Canadian Journal of Philosophy (2026), 56: e10, 1–18
doi:10.1017/can.2026.10107

ARTICLE

Moral Testimony, Moral Understanding, and "Rational" Prejudiced Beliefs

Anna Brinkerhoff

Philosophy, Concordia University, Montreal, Canada
Email: annakbrinkerhoff@gmail.com

Abstract
Moral deference strikes many as problematic. Some explain this problem by pointing to understanding: accepting a moral proposition on the basis of testimony is problematic because it's in tension with understanding that proposition, and understanding moral propositions is especially important. In this paper, I argue that there is a set of non-moral propositions that it is also especially important to understand, for many of the same reasons that defenders of The Understanding Explanation give to defend the claim that understanding moral propositions is especially important. These propositions are featured in cases at the center of the debate about moral encroachment.

Keywords: moral testimony; moral understanding; moral encroachment; prejudice; moral agency; affective alienation

1. Introduction

There seems to be something problematic about deferring to moral testimony—that is, accepting a moral proposition just on someone else's say-so. Consider the following case.

> **Gas Guzzler** Coco is in search of a new vehicle and is deliberating about which one to buy. She is drawn to SUVs – she likes the extra room, aesthetics, and power – but wonders whether opting for a vehicle that uses up so much gas is morally wrong, given the threat of climate change. Instead of thinking through things herself, she asks her friend Juno, whom she takes to be very reliable about moral matters. Juno tells her that opting for an SUV is wrong. Coco comes to believe that buying an SUV is wrong based on Juno's say-so.

Moral deference like Coco's strikes many as problematic. Many philosophers agree and seek to identify what, exactly, the problem is. One prominent explanation is The Understanding Explanation,[1] which involves two core claims:

> The Tension Claim: accepting p on the basis of testimony is in tension with acquiring understanding of p.

> The Relative Importance Claim: understanding moral propositions is distinctly or especially important (relative to understanding non-moral propositions).

[1]For versions of The Understanding Explanation, see Nickel (2001), Hopkins (2007), Hills (2016, 2013), and Callahan (2018).

© The Author(s), 2026. Published by Cambridge University Press on behalf of The Canadian Journal of Philosophy, Inc. This is an Open Access article, distributed under the terms of the Creative Commons Attribution licence (http://creativecommons.org/licenses/by/4.0), which permits unrestricted re-use, distribution and reproduction, provided the original article is properly cited.

Defenders of The Understanding Explanation need The Relative Importance Claim because accepting non-moral propositions on the basis of testimony seems fine: we accept non-moral propositions on the basis of testimony all the time without any semblance of a problem. The Relative Importance Claim allows them to say that—although accepting a non-moral proposition on the basis of testimony is in tension with acquiring understanding of that proposition—non-moral deference is not problematic because understanding non-moral propositions is not especially important. Some even contend that there is "a difference *in kind* between moral and non-moral matters when it comes to taking testimony" (Hopkins, 2007, 613; emphasis added).

Others, though, think that the posited asymmetry between moral and non-moral propositions is overstated. They argue that there are some non-moral propositions that it is especially important to understand, and, thus, it is problematic to accept those propositions on the basis of testimony, too.[2]

This paper focuses on one such set of non-moral propositions. These propositions feature in cases that have received an increasing amount of attention in another corner of contemporary epistemology—namely, the ethics of belief. Current debates in the ethics of belief largely center on moral encroachment on epistemic rationality, which is the view that moral considerations get a say in what is epistemically rational to believe.[3] Moral encroachment stands in contrast with more traditional "purist" views, according to which only truth-related considerations determine epistemic rationality.

Moral encroachers motivate their view, in part, by pointing to cases involving beliefs that seem to be both prejudiced (e.g., sexist, racist) but also epistemically rational. These cases involve some unfortunate fact about the world that tracks a historically marginalized social group—for example, the fact that some negative or relatively diminishing property is common among that group. The belief in question is about an individual member of that social group, which is inferred from knowledge of group-level demographic facts. Consider one such case, adapted from Basu (2019b, 2498), by Gardiner (2018):

> **The Server** Spencer works as a server at a restaurant. He sensed that white diners tipped more than Black diners. After a bit of googling, Spencer found a well-documented social trend that most Black diners tip substantially below average. He comes to believe that Black diners tend to tip below average. A Black diner, Jamal, enters Spencer's restaurant. Spencer believes that Jamal will probably tip below average.

My contention is that it is especially *morally* important to understand the relevant non-moral facts in cases like The Server, for many of the same reasons that defenders of The Understanding Explanation in the moral testimony literature give to defend The Relative Importance Claim. This vindicates the increasingly popular view that the once posited asymmetry between moral and non-moral deference is overstated. But, perhaps more interestingly, it undermines an important motivation for accepting moral encroachment.

2. Moral Testimony and The Understanding Explanation

The Understanding Explanation offers an answer to why moral deference is problematic: accepting a proposition on the basis of testimony is in tension with acquiring understanding of that proposition, and understanding moral propositions is distinctly important.

Before exploring The Understanding Explanation further, we first need a working account of understanding. There are several on offer, but—following many others in the literature on moral testimony—I will take on board Alison Hills' account. According to Hills, to understand p is to "grasp the reasons why *p*" (Hills, 2016, 102) where grasping the reasons why *p* is more than just

[2]See especially Groll and Decker (2014); Callahan (2018) also gestures at this claim.

[3]Bolinger (2020) gives a helpful overview and taxonomy of moral encroachment views.

believing or knowing that p because q. In addition, grasping the reasons why p involves a cluster of abilities.

> If you understand why p (and q is why p), then in the right sort of circumstances, you can successfully:
>
> (i) follow an explanation of why p given by someone else;
> (ii) explain why p in your own words;
> (iii) draw the conclusion that p (or that probably p) from the information that q;
> (iv) draw the conclusion that p' (or that probably p') from the information that q' (where p' and q are similar to but not identical to p and q);
> (v) given the information that p, give the right explanation, q;
> (vi) given the information that p', give the right explanation, q'.
>
> To understand why p, you have to have the abilities (i)–(vi) to at least some extent. Of course you can have these abilities to a greater or lesser degree, which may make it tempting to say that moral understanding comes in degrees. (Hills, 2016, 102–103)

With Hills' account in mind, let us consider each of the two core claims of The Understanding Explanation, starting with The Tension Claim.

> The Tension Claim: accepting p on the basis of testimony is in tension with acquiring understanding of p.

The basic idea behind this claim seems to be this: testimony can transmit propositional knowledge—knowledge that, for example, driving an SUV around town is morally wrong—but the cluster of cognitive abilities that are partly constitutive of understanding cannot be (or is often not) transmitted via testimony.[4]

Now let us turn to The Relative Importance Claim, which will be our focus.

> The Relative Importance Claim: understanding moral propositions is distinctly or especially important (relative to understanding non-moral propositions).

The idea behind this claim is that moral understanding is central to moral agency, and this is what makes it especially important. There are a couple of reasons for thinking that moral understanding is central to moral agency.

First, understanding in the moral domain enables us to form true moral beliefs in a wide variety of contexts, and these true moral beliefs allow us to take the morally right or best action in these contexts.[5] For example, if Coco understands why it is morally wrong to drive an SUV around town, then she has the cognitive resources to think well about a host of moral issues informing decisions about her carbon footprint—like whether to opt for the plant-based option at dinner or to shop at consignment stores—as well as collective harms. She can then form true moral beliefs in these contexts—for example, it is morally better to order the tofu rather than the steak—which allows her to take the morally right or preferable action. Coco would not be in this position if she merely knew that driving an SUV is morally wrong.

Second, it has been argued that understanding in the moral domain is necessary for morally worthy action. This holds according to some views of morally worthy action which require acting

[4]See Hills (2016, 2013) for a development of this idea. See Callahan (2018, 454-456) for further defense of The Tension Claim, and exploration of why other activities that do not transmit understanding are not thereby morally problematic.

[5]See Nickel (2001), Hopkins (2007), Hills (2016, 2013), and Callahan (2018) for discussion of this point.

for morality *de re* rather than *de dicto*.[6] On these views, φing is not morally worthy if the agent φs just because she takes φing as morally "right," and she wants to do the "morally right" thing. Rather, φing is morally worthy only if the agent φs because she takes it that φing prevents harm or respects persons (or whatever the reason is that makes φing morally right). According to this view, Coco's choice not to buy the SUV is morally worthy only if her decision is made out of concern for the reasons that make buying the SUV morally wrong—because, say, it contributes to a collective harm, and she does not want to contribute to a collective harm. So, if morally worthy action requires acting for the right reasons *de re*, then it seems that we must have a firm grasp on those reasons in order to act in a morally worthy way. That is, in order to be motivated to (not) φ for the reasons that make φing right (wrong), we must have a grasp on those reasons and the relationship between those reasons and φing. On these views of moral worth, it seems we must have moral understanding in order to act in a morally worthy way.

3. The Relative Importance Claim: Understanding & Affective Alienation

I think there is an additional argument to be made for the relative importance of moral understanding. This argument draws on insights from another corner of the moral testimony literature that identifies affective alienation as the problem with moral deference.[7] As David Enoch puts it, "… there seems to be something objectively *cold* about a moral judgment that is based solely on deference to the expert" (Enoch, 2014, 256; emphasis added). Moral deference, the thought goes, can change minds (beliefs) but it cannot change hearts (feelings)—leaving us "emotionally disconnected" from morally significant features of our environment (Sliwa, 2025, 23).

I want to suggest that considerations about affective alienation can largely be incorporated into The Understanding Explanation. That's because moral understanding and appropriate affects go hand-in-hand. In other words, moral understanding enables or promotes morally appropriate affects: if we understand why (and perhaps in part *by virtue of* understanding why) it is morally right or wrong to φ, we are more likely to have the emotions and motivations that are morally appropriate in light of both p and the correct explanation of why p.[8] For example, by understanding why driving an SUV around town is morally wrong, Coco will be more inclined to feel guilty if she opts for an SUV instead of a sedan, and frustration at others who opt for SUVs.

Callahan (2018) appeals to insights like these to motivate what she takes to be a richer, more holistic view of understanding, according to which understanding constitutively involves an affective dimension as well as a cognitive one. On Callahan's view, one fails to understand why p if one fails to have the appropriate emotions and motivations given p. For example, if we recognize φing as racist, but φing does not make us feel sad or angry, we fail to understand that φing is racist—even if we have all the cognitive abilities featured in Hills' view (Callahan, 2018, 452).

But I think that we can keep a traditional cognitive-centric account of understanding, like Hills', and still make sense of the connection between understanding and morally appropriate affects—and, by extension, the relationship between lacking understanding and affective alienation when it comes to moral matters. That's because, even on an account like Hills', it's plausible that understanding enables and promotes morally appropriate affects, thus preventing affective alienation.

For one, as Enoch points out, "typically, the emotional responses that are closely connected with moral judgments are responsive to the morally wrong… *de re*, and not merely *de dicto*…" (Enoch, 2014, 28). When we respond with sadness or anger to civilian suffering that results from a military invasion that fails to take proper precautions to protect non-combatants, our sadness or anger is

[6]For a development of this view of moral worth, see Arpaly (2002a, 2002b) and Markovits (2010); for a discussion of how these views relate to moral deference, see Nickel (2001) and Callahan (2018).

[7]For developments and discussions of The Affective Alienation Explanation, see Howell (2014), Enoch (2014), Fletcher (2016), and Sliwa (2025).

[8]I use the term "affect" broadly to cover both emotional and motivational states.

elicited *not* by our viewing the invasion as "morally wrong" but rather by recognizing the preventable gratuitous suffering of innocent human beings that makes it wrong. Similarly, when Coco feels frustrated by others who choose to drive SUVs around town, her frustration is not likely elicited by her viewing their actions as "morally wrong" but rather by recognizing their contributions to the severe collective harm of climate change (that, perhaps, she and others have made substantial sacrifices to avoid contributing to). Moral understanding on traditional accounts constitutively involves cognitive engagement with morality *de re*—that is, the moral reasons that tend to elicit the morally appropriate emotions and motivations when recognized.

A second point of connection between moral understanding and moral affects arises from "the reflective and experiential practices" (Callahan, 2018, 454) that we engage in when we seek out understanding. The idea here is that the same practices that foster understanding about moral matters are the very same practices that promote morally appropriate affects. For example, in order to better understand the moral horrors of the Holocaust, we might read *Night* by Elie Wiesel or Anne Frank's diary; or watch *The Boy in the Striped Pajamas*; or visit Auschwitz or a Holocaust memorial museum. These practices not only tend to foster understanding, but also elicit sadness or anger at the relevant moral horrors and injustices, and a desire to prevent or rectify other horrors.

Relatedly, understanding constitutively involves having certain abilities—such as explaining why p in your own words—and exercising or cultivating these abilities involves processing or "sitting with" the relevant information in ways that allow it to sink in, so to speak. It is when the information sinks in that it is poised to stir up fitting emotions and motivations.

4. Understanding & Cases of "Rational" Prejudiced Beliefs

My goal in this paper is to highlight a set of non-moral propositions that, I take, are especially morally important to understand for reasons that are similar to those that defenders of The Understanding Explanation (can) give to defend The Relative Importance Claim. The propositions are those featured in cases like The Server, which have received a lot of attention recently in the ethics of belief.

In these cases, the propositions in question constitute the content of beliefs that seem to be both prejudiced and supported by the believer's evidence. Here is another example case, which has the same basic structure as The Server.

> **The Teacher** Stacy is a fifth-grade teacher at a public elementary school. It is the first day of school, and she is meeting her students for the year for the first time. A new student, Jenna, walks in. Stacy knows that, on average, girls score lower than boys on standardized math exams. In light of this, Stacy comes to believe that Jenna probably scored lower than average on last year's statewide standardized math exam.

The belief in question is Stacy's belief about Jenna, which she infers from her knowledge of a demographic fact about girl students. Call the belief about Jenna *the inferred belief*, and the group-level belief *the demographic belief*. Moral encroachers contend that Stacy's inferred belief about Jenna (like Spencer's inferred belief about Jamal) is prejudiced, even though it is supported by evidence that is sufficient to rationalize beliefs according to traditional purist standards. But admitting that prejudiced beliefs can be epistemically rational is troubling for a couple of reasons. For one, it is characteristic of prejudiced beliefs that they are epistemically defective; second, there are problems with accepting that morality can forbid a belief that epistemology permits.[9] By accepting a revisionary view of epistemic rationality that makes room for moral considerations —moral encroachers argue—we can maintain that prejudiced beliefs are epistemically defective,

[9]See Basu and Schroeder (2019) for discussion of these points.

and avoid these troubling conflicts between epistemology and morality, all while rendering the intuitive verdict that Spencer's belief about Jamal and Stacy's belief about Jenna are prejudiced. According to moral encroachment, these beliefs are epistemically irrational along with being morally wrong.[10]

In order to make my case that understanding the non-moral propositions in cases like The Server and The Teacher is especially morally important, I need to first say a bit about what understanding would amount to in these cases. And, in order to do that, I need to say a bit about the socio-epistemic landscape in play: in these cases, a believer S (e.g., Spencer, Stacy) knows that most members of some historically marginalized group G (e.g., Black Americans, girls/women) has some negative or relatively diminishing property X (tips or scores worse than average, respectively) and then comes to believe, inferentially, that J (e.g., Jamal, Jenna)—who is a member of G—is probably X. The belief in question is the inferred belief about J.

So, what would it take for the believers in these cases to understand the propositional content of the belief in question? Such an understanding would amount to having cognitive mastery over the reasons why it is the case that Black diners tip less than average and that girl students score worse than boy students on math exams (and, thus, why it is that a randomly selected Black diner is likely to tip less than average or why it is that a randomly selected girl student likely scored worse than average on a math exam).

In these cases, there is a culturally prominent *mis*understanding of the subject matter: according to this misunderstanding, X is common among G because of some inherent flaw, vice, or deficiency characteristic of G. In The Server, the relevant misunderstanding would be that Black diners tip less than average because Black diners tend to be, say, less generous. In The Teacher, the relevant misunderstanding would be that girls tend to score worse than boys because girls and women tend to be less intelligent and rational or else worse at qualitative reasoning. The prominence of these misunderstandings may be partly explained by the tendency to "essentialize" socially marginalized groups, where a group is essentialized "just in case its members are viewed as sharing a fundamental nature that causally grounds a substantial number of their outwardly observable properties" (Leslie, 2017, 409). For example, being "hyperemotional" or "dramatic" or "intuitive"—and, thus, prone to irrationality—may commonly be seen as part of the essence of womanhood and thought to causally explain the stubborn gender achievement gap in math.

But, again, this is a *mis*understanding. It is false that Black diners tip less than average or that girl students score worse than average on math exams because of some inherent negative feature characteristic of Black Americans and females, respectively. Rather, the reasons why X is common among G trace back to the fact that G has historically been and continues to be marginalized, subject to a legacy of oppression. When it comes to the difference in tipping patterns between Black and white Americans, there are two prominent explanations in the sociological literature: one, servers systematically discriminate against Black diners, and two, Black diners tend to be less familiar with the tipping norms that evolved in higher-income, white-dominated spaces.[11] When it comes to the gender achievement gap in math, many sociologists have converged on a sociocultural explanation, attributing the gap to "various societal factors…. including differential treatment by others… stereotype threat… differential treatment, parental investment, or expectations…. attitudes towards competition… and societal gender inequity" (Gevrek et al., 2020, 1). One prominent sociocultural explanation is that girls are negatively impacted by gendered stereotypes impugning their mathematical abilities, which tend to be triggered in competitive testing environments (Niederle & Vesterlund, 2010).

[10]See Section 7.a for further discussion of moral encroachment.

[11]See Brewster and Mallison (2009) for an overview and critical discussion of the relevant sociological literature on racial disparities in tipping patterns.

So, in the relevant cases, an understanding of the propositional content of the belief in question —that J is likely X—involves cognitive mastery of reasons constituting the correct explanation of why X is common among G and, thus, why J is likely X. The correct explanation ultimately traces back to facts related to the historical and continuing oppression of G (e.g., facts about systematic discrimination, segregation, relatively low incomes, and negative stereotypes and gender norms). Culturally prominent explanations that attribute the commonness of X among G to an inherent deficiency or flaw characteristic of G are false. But, given their cultural prominence, these incorrect explanations are ready to fill in explanatory gaps whenever someone believes that X is common among G but does not understand why this is the case.

Before moving forward, it needs to be emphasized that the understanding in these cases is *non-moral*: the reason why X is common among G is an empirical or descriptive matter, not a normative one. This is worth emphasizing because it might be thought that—if the correct explanation of why X is common among G centrally involves a social injustice—then the relevant understanding is thereby moral. But this does not follow: just because there are salient normative implications of some non-moral truth does not transform propositions about that truth (or the related understanding) into moral propositions (or moral understanding).

To see this, consider a clear example of a non-moral proposition and related understanding. Suppose my friend is very late to our long-planned lunch date, after promising to be on time. The reason why he is late is an empirical matter, not a moral one. That said, the reason why he is late may have moral implications. Suppose there was a car accident on the only route to the restaurant, and he was stuck in unexpected traffic. In this case, my friend has a legitimate excuse for being late that mitigates the moral blame he would otherwise be subject to. In fact, it might be said that the reason why my friend is late is ultimately because of something morally excusable. Nevertheless, understanding why my friend is late is non-moral, even if the reason why he is late ultimately involves facts that are related to his having a moral excuse. Similarly, understanding why X is common among G is a form of non-moral understanding, even if the reason why X is common among G ultimately involves facts related to social injustices suffered by G. Understanding why p is non-moral to the extent that p is a non-moral proposition, even if p's truth has salient moral implications.

5. The Server, The Teacher, & the Moral Significance of Understanding

My contention is that it is especially morally important to understand the propositional content of the beliefs in question in cases like The Server and The Teacher. What's more, the reasons why it is especially morally important to understand these propositions are similar to those reasons that defenders of the Understanding Explanation (can) give in order to defend The Relative Importance Claim: understanding these propositions is central for morally right actions in a range of contexts; may be essential to morally worthy action; and enables and promotes emotions and motivations that are fitting in light of morally significant features of these cases.

Further, the non-moral understanding in cases like The Server and The Teacher serves a morally important role that moral understanding does not obviously serve in standard cases featured in the moral testimony literature—namely, it blocks S from having the culturally prominent *mis*understanding that attributes X's commonness among G to some inherent negative trait characteristic of G. For example, if Spencer understands that it is ultimately because of anti-Black racism that Black diners tend to tip less than average, then he will not mistakenly think that it is because of some inherent flaw or vice characteristic of Black folks.[12] The relevant *mis*understanding is morally pernicious for a series of familiar reasons: it promotes morally wrong actions, frustrates morally worthy action, and invites affects that are inappropriate in light of the morally significant features of

[12]At least, he is much less likely to have the common misunderstanding if he already understands the correct explanation.

the case. So, I will argue that understanding why X is common among G is not only important for each of the three familiar reasons explored in the moral testimony literature but also, in part, because it blocks a pernicious non-moral *mis*understanding.

5.a A simple argument

First, though, let us consider a simple argument for the moral importance of non-moral understanding that appeals to the metaphysical relationship between moral and non-moral facts. Contingent moral facts, this argument goes, are always at least partially explained by contingent non-moral facts. The moral fact that it is wrong for Coco to buy an SUV, for example, is explained by mechanical facts about how much gas SUVs use relative to other vehicles and scientific facts about how gas contributes to climate change. Given this, and given that understanding p entails grasping the reasons why p, then understanding a contingent moral fact entails understanding the non-moral facts that explain it. And if understanding some moral fact entails certain forms of non-moral understanding, then *of course* those forms of non-moral understanding are morally important for the same reasons that moral understanding is.

This simple argument offers a simple account of why it is morally important to understand why X is common among G in cases like The Server and The Teacher: this non-moral understanding is entailed by some salient form of moral understanding.

The simple argument, however, does not work. That's because understanding contingent moral facts does not always entail understanding the non-moral facts that explain them. Yes, understanding why it is wrong to buy an SUV entails grasping non-moral facts about gas and climate change and appreciating those facts as reasons that explain the wrongness of the action in question. But this is not the same as *understanding* those non-moral facts. Understanding the non-moral facts would involve grasping the mechanics of why SUVs use more gas than other vehicles and the science of why using gas contributes to climate change. Coco does not need this sort of technical understanding to understand why it is wrong for her to buy an SUV. So moral understanding—even when it involves grasping non-moral facts—does not entail non-moral understanding.[13]

The arguments in the rest of this section show that the moral importance of understanding why X is common among G is partly, but not solely, dependent on its relationship with moral understanding—in particular, sometimes the relevant form of non-moral understanding is *contingently* necessary for understanding why some salient action is wrong. These arguments reveal that the nature of the relationship between moral and non-moral understanding is more complicated than what is suggested by the simple argument just considered. They also show that non-moral understanding is morally important in its own right—that is, in ways not explained by its relationship with moral understanding. The upshot is that the moral importance of non-moral understanding is not merely parasitic on the moral importance of moral understanding.[14]

5.b Morally right action

The first reason that defenders of the Understanding Explanation give to defend the Relative Importance Claim is that moral understanding is distinctly instrumentally important insofar as it enables morally right action in a wide variety of contexts where related moral matters are in play. In cases like The Server and The Teacher, understanding the relevant non-moral proposition also enables morally right action in the cases at hand as well as a variety of related contexts. In addition, it prevents morally wrong actions.

[13]More generally, understanding why p (where p is because of q) does not entail understanding why q.

[14]Special thanks to an anonymous referee and the editor for pushing me to consider the simple argument and clarify the relevant points.

The underlying idea here is that which course of action is morally right or preferable depends not only on moral truths but also on garden-variety non-moral facts. By extension, forming true beliefs about which is the morally right or best option often depends on non-moral understanding. For example, a professor may need to understand why a student failed to show up for the midterm exam in order to form true beliefs about the morally appropriate course of action. These true beliefs about the morally appropriate course of action, in turn, enable morally right and desirable action. If the student just tested positive for COVID-19, it would be morally appropriate for the professor to offer a makeup exam (and morally inappropriate to give the student a 0).

It is not especially revelatory that understanding some one-off non-moral facts enables and promotes morally right action. What is revelatory is that there is a whole class of such facts (akin to the whole class of moral facts) the understanding of which reliably and systematically enables and promotes morally right action given the socio-epistemic landscape in play. Specifically, understanding why it is that X is common among G in cases like The Server and The Teacher enables (and may be necessary for) acting in a morally appropriate way towards an individual member of G who is likely X, not only in the context of the case but also in other contexts where X being common among G is salient. For example, Stacy may need to know why it is that girls score worse than average on standardized math exams in order to help her girl students reach their full academic potential: it would be morally good for Stacy to actively combat the gendered stereotypes around math that tend to be triggered in competitive testing environments. Or, to borrow an example from Brinkerhoff (2025, 967), a member of an admissions committee at an elite STEM-focused charter school may need to have the same understanding in order to give the dossier of female applicants a fair evaluation. After all, if the gender gap in standardized math exam scores does not reflect differences in mathematical ability, then it would be unfair to assume that a boy student with a higher score than a girl student has more mathematical promise.

What's more, understanding in these cases prevents morally wrong actions by blocking S from accepting the culturally prominent *mis*understanding of why X is common among G. The relevant *mis*understanding, I want to suggest, makes believers susceptible to morally wrong actions. That's because actions that are, in fact, morally wrong may be justified to some extent if the misunderstanding were correct. For example, it is morally wrong for Stacy to steer her female students away from fields and careers that emphasize quantitative skills like engineering, computer coding, and finance. But this may be justified if (counterfactually) women are inherently less competent when it comes to qualitative reasoning.

So understanding the propositions in these cases is morally important because it allows believers to form the true beliefs about which course of action is morally right or appropriate in the immediate context, as well as relevantly similar contexts. These beliefs, in turn, enable morally right action. In addition, having the relevant understanding prevents morally wrong actions by blocking the culturally prominent *mis*understanding, which may justify or otherwise support actions that are, in fact, morally wrong.

5.c Morally worthy action

The second reason that defenders of The Understanding Explanation give to defend The Relative Importance Claim is that moral understanding is necessary for morally worthy action, at least according to some views of moral worth that require acting out of concern for morality *de re*. In cases like The Server and The Teacher, understanding the relevant non-moral proposition may also be necessary for morally worthy action, if those same views about moral worth are correct.

Consider some of the morally wrong ways that the believers may act in the relevant cases, and how we might evaluate their decision not to act in those ways. For example, Stacy may discourage Jenna from pursuing STEM careers and fields. These actions are wrong at least in large part because they are unfairly discriminatory: it is unfairly discriminatory—it is racist and sexist—to treat people

differently on the basis of their race or sex unless their race or sex gives us a legitimate reason for the relevant form of differential treatment.[15]

In these cases, whether J's race or sex gives S a legitimate reason for differential treatment may hinge (at least in part) on whether likely being X is grounds for differential treatment. So, the question arises: is likely being X grounds for the relevant form of differential treatment toward J (and other members of G)? Does the fact that Jenna's sex is associated with lower math exam scores give Stacy a legitimate reason to steer Jenna away from STEM?

The answer to these questions seems to depend on why it is that X is common among G (and, thus, why J is likely X). Whether the fact that someone's gender corresponds to lower math test scores is a legitimate reason for differential treatment depends on why it is that their gender corresponds to lower math test scores in the first place. If it's because females are just inherently worse than boys at qualitative reasoning, then it could make sense for Stacy to steer her girl students away from STEM. If, however, it is because of gendered stereotypes that tend to be triggered in competitive testing environments, then the fact that girls tend to score worse on math exams would clearly *not* give Stacy a legitimate reason to steer female students like Jenna away from STEM.

The upshot is that believers in these cases need to understand why it is that X is common among G in order to understand why being a G is not a legitimate reason to treat J differently in light of X[16] and, thus, why treating J and other individual Gs differently in light of X would be unfairly discriminatory.

This has implications for the relevant views of moral worth, according to which morally worthy action must flow from concern for morality *de re*. On these views, in order for Stacy's decision to not steer Jenna away from STEM to be morally worthy, this decision must be based on concern for those reasons that make it morally wrong to steer female students away from STEM—that is, her decision must flow from a desire to not unfairly discriminate against someone on the basis of their sex. And, in order for Stacy to recognize that steering Jenna away from STEM would be unfairly discriminatory despite the fact that girls tend to score lower on math exams than boys, she must have a grasp on the reasons behind the gender gap in mathematical achievement.[17] So, given the relationship between the particular forms of moral and non-moral understanding in play, the relevant non-moral understanding is essential for morally worthy action.

It is worth noting here that if the culturally prominent *mis*understanding of why X is common among G were correct—if X is common among G because of some inherent flaw or shortcoming characteristic of G—then likely being X may give S a legitimate reason for differential treatment of J and other members of G. So understanding why X is common among G is also important for morally worthy action because it blocks the common *mis*understanding according to which the morally wrong actions in question (e.g., steering Jenna away from STEM) may *not* be unfairly discriminatory. It is, of course, possible that S would decide not to do the morally wrong action even if S accepted the culturally prominent misunderstanding; but—since accepting this misunderstanding prevents S from recognizing the actions as unfairly discriminatory—this decision would not be morally worthy since it would not be made out of concern for morality *de re*.

To sum up: if morally worthy action requires acting for the right reasons *de re*, then it seems plausible that S must have a grasp of what those reasons are in order to act in a morally worthy way.

[15]I focus on unfair discrimination, but these actions may be wrong for a variety of reasons. For example, it may be generally wrong to act in ways that further disadvantage a member of a historically marginalized group, and it may be generally wrong to act in ways that disadvantage someone when the action is justified only by statistical information about their social group rather than individual-specific information. The following argument can still run given these other reasons that make the relevant actions wrong.

[16]By "treating J differently in light of X," I mean treating J differently in light of the fact that J is likely X given that J is a G and most Gs are X.

[17]I have focused on an example of morally worthy avoidance of wrong action. But a similar point can be made about taking a morally worthy *right* action.

In cases like The Server and The Teacher, this, in turn, requires S to understand non-moral facts about why X is common among G.[18]

5.d Moral affects

In Section 3, we explored a third reason that defenders of The Understanding Explanation can appeal to in order to defend The Relative Importance Claim: moral understanding enables and promotes morally appropriate affects. Here, I develop a parallel argument: understanding why X is common among G enables and promotes the affective responses that are fitting in light of morally significant features of cases like The Server and The Teacher.

To pave the way for this argument, it is helpful first to get a grasp on what affects would be morally appropriate in the cases in question. The appropriate affects are those emotions and motivations that are fitting given p, and also given the reasons why p is true. So, in the cases in question, the appropriate affects would be the ones that are fitting not only given X's commonness among G (and, by extension, J's likely being X) but also given that X is common among G ultimately because of facts related to social injustices suffered by G. They are *morally* appropriate to the extent that they are fitting in light of the morally significant features of the case—especially that it is because of an injustice suffered by G that X is common among G.

The appropriate affects in these cases may include anger at the injustice or barriers suffered by G that caused X's commonness among G in the first place; a motivation to avoid acting in ways that exacerbate the relevant injustice or that further disadvantage G as a group as well as J and other members of G; a motivation to help remove barriers to opportunity and success that hold G back, to the extent that one can; pride or joy when J or another G succeeds or flourishes despite these barriers; sadness when members of G are harmed or disadvantaged, and frustration at our inability to fix the institutional or systematic roots of disadvantage. In contrast, the morally *in*appropriate emotions and motivations include resentment or frustration at J for likely being X and a motivation to act in ways that disadvantage J and other members of G. Notably, these inappropriate affects would perhaps be (counterfactually) appropriate if X is common among G because of an inherent vice or shortcoming characteristic of G (for which, it may seem, it is proper to hold Gs accountable).

In Section 3, we considered two points of connection between moral understanding and morally appropriate affects: first, moral understanding constitutively involves engagement with morality *de re*, which tends to elicit morally appropriate affects and, second, the practices that foster understanding also foster morally appropriate affects. Both points also explain how understanding why X is common among G enables and promotes morally appropriate emotions and motivations in cases like The Server and The Teacher, and blocks the morally inappropriate ones.

Regarding the first point: understanding why X is common among G does not constitutively involve engagement with morality *de re*. However, in cases like The Server and The Teacher, the moral importance of this understanding arises partly in relation to how S treats J in light of X: Spencer, for example, needs to decide how to treat Jamal in light of his belief that Jamal likely tips less than average. As discussed above, understanding why X is common among G is central to understanding why certain salient actions are morally wrong—namely, they are unfairly discriminatory (racist and sexist). And understanding *this* amounts to engagement with morality *de re*—understanding that the relevant actions would be racist and sexist is the type of thing that is prone to elicit morally appropriate emotions and motivations, including a motivation to avoid doing the relevant action, guilt if one does the relevant action, anger at others who do the relevant action. So

[18]Note: the argument here is *not* that it is always or metaphysically necessary to understand the non-moral facts that explain some contingent moral fact in order to achieve moral understanding of that moral fact. Rather, the argument is tailored to the particular propositions and socio-epistemic landscape in play: it is (contingently) necessary to understand why X is common among G in order to achieve moral understanding of why certain salient actions are morally wrong (because unfairly discriminatory).

understanding why X is common among G is central to relevant forms of moral understanding—namely, understanding why certain salient actions would be morally wrong—which does constitutively involve engagement with morality *de re*. And engagement with morality *de re* is the sort of thing that promotes the morally appropriate affects.

Regarding the second point: in order to seek out an understanding of why X is common among G, one will engage in various reflective and experiential practices, and these practices tend to stir up morally appropriate emotions and motivations. After all, the correct explanation of why X is common among G centrally involves morally significant facts related to injustices suffered by G, and achieving cognitive intimacy with these facts through the relevant practices is prone to elicit the appropriate affects.

Just consider Stacy. In order to understand why girls tend to score worse than average on math exams, she might, of course, do a Google search and read about the sociology studies that come up first. But—in order to acquire the mastery of the relevant reasons characteristic of understanding, which goes above and beyond simply knowing that p because q—Stacy might also talk to girls about the ways they think and feel about their math skills and those of their male peers; she might seek out stories of women in STEM, and the barriers they have faced in male-dominated and -coded environments; she might read *Lessons in Chemistry* or other fiction books that center on women in STEM. It is plausible these practices tend to elicit the morally appropriate emotions and motivations.[19] After reading a first-hand account of barriers faced by women in STEM, and meditating on the visceral weight of gendered expectations and stereotypes, Stacy is liable to feel angry about the way that women in STEM have been treated, and motivated to actively combat gendered stereotypes related to mathematical ability. What's more, engaging in these experiential and reflective practices invites Stacy to sit with the relevant information, giving it time to "sink in," whereby it is poised to elicit the relevant affects.

6. The Contingency Objection & Reply

So far, I have argued that—in cases like The Server and The Teacher—non-moral understanding about why X is common among G (and, thus, why J is likely X) is especially morally important. What's more, the reasons why it is especially morally important are similar to those that defenders of The Understanding Explanation (can) give to support The Relative Importance Claim: understanding why X is common among G is central to morally right and worthy actions, and it enables and promotes morally appropriate affects; what's more, it blocks a culturally prominent but pernicious *mis*understanding.

In this section, I seek to defend this argument from a potential objection. The objection is that the relevant non-moral understanding in cases like The Server and The Teacher is morally important only in a *contingent* sense. Because of this, the objection goes, its moral importance is not as robust as I have been suggesting.

There are two strands to this objection. On the first strand, achieving the relevant form of understanding requires cooperation from our socio-epistemic environment. Imagine, for example, that S's socio-epistemic environment is one in which racial or gender essentialism is widely endorsed, even among those considered "experts." In this situation, it would be morally worse for S to seek out an understanding of why X is common among G than it would be for S to, say, suspend judgment on the matter. On the second strand, it is not the understanding itself that is morally important, but rather the accompanying moral goodies—and, per the first strand of the objection, this accompaniment is neither automatic nor guaranteed.

[19]This is an empirical claim that would need to be confirmed by empirical studies. See Nussbaum (1995) for a strong case for the moral value of engaging with literary fiction which appeals in part to the morally significant emotions it elicits—she argues that engaging with fiction is a training ground for our moral sentiments.

In response to the objection's first strand: it is first worth noting that the same charge can be levied against the moral importance of moral understanding. Suppose Coco's socio-epistemic environment is one in which Randian egoism is widely endorsed by those considered ethical experts. If Coco seeks to understand why it is wrong to drive an SUV around town in an environment like this, she would likely be led astray and go on to form harmful and false beliefs; in this environment, it would be morally better for Coco to suspend judgment.

More importantly, though, my claim has been that, in cases like The Server and The Teacher, it is the understanding itself that is morally important, and not (primarily) the seeking out of the understanding. It is by virtue of understanding why X is common among G that S is in a position to do the morally right and worthy actions and develop morally appropriate affects, and is blocked from assuming the relevant misunderstanding. That said, given the moral importance of this understanding, it is plausible that many of us do have a strong moral reason to seek it out, though, of course, this reason might be defeated by, say, considerations about the quality of the socio-epistemic environment. I take it that this (defeasible) moral reason to seek out an understanding flows from a general moral duty to approach and handle morally high stakes issues with extra care. The duty of extra care may also generate moral reasons for S to double-check the evidence or actively bear in mind the possibility that J may not be X.

In response to the objection's second strand: yes, having the relevant form of non-moral understanding is not always accompanied by the standard moral goodies. Understanding why girls tend to score worse on math exams, for example, may *not* motivate Stacy to take morally good actions, or prompt her to have morally appropriate affects, given certain contingent facts about her psychology. Stacy might be distracted, tired, and mentally overburdened by teaching in an under-funded school; alternatively, she simply might not care much about addressing the gender gap in math achievement.

But it is also true that understanding moral propositions is subject to similar efficacy limitations.[20] Understanding why it is wrong to drive an SUV around town may not effectively lead Coco to take the morally right actions or prompt the appropriate affects. After all, Coco might be distracted, tired, overburdened by hardships of daily life, or skeptical about the moral demands imposed on individuals by collective action problems; alternatively, her attention may be directed towards other salient moral demands like helping the unhoused or rearing her own children; or she may just not care much about how her private choices contribute to collective harms.

I contend, then, that the moral importance of the relevant form of non-moral understanding in cases like The Server and The Teacher is not *worryingly* contingent, at least no more worryingly contingent than the moral importance of moral understanding. So, to the extent that we accept the moral importance of moral understanding, the contingency objection gives us no extra reason to worry about the moral importance of the non-moral understanding in the cases at hand.

7. Concluding Thoughts: Upshots for Ethics of Belief

To conclude, I want to consider the upshots of my argument and its implications for the current debate in the ethics of belief about moral encroachment on epistemic rationality.[21]

7.a The challenge from moral encroachment

The debate is between defenders of moral encroachment and purists. According to moral encroachers, moral considerations have a say in what is epistemically rational to believe by setting

[20]This is so at least if we assume some version of anti-internalism about moral motivation when it comes to grasping moral reasons.

[21]I here lean on the version of moral encroachment developed by Basu (2019a, 2019b, 2019c) and Basu and Schroeder (2018). For arguments against moral encroachment, see Gardiner (2018) and Begby (2021).

the evidential threshold for epistemic rationality. When the moral stakes for believing that p are high, believers tend to need stronger evidence to rationally believe that p. In contrast, purists contend that the evidential threshold for epistemic rationality holds steady across contexts, regardless of the moral stakes.

Moral encroachers appeal to cases like The Server and The Teacher to argue against purist accounts and motivate their own revisionary account of epistemic rationality. One challenge for purist accounts posed by these cases is that they cannot easily accommodate the intuitive verdict that the beliefs in question—Spencer's belief about Jamal and Stacy's belief about Jenna—are prejudiced. After all, if prejudiced beliefs must be epistemically defective (as is widely accepted) and if the beliefs in question are supported by evidence that is sufficient to rationalize beliefs according to purist standards (as they seem to be), then the beliefs in question are not prejudiced. So, if they really are prejudiced, we either must reject the plausible claim that prejudiced beliefs are epistemically defective or else reject purist standards that entail the beliefs in question are epistemically rational. Moral encroachers, of course, take the latter option. Call this *the challenge from moral encroachment*.

I want to suggest that—if understanding the propositional content of the beliefs in question is especially morally important—purists can appeal to this moral importance to resist the encroacher's challenge.

7.b The purist response: a sketch

The challenge begins with the intuition that Spencer's belief about Jamal and Stacy's belief about Jenna are prejudiced—or, at the very least, morally problematic. One type of purist response to the challenge is to acknowledge the force of this intuition but locate the moral problem somewhere other than the belief itself. For example, Brinkerhoff (2021) argues that there is something morally problematic about the attention patterns manifested in occurently having the beliefs in question. It seems fishy, for example, that the first thing that pops into Stacy's mind when Jenna walks into her classroom on the first day is that she likely underperformed on last year's math exam.

The purist response I have in mind has a similar form: in cases like The Server and The Teacher, the moral problem is not that the belief itself is prejudiced, but rather the lack of a morally important form of non-moral understanding. Given a face-value reading of the cases at hand, the beliefs in question are held in isolation, unaccompanied by an understanding (or, presumably, any *mis*understanding) of why X is common among G. Because S lacks the relevant form of understanding, S is not epistemically well positioned—and may be unable—to take the morally right and worthy actions, or have the morally appropriate affects. What's more—if S has the beliefs in question but does not believe anything at all about why X is common among G—then S is at substantial risk of assuming the pernicious misunderstanding, which, given its cultural prominence, is ready and waiting to fill in the existing explanatory gap.

Pointing to the moral importance of understanding[22] allows purists to accommodate the verdict that there is something morally problematic with the S's broader doxastic state without having to say that the moral problem inheres in S's belief about J itself. And if the moral problem does not inhere in the belief itself—but somewhere else nearby—then the pressure on purist accounts of epistemic rationality is relieved. Although there is something morally problematic about S's broader doxastic state, S's belief about J is not, itself, morally wrong (because prejudiced), and so we keep the purist account of epistemic rationality without having to give up the claim that prejudiced beliefs must be epistemically irrational.

[22]Here, and for the remainder of the paper, "understanding" should be taken to refer to the relevant form of non-moral understanding of why X is common among G.

This purist-friendly attempt to accommodate the relevant intuitions is vindicated by what I take to be the intuitive verdicts in variations on The Server and The Teacher in which S is explicitly portrayed as understanding why X is common among G (and, thus, why J is likely X). Consider the following variation on The Teacher.

> **The Informed Teacher** Stacy knows that, statistically, girls tend to score lower on standardized math exams than boys and infers that Jenna likely scored less than average on last year's statewide standardized math exam. Stacy has recently done a lot of research about the gender gap in mathematical achievement. In addition to the relevant statistical information, Stacy knows that women and girls have been historically characterized in ways that impugn their mathematical abilities. From her research, Stacy knows that the gender disparity in math performance is explained – not by a lack of rationality or analytic prowess in girls and women – but rather by the ways that negative stereotypes about women and math negatively affect girls' math performance.

In this variation, my intuition that Stacy's belief about Jenna is sexist has faded drastically. For those with similar intuitions, this vindicates the purist-friendly response that the lack of understanding is the source[23] of the moral problem in the original cases. According to this response, the intuition that S's belief about J in the original cases is prejudiced is tracking a genuine moral problem with S's doxastic state, but the problem is the lack of understanding and not the belief itself. After all, Informed Stacy has the belief in question in The Informed Teacher: so, if the moral problem really inheres in the belief itself, then the belief should still seem prejudiced in The Informed Teacher, too.

7.c Important clarification: demographic versus inferred belief

Some clarifications are in order about the purist response sketched above. Cases like The Server and The Teacher involve two beliefs: a group-level *demographic* belief (that X is common among G) and an individual-level *inferred* belief (that J, as a G, is likely X). Moral encroachers are primarily concerned with the inferred belief—this is the belief that they deem prejudiced and wrong. But the relevant form of non-moral understanding I have been focusing on is about the group-level trends encapsulated most clearly by the demographic belief. There remain questions about the connection between the understanding of the group-level demographic information and the way it bears on the normative status of the inferred belief, and whether the moral importance of the relevant form of non-moral understanding transfers to the inferred belief.

There are a few things to say.

First, it is worth re-emphasizing that, in cases like The Server and The Teacher, understanding why X is common among G also serves, derivatively, as an understanding of why J is likely X, and is morally important to the treatment of J. For example, it is important that Spencer understands why Black diners tend to tip less because it puts him in a position to do the right and morally worthy actions, and to have the appropriate affects, *when it comes to Jamal in particular*. What's more, this understanding prevents Spencer from essentializing Jamal, in particular, in a negative way—it blocks him from assuming that Jamal likely tips less than average because he, as a Black American, is inherently less virtuous. So, although the understanding is primarily of demographic trends, the moral importance of this understanding trickles down in a meaningful way to the inferred belief.

That said, some moral encroachers contend that it is wrong to infer things about individuals on the basis of information about their apparent social group, regardless of whether the inferred belief is couched in the relevant understanding. Rima Basu, for one, argues that inferring that J is likely X from demographic information about G involves treating J as an object, and, thus, is wrong. Basu's

[23]Or at least one significant source – see Section 7.c for further discussion.

view entails that Spencer's belief about Jamal constitutes a way of "looking at (Jamal)… not as a person, but as an object that is determined by causal laws, as something whose behaviour is to be predicted" (Basu, 2019a, 922). We ought to treat people as people, not as objects, and this holds whether the inferences in question are surrounded by an understanding of the relevant demographic trends. The worry for the purist response, then, is that understanding why X is common among G does not vitiate the moral ills of stereotyping J on the basis of G-related demographic information by making the relevant inferences.

My view—and the view I think purists endorsing the response sketched in 7.b should take—is that it is not inherently wrong to infer things, even negatively-valenced things, about individuals on the basis of demographic information about their apparent social group, *whether or not the resulting inferred beliefs are couched in the relevant form of understanding*. For one innocuous example, there is nothing at all wrong with inferring that J is probably sleep-deprived from knowledge that most parents of newborns are sleep-deprived and that J has a newborn. Such inferences are indispensable to our cognition and ability to navigate the world. And the same generally holds for negatively-valenced properties: although there may be other moral problems in the cognitive vicinity, there is nothing inherently wrong, for example, with inferring that Jenna likely scored lower than average from knowledge that most girls scored lower than average and that Jenna is a girl.

The upshot: it is *not* my view that it is permissible to have such inferred beliefs only when those beliefs are couched within the relevant form of understanding, and impermissible otherwise. My view is a bit less precise, and relies on categories of moral appraisal—namely, moral importance/significance, and being morally problematic—that do not fit comfortably within a standard deontic framework. Thankfully, these other categories of moral appraisal are familiar. A virtue ethicist might say that it is morally important for me to have the virtue of courage—and insist that something very important is missing if I am not courageous—without thereby being committed to the view that it is morally impermissible not to be courageous. Alternatively, it might be said that it is morally problematic for medical students to forgo studying bioethics during med school, without being committed to the view that it is wrong for them not to take Bioethics 101.

This is the general tenor of my purist suggestion. In cases like The Server and The Teacher, it is morally important to have the relevant form of non-moral understanding alongside the inferred (and demographic) beliefs. When these beliefs are present while the understanding is absent (even when there is no *mis*understanding), something substantially morally important is missing, just like when one goes through medical school without studying bioethics, or life without developing the virtue of courage.

That said, there may be other significant moral problems lurking in S's broader cognitive state even if S understands why X is common among G: S's attention patterns manifested in occurently having the inferred belief may be morally problematic, or S may still be inclined to act in harmful ways towards J and other Gs despite the understanding. So understanding is no panacea. The moral-epistemic dynamics in cases like The Server and The Teacher are multifaceted, and the moral importance of the relevant form of non-moral understanding captures only one facet—albeit, I would add, a significant one.

7.d In sum

Part of what makes the beliefs in question in cases like The Server and The Teacher so philosophically interesting is that—although their propositional contents are non-moral and merely descriptive of the way the world is—they describe parts of the world that have been warped by sexism and racism. Basu has even claimed that, because of this warping, the evidence is "stacked in favor of" prejudiced beliefs (Basu, 2019b, 2513–2514). The purist can agree with Basu that prejudice's warping of our world is normatively significant by insisting that it is especially morally important to understand propositions that capture parts of the world that have been so warped. But the purist can also insist, *contra* Basu and other moral encroachers, that it is not racist and sexist to believe the

propositions that capture a bit of warped reality; rather, it is morally problematic to believe these propositions without also understanding the roles played by racism and sexism in rendering the world the way it is.

Acknowledgments. For very helpful feedback and discussion on this paper, I would like to thank the participants at the 2025 Philadelphia Normativity Philosophy Conference, especially Nicole Dular and Sumeet Patwardhan. I would also like to thank Ripley Stroud, Samuel Carlsson Tjernström, and the two anonymous referees for the *Canadian Journal of Philosophy*.

Anna Brinkerhoff is an Associate Professor of Philosophy at Concordia University in Montreal, Quebec. Her research focuses on questions at the intersection of ethics and epistemology.

References

Arpaly, N. (2002a). *Unprincipled virtue: An inquiry into moral agency*. Oxford University Press. https://doi.org/10.1093/0195152042.001.0001

Arpaly, N. (2002b). Moral worth. *Journal of Philosophy*, 99(5), 223–245. https://doi.org/10.2307/3655647

Basu, R. (2018). Can beliefs wrong? *Philosophical Topics*, 46(1), 1–17. https://doi.org/10.5840/philtopics20184611

Basu, R. (2019a). What we epistemically owe to each other. *Philosophical Studies*, 176(4), 915–931. https://doi.org/10.1007/s11098-018-1219-z

Basu, R. (2019b). The wrongs of racist beliefs. *Philosophical Studies*, 176(9), 2497–2515. https://doi.org/10.1007/s11098-018-1137-0

Basu, R. (2019c). Radical moral encroachment: The moral stakes of racist beliefs. *Philosophical Issues*, 29(1), 9–23. https://doi.org/10.1111/phis.12137

Basu, R., & Schroeder, M. (2019). Doxastic wronging. In B. Kim & M. McGrath (Eds.), *Pragmatic encroachment in epistemology* (pp. 181–205). Routledge.

Begby, E. (2021). *Prejudice: A study in non-ideal epistemology*. Oxford University Press. https://doi.org/10.1093/oso/9780198852834.001.0001

Bolinger, R. J. (2020). Varieties of moral encroachment. *Philosophical Perspectives*, 34(1), 5–26. https://doi.org/10.1111/phpe.12124

Brewster, Z., & Mallison, C. (2009). Racial differences in restaurant tipping: A labour process perspective. *Service Industries Journal*, 29(8), 1053–1075. https://doi.org/10.1080/02642060902764343

Brinkerhoff, A. (2021). Prejudiced beliefs based on the evidence: Responding to a challenge for evidentialism. *Synthese*, 199(5–6), 14317–14331. https://doi.org/10.1007/s11229-021-03422-y

Brinkerhoff, A. (2025). Sexist beliefs in a sexist world: Exploring the causal role of sexism in sexist beliefs. *Episteme*, 22(4), 952–970. https://doi.org/10.1017/epi.2024.61

Callahan, L. F. (2018). Moral testimony: A re-conceived understanding explanation. *The Philosophical Quarterly*, 68(272), 437–459. https://doi.org/10.1093/pq/pqx057

Enoch, D. (2014). A defense of moral deference. *Journal of Philosophy*, 111(5), 229–258. https://doi.org/10.5840/jphil2014111520

Fletcher, G. (2016). Moral testimony: Once more with feeling. In R. Shafer-Landau (Eds.), *Oxford studies in metaethics* (Vol. 11, pp. 45–73). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780198784647.003.0003

Gardiner, G. (2018). Evidentialism and moral encroachment. In K. McCain (Ed.), *Believing in accordance with the evidence: New essays on Evidentialism* (pp. 169–198). Springer. https://doi.org/10.1007/978-3-319-95993-1_11

Gevrek, Z. E., Gevrek, D., & Neumeier, C. (2020). Explaining the gender gaps in mathematics achievement and attitudes: The role of societal gender equality. *Economics of Education Review*, 76, 1–31.

Groll, D., & Decker, J. (2014). Moral testimony: One of these things is just like the others. *Analytic Philosophy*, 55(1), 54–74. https://doi.org/10.1111/phib.12033

Hills, A. (2013). Moral testimony. *Philosophy Compass*, 8(6), 552–559. https://doi.org/10.1111/phc3.12040

Hills, A. (2016). Understanding why. *Nous*, 50(4), 661–688. https://doi.org/10.1111/nous.12092

Hopkins, R. (2007). What is wrong with moral testimony? *Philosophy and Phenomenological Research*, 74(3), 611–634. https://doi.org/10.1111/j.1933-1592.2007.00042.x

Howell, R. (2014). Google morals, virtue, and the asymmetry of deference. *Nous*, 48(3), 389–415. https://doi.org/10.1111/j.1468-0068.2012.00873.x

Leslie, S.-J. (2017). The original sin of cognition: Fear, prejudice, and generalization. *Journal of Philosophy*, 114(8), 393–421. https://doi.org/10.5840/jphil2017114828

Markovits, J. (2010). Acting for the right reasons. *Philosophical Review*, 119(2), 201–242. https://doi.org/10.1215/00318108-2009-037

Nickel, P. J. (2001). Moral testimony and its authority. *Ethical Theory and Moral Practice*, 4(3), 253–266. https://doi.org/10.1023/A:1011843723057

Niederle, M., & Vesterlund, L. (2010). Explaining the gender gap in math test scores: The role of competition. *Journal of Economic Perspectives*, 24(2), 129–144. https://doi.org/10.1257/jep.24.2.129

Nussbaum, M. C. (1995). *Poetic justice: The literary imagination and public life*. Beacon Press.

Sliwa, P. (2025). Changing minds and hearts: Moral testimony and hermeneutical advice. In R. Shafer-Landau (Eds.), *Oxford studies in metaethics* (Vol. 20, pp. 20–44). Oxford University Press. https://doi.org/10.1093/9780198958109.003.0002

Cite this article: Brinkerhoff, A. 2026. Moral Testimony, Moral Understanding, and "Rational" Prejudiced Beliefs. *Canadian Journal of Philosophy* 56: e10, 1–18, doi:10.1017/can.2026.10107