticket to him."

If we assume, with Basu, that the club's discriminatory membership policy has resulted in a situation where only a small proportion of its members are Black, whereas all of its staff members are Black, then Agnes's belief at least *seems* to be epistemically justified. This statistical evidence provides reason for thinking that the Black man is indeed a staff member. Despite this, many think that the previous example should give us pause. There is *something* that goes wrong in SOCIAL CLUB. A significant literature has developed around the task of explaining what this something is.

On the one hand, we might think that Agnes's mistake is a purely epistemic one. Some have argued that statistical evidence is simply not the kind of evidence that is able to produce a justified belief.⁶ Others have suggested, along broadly similar lines, that what goes wrong in cases like SOCIAL CLUB has to do with the particular sort of statistical inference that is being drawn. Munton (2019) suggests that the problem with certain accurate statistical generalizations, like those that we appeal to when racial profiling, is that we use them to project a regularity observed among past instances of a kind to a novel, unobserved instance to which it does not apply. Still others have claimed that the tools of mainstream

epistemology can be used to explain the bad feature of SOCIAL CLUB. Gardiner (2020) suggests that one plausible interpretation of what goes wrong in such cases is that a relevant alternative—that the well-dressed Black man is a club member—has failed to be properly considered.⁷

On the other hand, some want to insist that it is the moral features of SOCIAL CLUB that undermine the justifiability of Agnes’s belief. Most notably perhaps, Basu (2019a) argues that more moral care is required when forming a belief on the basis of evidence that is the result of our racist history.⁸ She notes that in cases like SOCIAL CLUB, where Black people are mistaken for staff members, what goes wrong is that “their central self-descriptions, the descriptions from which they draw their sense of self-worth are deeply challenged” (Basu (2019b, p. 925)), and that such an error would not have been possible had it not been for the history of racism that produced the generalization in the first place:

Underrepresented groups are more often mistaken for employees because of the color of their skin and the racist intuitions that make their skin color a determining factor in their inability to gain access to more prestigious employment opportunities. Being mistaken in this context, namely one in which you’ve historically been excluded, is a greater harm and wrong than being mistaken in a space where that historical disadvantage is lacking. (Basu (2019a, p. 13))

These remarks suggest that the moral encroachment thesis can explain at least part of the wrongness of racial profiling in a way that draws upon its similarity with the pragmatic encroachment thesis. Plausibly, what happens in SOCIAL CLUB is that the high moral stakes that Basu describes raise the evidential threshold. Since Agnes’s evidence does not satisfy this higher evidential threshold, she fails to hold a justified belief.

There are, then, a number of ways of explaining what goes wrong in cases like SOCIAL CLUB. In what follows, I will assume, rather than argue, that the moral encroachment thesis provides the best of these explanations. The next section will consider a well-known problem for this use of the moral encroachment thesis. The remainder of this paper will argue that a broader interpretation of the moral encroachment thesis for racial profiling, which I develop in §5, can help us to resolve this problem. While Basu’s diagnosis of the harm that shifts the evidential threshold in cases like SOCIAL CLUB seems reasonable, we might want to give this harm an interpretation that is both more concrete and more general. In §6, I do just this.

4 A Puzzle About Moral Encroachment

The moral encroachment thesis is motivated by its similarity to the pragmatic encroachment thesis. It's further motivated by its ability to explain at least part of the wrongness of racial profiling. But these two motivations pull the thesis in different directions. The account of moral encroachment that seems to provide the best explanation of cases of racial profiling found in the literature has a different structure than standard accounts of pragmatic encroachment.

Following Moss (2018, p.195), we might call a belief 'costly' just in case that belief would lead to a significant harm. And we might call a belief 'risky' just in case that belief would lead to a significant harm if and only if the belief turned out to be false.⁹ As Moss notes, it's tempting to think that the harm of racial profiling is the harm of holding a costly belief. The moral harms gestured towards by Basu's comments seem to be harms that are suffered by members of marginalized groups, not in virtue of our being mistaken about them, but in virtue of our holding certain beliefs about them, regardless of whether or not these beliefs are true. However, this departs from cases of pragmatic encroachment, which maintain that it is the *riskiness* of beliefs that raise the evidential threshold and make justified belief (or knowledge) more difficult to come by. In the BANK CASE, it is our risk of being wrong about the bank being open on Saturday that entails that we may not have a justified belief in this proposition. There is, then, a puzzle in the form of three plausible, but jointly inconsistent claims that capture the relationships between moral encroachment, pragmatic encroachment and racial profiling:

A Puzzle About Moral Encroachment

1. Moral encroachment should be understood by analogy with pragmatic encroachment.
2. Pragmatic encroachment affects the epistemic status of beliefs in virtue of their riskiness.
3. Moral encroachment affects the epistemic status of beliefs in virtue of their costliness.

While the practical harms that shift the evidential threshold in certain cases arise in virtue of the falsity of these beliefs, then, cases of racial profiling suggest that the moral harms that shift the evidential threshold arise in virtue of the mere holding of these beliefs, making the pragmatic encroachment thesis and the moral encroachment thesis disanalogous.

While Moss acknowledges that this line of reasoning is tempting, she thinks that we ought to resist it. She argues that the most promising solution to the "challenge" posed by racial profiling for the moral encroachment thesis is to reject (3) by recognizing that "there

are moral harms that bear a distinctive connection to false racial profiling.” [196] She goes on to argue that such opinions can be harmful in virtue of forcing certain individuals to confront “general stereotypes that do not reflect [their] character” and, also, in virtue of the “looping effects that constitute distinctive harms for victims of false profiling.” [197]. While Moss does not elaborate on this last point, the looping effects she has in mind here plausibly involve, among other things, internalizing and coming to act in accordance with the expectations fostered by the negative stereotype one has been mistakenly ascribed.¹⁰ While there are many possible ways one might respond to being mistaken for a criminal, one of the possible consequences of internalizing this stereotype is coming to act it out.¹¹

Moss assumes, then, that identifying *some* harms that arise only in cases of false racial profiling provides ‘adequate grounds’ for applying a version of moral encroachment that makes the evidential threshold sensitive to the risks of believing, rather than the costs, to cases of racial profiling. [197] However, there are worries about her solution to the puzzle. It’s unclear that the harms Moss identifies will arise very often in cases of false racial profiling. Whether or not an individual is forced to confront general stereotypes that do not reflect their character—and whether or not these looping effects occur—will depend, at the very least, on whether the individual in question is able to recognize that they are being profiled. Arguably, they will not be in a position to recognize this very often. While cases like SOCIAL CLUB dramatize instances of racial profiling, it’s reasonable to think that in many (though not all) mundane cases where an individual holds a belief about a Black person being a staff member, or about a Black person being a criminal, this will go unnoticed by the Black person in question, even if one acts on this belief by crossing the street to avoid the approaching Black man, or by following the Black woman around the department store.

The irregular nature of the harms Moss identifies undermines the claim that these harms bear the sort of distinctive connection with false racial profiling capable of supporting the claim that it’s the risk of being wrong in such cases that shifts the evidential threshold. It would be odd to claim that these risks are important enough to shift the evidential threshold while conceding that the harms they involve almost never occur. While I think Moss’s observations are on the right track, then, they could use some more support. Like Moss, I’ll assume we can reject (3) provided that we can identify at least some significant risks so closely connected with cases of racial profiling that they make plausible that cases of mistaken profiling alone have the moral features required to shift the evidential threshold.

A desiderata for an account that would allow us to reject (3), then, is that it include at least some harms that arise reliably in cases of profiling where we have made a mistake, ideally because such harms are grounded in the structure of such cases. The final section of this paper will offer an account of this type of harm.

Before moving on to that, it's worth considering whether there might be some other way of resolving our puzzle. A different approach might be to reject (1), or at least to rethink what it implies. One strategy along these lines might be to claim that pragmatic encroachment's appeal to risky beliefs is a feature of the *pragmatic* nature of pragmatic encroachment rather than a feature of encroachment more generally. One might argue that the reason that pragmatic encroachment, in particular, involves risky beliefs is because the sort consequentialism that pragmatism involves make our obligations sensitive to which state of the world comes about when one acts upon one's belief. One might further point out that since the moral encroachment thesis is a metanormative thesis, it should be able to be paired with a deontological conception of moral harm instead.¹² And perhaps a deontological conception of moral harm—the sort alluded to above by Basu—needn't be sensitive to the truth or falsity of the belief in question. If this is right, it could explain why the moral encroachment thesis is a thesis about costs rather than risks. A deontological conception of harm that grounds the harm of profiling in our relationships with others, and perhaps with ourselves, rather than in the outcomes of our actions, is less likely to be sensitive to what state of the world we happen to be in.

Of course, a variation of the strategy that involves rejecting the first claim as false is to maintain that this claim actually is *true*, and consistent with the other claims in our triad, upon a weaker way of conceiving the analogy. While moral encroachment and pragmatic encroachment might be disanalogous in the senses brought out by the second and third claims in our puzzle, they are perfectly analogous in making the justification of an opinion depend upon some non-epistemic features. Why think that moral encroachment should be analogous to pragmatic encroachment in the stronger way that gives rise to our puzzle?¹³

The fact that this puzzle has already been raised is perhaps reason enough to want to resolve it. Moreover, there's good reason for the concern it expresses. While it's not unreasonable to try to justify either a disanalogy between moral and pragmatic encroachment, or a weaker version of this analogy, it forces the proponent of the moral encroachment thesis to give up quite a bit. In abandoning an important part of the structure that defines and

motivates pragmatic encroachment, one might worry that the moral encroachment thesis would sacrifice too much of the plausibility it borrows from its connection to a much more established view. While I don't think the lack of a stronger analogy between the moral encroachment thesis and the pragmatic encroachment thesis would give us decisive reason to reject the former, it seems reasonable to think that the former thesis becomes at least a little weaker to the extent that the analogy becomes weaker. This is all that is needed to make a different solution to the puzzle worth exploring. At the end of the day, whatever account of moral encroachment we endorse will need to be attractive enough to stand on its own, whether or not it exemplifies the features of paradigm cases of pragmatic encroachment. But that doesn't mean that appealing to paradigm cases isn't the right place to start to look for such an account. Ideally, the account we end up with will be the result of balancing the features of such accounts with our intuitions about how the practical and the moral come apart, so that our account of moral encroachment is informed, though not determined, by our account of pragmatic encroachment and, more contentiously, perhaps also vice versa.

With this in mind, in the next section, I introduce a different type of profiling error that broadens our understanding of the moral encroachment thesis. While I take this to be an interesting result on its own, in the last section of this paper, I'll also argue that this profiling error supports the idea that it is the risks of believing, rather than the costs, that make a difference to whether or not a belief is justified.

5 Negative and Positive Profiling

Our socio-historical context gives rise to all sorts of true negative racial generalizations. It grounds the prevalence of Black staff members and that of Black criminals. But our socio-historical context also gives rise to a different type of true racial generalization. It grounds generalizations about the mistreatment of individuals belonging to the marginalized racial group in question. Along with the fact that members of certain marginalized groups are more likely to have certain putatively negative traits, like being a staff member or being a criminal, it is almost trivially true that members of marginalized groups are more likely to be *victims*. The ways that they are known to be victimized also make them worthy of a special kind of moral concern that affects the epistemic status of our beliefs.

This section considers the class of racial generalizations that mark these individuals in this way and argues that the moral encroachment thesis can equally provide an explanation

for a different type of wrongness. This wrongness arises in virtue of the fact that the nature of these generalizations ought to make the inferences that they license *easier* rather than *harder* to justify. The nature of these generalizations lower, rather than raise, the evidential threshold for justified belief.

As mentioned, it's almost trivial that there are true generalizations that ascribe to members of marginalized, racial groups characteristics that suggest that they fare less well than members of the dominant group. One example of this that has recently received a lot of attention is the disparity in the treatment of pain of racial and ethnic minorities. Here are some of the statistics:

- Racial/ethnic minority patients are less likely than White patients to receive any pain medication.¹⁴
- Racial/ethnic minority patients are more likely to receive lower doses of pain medications.¹⁵
- Racial/ethnic minority patients are more likely to have longer wait times to receipt of pain medication in the emergency department.¹⁶
- Racial/ethnic minority patients are less likely to receive opiates as treatment for pain.¹⁷
- Hispanics and African Americans more frequently than Whites reported fear of discrimination based on race/ethnicity when seeking relief for chronic pain.¹⁸
- Racial/ethnic minority patient requests for pain relief are perceived to be ‘drug seeking’ behavior more often than such requests from non-minorities.¹⁹

Shavers et al. (2010, p. 179) conclude that, “there is fairly consistent evidence that racial/ethnic minorities suffer disproportionately from unrelieved pain compared with Whites”. I want to focus on this last claim. To be even more precise, let us focus on the claim that *Black Americans are more likely to have their pain ignored and undertreated*. While the causes of this statistic are varied and complex, it's more than plausible that discrimination plays at least some role here.²⁰ As Shavers et al. (2010, p.207) put it, when it comes to investigating contributions to racial/ethnic disparities in pain management, “there is ample evidence in the literature that racial/ethnic minorities are often less favorably viewed than Whites.” With these facts in mind, consider the following case:

PAIN. James, a Black American, goes to the emergency room complaining of shoulder pain. He's been in the emergency room three other times before in the past month, each time complaining of a different ailment. These sorts of complaints are common

in this particular ER, which is a hot spot for addicts in search of opioids. The doctor examines James's shoulder. Unable to find anything wrong with it, he sends him on his way without any pain medication. The doctor does not believe that James is in pain.

Is the doctor justified in failing to believe James? Let's take things in steps. First, the doctor in PAIN seems to have at least some evidence, in virtue of his testimony, that James is telling the truth about his pain. But even if James's testimony is very convincing, the fact that a physical exam revealed no further evidence of James's pain (broken bones, physical deformities, etc.) is a reason that speaks against believing that he is in pain. In addition to these more individualized considerations, there is statistical evidence that pulls in opposite directions. There is statistical evidence that indicates that a non-trivial number of patients who are frequent visitors to the ER, in the way that James is, are not being truthful when they claim to be in pain. But there is also statistical evidence that indicates that James is a member of a marginalized group that has a history of having its pain ignored.

With so many considerations pulling in different directions, it's difficult to know whether the doctor is justified in failing to believe James. Here's one thing I think we can say for sure. The fact that James belongs to a racial group that has a history of having its pain ignored seems to provide, not merely an epistemic consideration, but also a *moral* consideration that must somehow figure into the doctor's calculus about what to believe. A little more carefully: to whatever extent we feel that cases like SOCIAL CLUB involve some sort of moral factor that ought to be taken into account when we decide what we ought to believe, we should feel that cases like PAIN do so as well. Just as many feel that Agnes should take special care before coming to conclude that the Black man in SOCIAL CLUB is a staff member, in virtue of the sort of moral harm that Basu alludes to, the doctor in PAIN should take special care before *failing* to conclude that James is in pain, in virtue of the possibility of inflicting upon James a similar kind of harm.

This last appeal to intuition marks an important difference between SOCIAL CLUB and PAIN. Earlier we defined racial profiling as the practice of forming opinions about an individual on the basis of statistics about members of their racial group. The sorts of cases that are standardly discussed in the literature are what we might call cases of *negative profiling*. They are cases where we have a true racial generalization that seems to license a negative opinion about some individual, despite the intuition that this amounts to depriving them

of the special moral concern that we owe them. The sort of true racial generalization that we have about the pain treatment of Black Americans also licenses an inference, though one that attributes to an individual a putatively non-negative trait. In this case as well, we have the intuition that the individual in question is worthy of special moral concern. But, unlike in SOCIAL CLUB, this special moral concern should not prompt us to *refrain* from drawing this inference in certain situations. Instead, this special moral concern requires that we be extra prone—more *positively inclined*—to draw this inference. In the case of Black Americans like James, we ought to be more inclined to believe their claim that they are in pain, on the basis of a statistic that tells us that such claims are often ignored.

We might call the sorts of cases just described instances of *positive profiling*, by contrast with the more familiar instances of *negative profiling* we get in SOCIAL CLUB and elsewhere, to flag this difference. Corresponding to these two types of profiling practices are two types of *profiling errors* we might make. There is the *positive profiling error* that we make when we fail to shift the evidential threshold in response to the harm of failing to believe something about an individual on the basis of their membership in a racial group. This is the sort of error that we sometimes make when we fail to infer that an individual is in pain given the statistical evidence about Black Americans described above. And there is the *negative profiling error* that we make when we fail to shift the evidential threshold in response to the harm of believing something about an individual on the basis of their membership in a racial group.²¹ This is the sort of error that many have claimed we get in SOCIAL CLUB and elsewhere. I want to suggest that insofar as the moral encroachment thesis is appealed to in order to explain the wrongness of costly or risky belief states, it ought to be able to explain belief states that are costly or risky in virtue of *both* these types of profiling errors.²²

To develop this idea a little further, recall that in §2, we saw that the BANK CASE supports the following principle:

HRHJ: The higher the risk of acting on p , the higher the evidential threshold and the harder it is to have a justified belief in p .

Since an individual who makes a positive profiling error doesn't harm the individual in question by holding a particular belief, but instead harms this individual by *failing* to hold a particular belief, such cases make plausible the principle that the higher the risk or cost of failing to hold a belief, the easier it is for this belief to be justified:

HREJ: The higher the risk or cost of failing to believe p , the lower the evidential threshold and the easier it is to have a justified belief in p .

The view that encroachment is symmetrical in this way is not without precedent. The discussion in Pace (2011), one of the very first to use the term ‘moral encroachment’, suggests that the importance or value of having a belief can lower the threshold of epistemic support needed to justify that belief. Pace illustrates this through the example of theistic belief. Bolinger (2020b) also describes a view, which she calls “robust encroachment”, that is symmetrical in this way, and notes several more recent accounts of moral encroachment that are instances of it.²³ The main contribution this paper makes is to introduce this symmetry into the discussion of racial profiling—to introduce the concept of a positive profiling error—and to argue that it can teach us something interesting about the structure of moral encroachment by helping us to resolve the puzzle we were left with in §4.

One might resist the distinction I am attempting to draw between the negative and positive profiling errors that arise when we fail to be more negatively or positively inclined to ascribe a certain trait. Consider that some of the attributes one ascribes in paradigmatic cases of so-called negative profiling aren’t obviously negatively valued traits. While being a criminal is clearly such a trait, it’s less clear we can say the same thing about being a staff member, absent classist assumptions we clearly shouldn’t be making. Conversely, being a victim of racism and having one’s pain ignored seem like intuitively negative traits. And yet here I am suggesting that we can use them to motivate the concept of positive profiling. One might worry, then, that drawing a distinction between positive and negative profiling will be impossible. One might further worry that, even if such a distinction isn’t impossible in principle, the sorts of classist assumptions exhibited in examples like SOCIAL CLUB suggest that our intuitions are ill-suited for determining whether some property falls into either of these categories.²⁴

But we needn’t endorse particular assessments as intrinsically negative or positive to talk in a meaningful way about negative and positive profiling. Paradigmatic cases of negative and positive profiling are intended to invoke the intuition that certain property ascriptions would be perceived of as undesirable by members of the marginalized group in question, while others would be perceived of as desirable. However, we needn’t think of these properties as intrinsically negative or positive to acknowledge that certain historical and contextual facts make their ascriptions to members of these groups more or less desirable or appropriate.

We needn't think that having one's pain ignored is a good thing to concede that being *recognized* as being vulnerable in this way is desirable in certain contexts. Among other things, this amounts to recognizing that one is a truth-teller in a context where, because of racism, this sort of recognition has historically been denied. While it's useful, then, to motivate the distinction between negative and positive profiling by appealing to traits whose valences are more objective, like being a criminal, such traits aren't necessary to define this distinction. What upholds this distinction is the fact that the ascription of certain traits would be harmful, when bestowed or withheld, for historical or contextually determined reasons. Such reasons play a dual role in both providing the valence of such traits, and also giving us more or less reason to ascribe them.

What about the worry that our intuitions are ill-suited for determining what these reasons recommend? Even if we think that the justification of our beliefs should be responsive to a history that has resulted in Black people being relegated to less desirable employment opportunities, one might insist that a well-paid waiter's job shouldn't fall into this category. However, I want to suggest that while the appeal to, for instance, certain elitist character traits may be problematic for what it says about the audience whose intuitions this appeal is attempting to elicit, it isn't problematic for the moral encroachment thesis itself. Since the thesis itself is neutral about the moral properties that underwrite encroachment, we can endorse it without endorsing all the intuitions that motivate it. This isn't to say that the thesis would stand if there were no clear cases. My own argument relies on the idea that we can determine the contours of the thesis by appealing to particular examples. However, the moral encroachment thesis can take into account moral progress. It can withstand uncertainty and disagreement about the cases to which it applies.

Even if one isn't worried about the coherence of the concept of positive profiling, one might still object that the central example I've used to motivate it is descriptively inadequate. One might object that we have no reason to assume that the failure to believe James will result in the doctor failing to administer any pain medication. Instead, we might reasonably think that the doctor would give patients in this situation the benefit of the doubt. If so, it's less clear there is a moral harm in the offing for the evidential threshold to be responsive to.

Whether the resolution my example describes seems plausible will depend upon the details of the case. One might imagine filling out PAIN in a way that casts more doubt on James, to the extent that giving him the benefit of the doubt no longer seems reasonable. One

might imagine, for instance, that James is well-known to have deceived these doctors in the past. Or one might imagine a hospital low on supplies having to triage—in such a situation, a doctor might very reasonably give James a lower dose of pain medication than he requires. To motivate the idea of a positive profiling error, it's enough that there are cases along the lines of PAIN that would reasonably lead the doctor to withhold adequate treatment in cases where we fail to lower the evidential threshold. Since the statistics described at the beginning of this section confirm that this happens all the time, the only interesting question is how to fill in the details of the case that would lead to this result.

One might still press that failing to believe that James is in pain is consistent with having a high credence that he is in pain and prescribing the drugs anyway. If this is right, we don't need to appeal to the moral encroachment thesis to provide an explanation for what the doctor ought to do.²⁵ But it's not clear that whatever considerations lead the doctor to fail to believe James should not also lead him to fail to maintain a high credence that James is telling the truth. More importantly, as we will see shortly, my diagnosis of the harm involved in such cases lies not just in wronging in virtue of acting, but in wronging in virtue of failing to believe. This suggestion would not, then, alleviate the need to interpret the moral encroachment thesis in the way that I am proposing.

Supposing that one is on board with the example, one might still be skeptical that positive profiling is an interesting concept at all. For one might suspect that any so-called positive profiling error can be redescribed as a negative profiling error that we are able to explain using the moral encroachment thesis in the usual way. Instead of describing PAIN as a case where the high moral stakes make the justification of the belief that *James is in pain* easier to come by, one might describe it as a case where these high stakes make the justification of the belief that *James isn't in pain* harder to come by. The latter belief might be thought to amount to ascribing to James the negative characteristic of being a liar, and we might think that we would want to avoid this, in the same way that we would want to avoid pegging him as a criminal. If this is correct, there is no need to expand the concept of moral encroachment in the way that I am proposing.

But we clearly can't redescribe PAIN in this way. For then in order to be the sort of case of racial profiling that the moral encroachment thesis plausibly explains, it would need to be the case that forming the belief that *James isn't in pain* amounts to forming this opinion about him based on a true racial generalization. Consider that if it were a false racial

generalization, this case would not exhibit what Bolinger (2020b) calls “sufficiency failure”. The falsity of the statistical evidence would be sufficient to explain why we should not form this belief, and the appeal to the moral encroachment thesis would not be necessary. To reinterpret PAIN as involving a negative profiling error, then, the generalization that *Black Americans are disproportionately prone to lying about their pain* or that *Black Americans are disproportionately unlikely to be in pain*, or something in the neighborhood of these claims, would have to be true. But they aren’t true. The concept of a positive profiling error is, therefore, not redundant.

Finally, in the spirit of this last response, one might be tempted to object that what appear to be positive profiling errors are actually errors that don’t involve a moral flaw at all. Notice that the true racial generalization in such cases is one that *supports* the belief that we, intuitively, ought to adopt. Therefore, where we fail to believe that James is in pain, perhaps we can simply say that the problem with this belief is that it ignores the statistical evidence that indicates that failing to believe in such cases is often a mistake. It might be thought that what we have in this case, then, is a classic example of a higher-order belief that undermines a first-order belief. The fact that the doctor knows that the pain of Black Americans often goes undertreated provides him with higher-order evidence that his first-order assessment of James is not to be trusted, with no moral remainder for the moral encroachment thesis to explain.

However, I want to suggest that *in addition* to the epistemic reason that this statistical evidence provides, this evidence also provides a moral consideration that must be taken into account. This moral consideration is more difficult to discern in this case because, unlike in the case of negative profiling, it runs in the same direction as the epistemic consideration. But, to see that it does indeed exist, consider the following case:

PAIN 2. Jack, a white American, goes to the emergency room complaining of shoulder pain. He’s been in the emergency room three other times before in the past month, each time complaining of a different ailment. These sorts of complaints are common in this particular ER, which is a hot spot for addicts in search of opioids. There is a statistic that says that individuals who enter the ER at the particular time of day that Jack did are more likely than the average person to have their pain undertreated. The doctor examines Jack’s shoulder. Unable to find anything wrong with it, he sends

him on his way without any pain medication. The doctor does not believe that Jack is in pain.²⁶

Notice that PAIN and PAIN 2 differ only in what grounds the statistic that says that the patients in these cases belong to a demographic group whose pain is disproportionately ignored. However, while the doctors who fail to believe James and Jack may both be guilty of neglecting a relevant piece of higher-order evidence, I think we would want to say that the former doctor's reasoning is susceptible to an *additional* flaw that the latter doctor's reasoning is not susceptible to. There is a moral consideration that needs to be incorporated into the former doctor's reasoning about what to believe, which yields the result that the doctor's belief that James is in pain is more easily justified than the doctor's belief that Jack is in pain. If the evidence in both cases is reasonably weak, then it may be the case that the doctor *is* justified in failing to believe Jack, even though he is *not* justified in failing to believe James. At least we should think this to the extent that we believe that Alice does nothing wrong in the following modified version of SOCIAL CLUB, where the owner has a strong and completely idiosyncratic preference for hiring red-headed staff members:

SOCIAL CLUB 2. Alice and Esther are members of a swanky D.C. social club with a strict dress code of tuxedos for both male guests and staff members, and dresses for female guests and staff members. While preparing for their evening walk, the two women head toward the coat check to collect their coats. As they approach the coat check, they both look around for a staff member. As Alice looks around she notices a well-dressed red-headed man standing off to the side and tells Esther, "There's a staff member. We can give our coat check ticket to him."

Presumably, Alice's belief about the red-headed man being a staff member is justified, even though the justification of Agnes's belief in SOCIAL CLUB is questionable—for the same general reason that the doctor's failing to believe that Jack is in pain may be justified, even if we think that his failing to believe that James is in pain may be questionable. If we assume that each pair of cases are identical to each other, in all epistemically relevant respects, except for the characteristic that grounds the statistical evidence, then this characteristic ought to shift the evidential threshold in a way that makes it harder for Agnes's belief about the staff member to be justified and easier for the doctor's belief about James to be justified, than it is for the judgments in their counterpart's cases to be justified.

Summing up, in cases of negative profiling, there is moral pressure to discount the probative value of evidence, sometimes very strong evidence. By contrast, I've suggested that we can identify a class of cases—cases of positive profiling—where there is moral pressure to take into greater account the probative value of evidence, perhaps quite weak evidence. Common to both cases is the idea that the moral stakes change the standard for what it takes to have a justified belief.

6 Risk Revisited: Perpetuating the Status Quo

I've argued that there are two different types of profiling errors we might make and that there is a natural way of expanding our account of moral encroachment that leaves it able to explain both of them. One reason for expanding our account of moral encroachment is this increased explanatory power. A second reason is to get an account of moral encroachment that more perfectly mimics pragmatic encroachment. Recall our inconsistent triad from §4:

A Puzzle About Moral Encroachment

1. Moral encroachment should be understood by analogy with pragmatic encroachment.
2. Pragmatic encroachment affects the epistemic status of beliefs in virtue of their riskiness.
3. Moral encroachment affects the epistemic status of beliefs in virtue of their costliness.

By analogy with negative profiling, call the failure to positively profile *costly* just in case failing to hold the belief in question would lead to a significant harm. And call it *risky* just in case failing to hold the belief in question would lead to a significant harm if and only if the belief turned out to be true. This section will defend the idea that there is a harm distinctive of the risk involved in failing to believe. It will do this by arguing that failing to believe what is true in cases like PAIN is *constitutive* of a significant moral harm. If we assume that costly and risky belief states represent unified classes across negative and positive profiling errors, the argument of this section undermines the challenge posed by racial profiling for the moral encroachment thesis by providing us with reason to reject (3).

There are, on the one hand, strong consequentialist reasons for thinking that James will be harmed in PAIN if and only if the doctor fails to hold a true belief. Most obviously, if the belief that *James is in pain* is false, then James is not actually in pain. Only if the belief is true will James's pain go on untreated. Exacerbating this, for Black Americans like James, are distinct harms that accompany untreated pain that is the result of discrimination. As

Shavers et al. (2010, p.207) note, “anger, helplessness, and depression are frequent consequences of experiences with racial/ethnic discrimination and have been known to affect the experience of pain negatively.”

There are, then, clear consequentialist reasons for thinking that, insofar as the moral encroachment thesis can explain the badness of failing to positively profile, it does so in virtue of explaining the harm of failing to hold certain true beliefs. But we need not conceive of the harm involved in cases like PAIN in consequentialist terms in order to reach this conclusion. There’s a more fundamental reason for thinking this. Recall that Basu suggests that more moral care is required when forming a belief on the basis of evidence that is the result of our racist history. One way of interpreting this duty of moral care is as the duty to make sure we don’t contribute to perpetuating an unjust status quo. Plausibly, the characteristic harm that we cause when we believe so casually that Black people are staff members is that of supporting, in an indirect way, a state of affairs where Black people will continue to be overrepresented in these types of roles.

But if our duty of care is not to believe in a way that perpetuates the status quo, then where we make a positive profiling error, the content of the statistic *itself* seems to entail that the strongest form of the duty that we violate is the duty not to fail to hold a true belief. The strongest form of the duty not to perpetuate the status quo in PAIN is the duty to believe James when he is telling the truth. To see this, notice that Black Americans who have had their pain ignored are *by definition* individuals whose *true* claims we have failed to believe. This means that, in cases like PAIN, the doxastic act of the doctor reinforces the status quo by *directly contributing* to the very statistic that grounds the profiling error of failing to believe that *James is in pain* if and only if this belief is true. There will be one more Black American who has had their pain ignored *if and only if this belief turns out to be true*. Moreover, there will be one more Black American who has had their pain undertreated if and only if this belief turns out to be true. For consider that if the belief that *James is in pain* is false, then James is not actually in pain. Only if the belief is true can James’s pain go on untreated as a result of the doctor’s failure to believe him. Thus, the doctor directly contributes to the status quo by contributing to the statistic that *Black Americans are more likely to have their pain ignored and undertreated* if and only if he fails to hold a true belief. He contributes to the status quo both in virtue of failing to hold the belief in question, and also in virtue of the state of affairs his failing to act on this belief brings about.

This interpretation of the distinctive harm of failing to positively profile in cases where the belief in question is true overcomes the concerns raised earlier about Moss's defense of the moral encroachment thesis as a thesis about risky belief. Recall I criticized Moss's defense on the grounds that the harms she identifies don't arise as systematically as we would need in order to maintain that the risks involved in bringing them about bear a distinctive connection to false racial profiling. I suggested that a desiderata for an account of the harms distinctive of risky belief is that they arise reliably in cases where these harms are the result of having made a mistake, ideally because they are grounded in the structure of the cases in question. My account clearly satisfies both parts of this desiderata. Because the positive profiling error of failing to believe what is true is partly *constitutive* of the harm it involves—the harm of contributing to the status quo by ignoring James's pain—there is no doubt about whether this harm will come about as a result of this profiling error. This harm is *entailed* by this profiling error. This contrasts with the harms that Moss identifies, which require the agent to recognize that they are being profiled. In addition to this, the undertreatment of James's pain relies upon what seems reasonable, which is that failing to believe that James is in pain will likely lead the doctor to withhold treatment—in much the same way that believing that the bank is open tomorrow will likely lead Hannah and Sarah to miss paying their bill by failing to make their deposit today. In both cases, the action that most immediately follows from holding/withholding the belief in question will harm Hannah and Sarah/James if and only if the belief turns out to be false/true.

My account offers a reason, grounded in the very structure of positive profiling errors, to think that it is the risk of being wrong that grounds the moral encroachment thesis. If we are looking for a principled way to identify the harm that is distinctive of cases of moral encroachment, I want to suggest that we can't do much better than to identify it as the harm of contributing to the very injustice that grounds these cases to begin with. The fact that the statistic that grounds this profiling error is reinforced if and only if we are wrong about James gives us a principled reason to hold that the harm that grounds moral encroachment in this case is the risk of being wrong about James, and not the cost of failing to believe him. While there may be other harms incurred by those whose false claims of pain we fail to believe, the fact that these harms do not directly contribute to the statistic at issue gives us reason to think that they are derivative of the harm of risky belief, which does precisely this. Plausibly, the reason failing to believe a lying James is psychologically harmful is because

of the realization that this would both perpetuate the status quo, and also harm James, if James were telling the truth. My diagnosis of the harm of positive profiling errors provides a plausible explanation for why the harms of merely failing to believe in such cases don't shift the evidential threshold, while allowing us to appreciate their connection to those harms that do.

One might object that the fact that negative and positive profiling errors involve different types of risks undermines my solution to the puzzle. The negative profiling errors discussed earlier in this paper do not entail the particular harms described in this section. More generally, negative profiling errors do not inevitably contribute to any statistic that grounds encroachment. We don't inevitably contribute to the statistic that Black men in swanky social clubs are more likely to be staff members, or that Black men in particular neighborhoods are more likely to be criminals, merely by falsely believing that a particular Black man is a staff member or a criminal. However, in highlighting the harm of contributing to the status quo, my proposal lends support to the idea that the harms that negative profiling errors involve might also be best interpreted as those of risky belief. It does this by offering a more general, and ultimately more satisfying, interpretation of the harm generated by the looping effects referred to in Moss's discussion. Recall we noted earlier that these looping effects may involve, among other things, internalizing and coming to act in accordance with the expectations fostered by the negative stereotype one has been mistakenly ascribed. While what this looks like will differ from individual to individual, and while these effects aren't the inevitable result of negative profiling, one might suppose that what makes these effects so pernicious is that they are likely to contribute to the racial generalization in question, where they do occur. The nature of these effects makes plausible that where the Black man on the street *does* recognize that we falsely believe that he is a criminal, we may be contributing to the very statistic that initially led us to draw this negative inference by making someone who was not previously a criminal more likely to act in accordance with this ascription. In mistaking this man for a criminal, or a waiter, we make him more likely to become one. There is more likely to be one more Black criminal or waiter if and only if our belief about this individual was initially false.²⁷

My account of the harm that is distinctive of positive profiling errors in cases of racial profiling is, then, equally an argument for why negative profiling errors involve risky, rather than costly, belief. While cases of mistaken racial profiling don't inevitably contribute to

the status quo by contributing to the statistic at issue, they do contribute to the statistic at issue where the looping effects described above come about as the result of negative profiling. While my diagnosis of the harm involved in positive profiling errors reveals why Moss's solution to the puzzle was inadequate, it also allows us to see why this solution was on the right track, and why, when supplemented with my own proposal, we might reasonably conclude that the risk of making a mistake is what shifts the evidential threshold in all of these cases.

Summing up, this paper has argued that there exist harms tied closely enough to cases of racial profiling involving risky belief to undermine the challenge posed by racial profiling for the moral encroachment thesis. The existence of positive profiling errors brings this thesis closer to the ideal of pragmatic encroachment that made it seem so initially plausible. In closing, it's worth emphasizing again that the increased explanatory power such errors lend to the moral encroachment thesis is due, in the first instance, to the fact that the moral harm of the belief withheld in PAIN and the belief bestowed in SOCIAL CLUB have the same source. Our racist history gives rise to both types of cases and to the profiling errors that these cases involve.

7 Acknowledgments

Thanks to Georgi Gardiner and to participants at a session of the 2022 Eastern APA for feedback on an earlier version of this paper. I'm especially grateful to several anonymous referees for extensive and very helpful comments and suggestions that helped improve this paper considerably.

Notes

¹Arguably, the term 'evidentialism' picks out a cluster of related positions that intersect in different ways with the views this paper discusses. However, I follow early proponents of pragmatic views in taking evidentialism to be the name of the position they oppose. For a canonical evidentialist view, see Feldman and Conee (1985).

²For a contemporary proponent of the view that all reasons for belief are practical reasons, see Rinard (2019). There are also more moderate views, which say that practical reasons are among some of the reasons for our beliefs. See, for instance, Howard (2016).

³While early proponents of pragmatic encroachment formulate it as a view about knowledge (e.g., Hawthorne et al. (2004), Stanley (2005)), I will formulate the view as one about justified belief since most

proponents of the knowledge view would accept the justified belief view as well (see, for instance, Fantl and McGrath (2002), Fantl and McGrath (2007), and Fantl et al. (2009)). Moss (2018), who we will discuss in §4, is a notable exception to this, though nothing of importance turns on this.

⁴This definition can be found in Moss (2018, p.177) and Basu (2021, p.1), and something very close to it is widely assumed throughout the literature. It is deliberately general since a large part of this discussion will involve arguing for a particular way of interpreting it.

⁵This definition is from Moss (2018, p.178) though, again, something like it is widely assumed in the literature.

⁶This idea is referred to as the problem of naked statistical evidence. For the locus classicus of this problem in the legal scholarship, see Nesson (1979).

⁷See Dretske (1970) and Lewis (1996) for two canonical versions of the relevant alternatives approach.

⁸For reasoning along similar lines, see Bolinger (2020a).

⁹These formulations are generalizations of Moss’s definitions, which appeal to the harm of *acting* on beliefs rather than the more general harm of holding beliefs. While this more general account is needed to make these definitions relevant to the wider range of discussions this paper addresses, it’s worth noting that since my own diagnosis of the harm of racial profiling has both a belief-theoretic and action-theoretic interpretation, it’s consistent with Moss’s more narrow definitions.

¹⁰The concept of ‘looping effects’ being assumed here is from Hacking (1995).

¹¹For a general discussion of the association between internalized racism and violence, see, for instance, Bryant (2011).

¹²Basu and Schroeder (2019) seem to adopt something like this approach.

¹³Thanks to an anonymous referee for raising this concern.

¹⁴Kposowa and Tsunokai (2002).

¹⁵Cleeland et al. (1997).

¹⁶Epps et al. (2008).

¹⁷Kohn et al. (2008).

¹⁸Shapiro et al. (1997).

¹⁹Rosenberg (2001).

²⁰Shavers et al. (2010), Kohn et al. (2008), Kposowa and Tsunokai (2002).

²¹Whether or not these profiling errors ought to be further refined as “failing to believe what is *true*” and “believing what is *false*” is the topic of the next section. This section will remain agnostic about this difference.

²²One might insist that we ought to individuate things even more finely and distinguish two different ways that we might fail to believe. On the one hand, we might disbelieve; on the other hand, we might merely suspend judgment. However, this distinction isn’t crucial for our discussion. The crucial point is that the harm of failing to believe in certain cases—whether this involves suspension of judgment or disbelief—ought to make justified belief in these cases easier to come by.

²³See, for instance, Marušić and White (2018), Crewe and Ichikawa (2019) and Basu (2020).

²⁴Thanks to an anonymous referee for pressing me to address these concerns.

²⁵Thanks to an anonymous referee for this objection.

²⁶We might further stipulate that there is no reason to think that entering the ER at a particular time of day is correlated with anything that could be construed of as giving rise to some form of morally problematic bias or discrimination.

²⁷As Liebow (2016, p.713) puts it, “Since those who have internalized oppressive prejudices often engage in behavior that further perpetuates these biases, internalized oppression is not only a symptom of an oppressive social climate, but it also represents a mechanism for its continued existence”.

References

- Basu R (2019a) Radical moral encroachment: The moral stakes of racist beliefs. *Philosophical Issues* 29(1):9–23
- Basu R (2019b) What we epistemically owe to each other. *Philosophical Studies* 176(4):915–931
- Basu R (2020) *The Specter of Normative Conflict: Does Fairness Require Inaccuracy?*, Routledge
- Basu R (2021) A tale of two doctrines: Moral encroachment and doxastic wronging. In: Lackey J (ed) *Applied Epistemology*, Oxford University Press
- Basu R, Schroeder M (2019) Doxastic wronging. *Pragmatic Encroachment in Epistemology*
- Bolinger RJ (2020a) The rational impermissibility of accepting (some) racial generalizations. *Synthese* 197(6):2415–2431
- Bolinger RJ (2020b) Varieties of moral encroachment. *Philosophical Perspectives* 34(1):5–26
- Bryant WW (2011) Internalized racism’s association with african american male youth’s propensity for violence. *Journal of Black Studies* 42(4):690–707
- Cleeland C, Gonin R, Baez L (1997) Pain and treatment of pain in minority patients with cancer: The eastern cooperative oncology group minority outpatient pain study. *Ann Intern Med* 127(9):813–816
- Crewe B, Ichikawa JJ (2019) Rape culture and epistemology. *Rape culture and epistemology*
- DeRose K (1992) Contextualism and knowledge attributions. *Philosophy and phenomenological research* 52(4):913–929
- Dretske FI (1970) Epistemic operators. *The journal of philosophy* 67(24):1007–1023
- Epps C, Ware L, Packard A (2008) Ethnic wait times differences in analgesic administration in the emergency department. *Pain Manag Nurs* 9(1):26–32
- Fantl J, McGrath M (2002) Evidence, pragmatics, and justification. *The Philosophical Review* 111(1):67–94
- Fantl J, McGrath M (2007) On pragmatic encroachment in epistemology. *Philosophy and Phenomenological Research* 75(3):558–589
- Fantl J, McGrath M, et al. (2009) *Knowledge in an uncertain world*. Oxford University Press
- Feldman R, Conee E (1985) Evidentialism. *Philosophical studies* 48(1):15–34
- Fritz J (2017) Pragmatic encroachment and moral encroachment. *Pacific Philosophical Quarterly* 98:643–661
- Gardiner G (2020) Relevance and risk: How the relevant alternatives framework models the epistemology of risk. *Synthese* 199(1-2):481–511, DOI 10.1007/s11229-020-02668-2
- Gendler TS (2011) On the epistemic costs of implicit bias. *Philosophical Studies* 156(1):33
- Hacking I (1995) The looping effect of human kinds. i: D. sperber, d. premack d, & aj premack (red.) *causal cognition: An interdisciplinary approach*

- Hawthorne J, Hawthorne JP, et al. (2004) Knowledge and lotteries. Oxford University Press
- Howard C (2016) In defense of the wrong kind of reason. *Thought: A Journal of Philosophy* 5(1):53–62
- Kohn M, Pletcher M, Kertesz S (2008) Trends in opioid prescribing by race/ethnicity for patients seeking care in u.s. emergency departments. *JAMA* 299(1):70–78
- Kposowa A, Tsunokai G (2002) Searching for relief: racial differences in treatment of patients with back pain. *Race and Society* 5(2):193–223
- Lewis D (1996) Elusive knowledge. *Australasian journal of Philosophy* 74(4):549–567
- Liebow N (2016) Internalized oppression and its varied moral harms: Self-perceptions of reduced agency and criminality. *Hypatia* 31(4):713–729
- Marušić B, White S (2018) How can beliefs wrong?—a strawsonian epistemology. *philosophical topics* 46(1):97–114
- Moss S (2018) Moral encroachment. In: *Proceedings of the aristotelian society*, Oxford University Press, vol 118, pp 177–205
- Munton J (2019) Beyond accuracy: Epistemic flaws with statistical generalizations. *Philosophical Issues* 29(1):228–240
- Nesson C (1979) Reasonable doubt and permissive inferences: The value of complexity. *Harvard Law Review* 92:1187
- Pace M (2011) The epistemic value of moral considerations: Justification, moral encroachment, and james’will to believe’. *Noûs* 45(2):239–268
- Rinard S (2019) No exception for belief. *Contemporary Epistemology: An Anthology* pp 44–57
- Rosenberg N (2001) Two levels of treatment: pain control ages for minorities and woman. *Milwaukee Journal Sentinel Online*
- Shapiro B, Benjamin L, R P (1997) Sickly cell-related pain: perceptions of medical practitioners. *J Pain Symptom Manage* 14(3):168–174
- Shavers VL, Bakos A, Sheppard VB (2010) Race, ethnicity, and pain among the u.s. adult population. *Journal of Health Care for the Poor and Underserved* 21(1):177–220
- Stanley J (2005) Knowledge and practical interests. Clarendon Press