

Radical Internalism

ZOË JOHNSON KING
New York University

*'Cause the circle of hatred continues unless we react
We gotta take the power back!
– Rage Against The Machine*

1. Introduction

We live in dark times, morally, politically, and epistemically. Oppressive social structures create hierarchies of dominance and marginalization between social groups — hierarchies based on gender, race, class, sexual orientation, and disability, for example. And these social hierarchies each reinforce and are reinforced by ideologies that both limit the evidence we receive about them and distort the ways we learn to respond to the evidence that we have.¹ It is thus difficult to figure out the truth about morally and politically important matters, not only because these matters are complex, but also because ideology prevents us from acquiring and appropriately reacting to accurate information about them. In this way, ideology sustains oppressive social structures by preventing us from learning about them.

If you disagree with some of what I've said so far, then you can stop reading now. This paper is addressed to those who agree with this rough description of the times we live in and would like to figure out what we ought to do — morally and epistemically — in light of it.

Amia Srinivasan (2019) introduces the term “radical worldview” for the sort of worldview I just articulated. Srinivasan’s paper is a novel and powerful defense of externalism about epistemic justification. She argues that externalism should be preferred to internalism by those who hold a radical worldview (like the one I just articulated). This paper is a reply to Srinivasan. I share her radical worldview. But do not agree that externalism about justification is the preferable approach in light of the radical worldview that we share. On the contrary, I think that an internalist view of justification is attractive in these dark times, and in this paper I will articulate and defend a first-pass version of this “radical” internalism. So, while the paper should be of interest to those who want to assess Srinivasan’s recent argument, I hope that the positive view explored here is also of interest in its own right.

2. Cases, cases, and more cases

¹ In this paper I will not take a stance on the metaphysics of ideology. For work on this see, for example, the essays in *Analyzing Ideology* (forthcoming), ed. Celikates, Haslanger, and Stanley.

Here is how Srinivasan characterizes internalism and externalism (pp.6-7):

According to epistemic internalism, whether S's belief is justified supervenes solely on facts to which S has introspective access: 'internal' duplicates, the internalist says, do not differ in justification. A typical internalist might say that epistemic justification is a matter of fit with one's evidence, or with one's epistemic reasons, or more generally with how things look from one's own perspective on the world – where it is presumed that such facts are introspectively available to one. Epistemic externalism, meanwhile, denies that epistemic justification supervenes solely on such introspectively accessible facts: 'internal' duplicates might well differ in justification. The externalist says that epistemic justification is a matter of facts to which there is evidently no guarantee of introspective access – for example, whether one's belief exhibits an appropriate causal connection to its content, or is a product of a reliable or safe method.

Notice that, on this characterization, internalism is simply a supervenience claim and externalism is simply the denial of that claim. Internalism is the claim that justification supervenes on introspectively accessible features. Externalism is the claim that it doesn't. I will return to this point in the next section.

Srinivasan's argument for externalism is based on three cases. The cases are realistic and sophisticated. But, for present purposes, we just need a summary of their central features. Here is one:

Domestic Violence (p.5): Radha lives in a community suffused by patriarchal ideology, which leads her and everyone around her to believe that she deserves the brutal beatings that she receives from her husband whenever he judges that she has been insufficiently obedient or caring. She has never doubted this, nor has anybody given her any reason to.

Racist Dinner Table (pp.1-2): Nour, a British woman of Arab descent, goes to dinner with a friend's family. She leaves the table with the strong sense that her friend's father is racist against Arabs, though she can point to nothing specific that happened at dinner to support this sense. In fact, the friend's father is indeed racist, and Nour was subconsciously picking up on subtle behavioral cues to which she is reliably sensitive.

Classist College (pp.3-4): Charles is a junior fellow from a working-class background at an Oxbridge college publicly committed to promoting diversity. Charles is struck by instances of classist behavior from the other fellows, which he in fact detects reliably. But, when he raises this issue with the college Warden – who is also from a working-class background – the Warden tells Charles that this behavior is harmless fun and he is being overly sensitive. Nevertheless, Charles retains his belief that the college is classist, though he lacks the concept of false consciousness and has no idea that the Warden is suffering from it.

Srinivasan says that the internalist verdict on *Domestic Violence* must be that Radha is justified in believing that she deserves to be beaten, since “she is doing as well as she can, given how things seem from her own limited view on things” (p.7-8). Meanwhile, according to Srinivasan, the internalist verdict on *Racist Dinner Table* and *Classist College* is that Nour and Charles are unjustified in believing that they have been victims of racism and classism, since “neither has any introspective access to the reliable grounds of their beliefs” (p.7). But the externalist verdicts on these cases, Srinivasan says, are that Radha is *unjustified* while Nour and Charles are justified, since Nour and Charles' beliefs are “reliably and safely connected to the truth” and thus “straightforwardly satisfy the typical externalist conditions on justification”, whereas Radha “fails

to exhibit such a connection between her belief and the truth” (p.7). And, Srinivasan says, these externalist verdicts on the cases are far more intuitively compelling than the internalist ones. This provides intuitive support for externalism.

Srinivasan also argues that her cases are “radical” analogues of three classic cases widely thought to favor internalism. She argues that *Domestic Violence* is analogous to a traditional brain-in-a-vat case, in which a brain is kept in a vat attached to a supercomputer and fed sensory input that makes it seem as though it lives in the real world. She argues that *Racist Dinner Table* is analogous to a traditional clairvoyant case, in which someone’s clairvoyant power gives rise to a hunch that the President is in New York City, but he has no evidence for or against the existence of clairvoyant power and no evidence that he has this power.² And she argues that *Classist College* is analogous to a traditional case of dogmatism, in which a visitor to an art gallery sees a red object, is told, falsely, by a gallery assistant that the object is illuminated by trick lighting that makes it look red no matter what color it is, but ignores this testimony and continues to believe that the object is red just because it looks red.³ When faced with these traditional cases, most people have strong intuitions that the brain-in-a-vat is justified in believing that it has hands but the clairvoyant is unjustified in believing that the President is in New York City and the dogmatist unjustified in continuing to believe that the object is red (after talking to the gallery assistant). These intuitions are widely thought to favor internalism about justification. (I will say more about the classic cases in sections 3 and 4.)

Srinivasan intends for these analogies to twist the knife in her case against internalism. She argues that, since externalist verdicts are intuitively compelling in her “radical” cases and these are structurally analogous to the classic cases, internalists face an explanatory challenge. They must either explain away our externalist-friendly intuitions in her cases or find a difference between the radical cases and the classic ones that explains why internalist verdicts might be correct in the latter even if externalist verdicts are correct in the former. If internalists cannot meet this challenge, then, according to Srinivasan, this gives us reason to suspect that there was something amiss with our intuitions in the classic cases all along. For we should think it more likely that our intuitions are reliable in response to realistic cases like hers and unreliable in response to the “recherché” cases of philosophical imagination than the other way around (p.24). Thus, Srinivasan not only offers three cases in which the externalist verdict is compelling, but also offers grounds to doubt the reliability of our internalist-friendly intuitions in the classic cases.

Srinivasan then offers her own explanation of why our intuitions about her radical cases differ from those in the classic cases. Her explanation is that the radical cases are radical; they are “what we might call ‘bad ideology’ cases – that is, cases that feature subjects who exist in conditions in which pervasively false beliefs sustain and are sustained by systems of social oppression” (p.15). Srinivasan says that, in such cases, “the salient epistemological question becomes... how [subjects’] beliefs relate to systems whose function is to distort [their] access to the truth” (*ibid.*). And she argues that externalism is well-equipped to answer this question, being a “structural” epistemology focused on individuals’ relationships to their surrounding environments (p.18). Internalism, meanwhile, is ill-equipped to answer the question as it is a “meritocratic” epistemology focused on individuals’ epistemic “conscientiousness” and thus blind to structural concerns (*ibid.*). This is a formidable challenge for internalism; if it cannot handle bad ideology cases, that is a major problem.

However, as bad ideology cases go, Srinivasan’s are a little simplistic. Each one focuses on a single axis of oppression; patriarchy, racism, classism. And they are described in such a way that, in each case, there is a

² This case is originally from Bonjour (1980).

³ This version of the case is from Lasonen-Aarnio (2010).

clear epistemic hero who is a victim of oppression and either succumbs to the distorting influence of bad ideology or somehow manages to resist it and ascertain the moral truth, while everyone else is an epistemic villain who perpetuates oppression by reinforcing the effects of bad ideology on our hero. This is not what real epistemic life is like. There is never just one axis of oppression operating in a context. On the contrary, our social world is pervaded by intersecting axes of oppression, whose associated ideologies interact in complex, subtle, and hard-to-discern ways. Thus, there are few real cases in which it is clear who are the epistemic heroes and who the epistemic villains. And there are no real cases in which it is helpfully stipulated to us that our hunches and gut feelings are or are not reliable. Srinivasan's three cases thus drastically oversimplify the phenomena. To be sure, thinking about bad ideology at all is more realistic than thinking about clairvoyants and brains in vats. This is a significant step in the right direction. But, to be even more realistic, we should consider cases involving multiple intersecting ideological forces.

Here are two such cases:

Microaggression. Alice, a white woman, has a colleague, Liam, a black man. In repeated interactions at their department over a prolonged period of time, Alice gets the subtle sense that Liam is speaking to her in a dismissive and slightly aggressive way – more dismissive and aggressive than the way he speaks to his male colleagues. So it seems to her as if Liam is treating her in a somewhat sexist way. However, when she confronts him about this, Liam says “Oh Lord. Here we go again. People always accuse me of treating them in a dismissive and aggressive way. That’s because I’m a black man, and our perception of social reality is permeated with racial stereotypes, including the stereotype that black men are aggressive. People – in general, but particularly white women like you – are taught to fear black men and see us as dangerous. This is why I end up being accused of dismissive and aggressive behavior no matter what I do to try to give off the exact opposite impression!” Liam then reminds Alice of several occasions on which it is very clear that he has gone out of his way to be exceptionally gracious and polite.

Prejudice. An interdisciplinary university committee is discussing the merits of a proposed policy stating that departmental social events should not be held in bars. Two committee members get into a heated disagreement about the policy. One of them, a Muslim woman, says that she often feels deeply uncomfortable in the “boys’ club” atmosphere that she finds at a lot of bar-based social events in her department, and that such events are exclusionary for people like her who choose not to drink for religious reasons (or for any of a whole host of other potential reasons), but are then pressured to do so in order to access the social and professional opportunities available to junior academics in her field. Another committee member, a trans woman, says that her deep discomfort is often strongest at social events held within academic spaces, which are spaces that she has come to associate with hostility and exclusion and which often have their own “boys’ club” atmosphere, and that she would be harmed if the opportunity to relax in a more informal setting with colleagues and others in her field were removed. At the end of the discussion, each of the women has the subtle sense that the other is prejudiced against them and is interpreting their remarks in line with false and pernicious social stereotypes about them – stereotypes associating Muslim women with meekness and conformity and stereotypes associating trans women with malice and debauchery. They each contact the committee’s Chair to air their grievances, and each are surprised to learn from the Chair that a similar grievance has been aired against them.

These cases involve multiple axes of oppression; sexism and racism, Islamophobia and transphobia. And in these cases it is accordingly much harder to tell who is the epistemic hero and who the epistemic villain. Moreover, it is possible — indeed, plausible — that the protagonists in these cases are both heroes and villains, in that Liam’s responses to Alice are somewhat sexist *and* her reaction to him is somewhat racist, and the women on the committee are affected by both transphobia and Islamophobia. Thus, the people in these cases might each be getting something right and something wrong. However, there is also a way of filling out the details of *Microaggressions* such that Alice and Liam cannot both be right. For there is a version of the case in which Liam flat-out denies that he has treated Alice in a sexist way and stridently maintains that her subtle sense to the contrary is solely the product of racist ideology, contrary to what she claims. This is the version of the case that I want the reader to have in mind going forward.⁴

Srinivasan says that, in thinking about bad ideology, we should take seriously the Marxist insight that “the position of the oppressed can afford a dispensation from some of the epistemic ills of the oppressors” (p.17). By this she means that being a member of a marginalized social group can affect one’s epistemic position in complex ways. It can limit one’s access to a lot of information, sometimes including information about the nature of one’s oppression (as in Radha’s case). But it can also make one unusually good at recognizing the forms of oppression to which one is regularly subject and the associated forces of bad ideology with which one is all too familiar, even if this recognition occurs in a subtle and inchoate way (as in Nour and Charles’ cases). I agree with all of this. But the Marxist insight does not mean that any member of a marginalized group who gets the subtle sense that she is being treated in a prejudiced way should trust this gut feeling unwaveringly — particularly not in realistic cases involving multiple intersecting axes of oppression. The trouble is that, even if being marginalized in some respect makes you good at identifying ideological forces associated with the form of oppression to which you are subject, it does not necessarily make you any good at identifying other ideological forces associated with other forms of oppression (to which you are not subject). And it may even blind you to them. Some people see all social interactions through the lens of the particular marginalized identity(ies) they occupy, and thus ignore or underestimate available evidence about other forms of oppression operating in their context. This can leave them with a view of their situation that is badly one-sided and unduly self-serving.

So, suppose that Alice dismisses Liam’s testimony out of hand and continues to believe that he treats her in a sexist way, fighting dogmatism with dogmatism. And suppose that Alice does this just because she is convinced that the way things seem to her must be the way they are, as a result of which she dismisses all testimony to the contrary as mistaken (or malevolent, or both). Intuitively, her belief that Liam treats her in a sexist way is then not very well-justified. On the contrary, it is epistemically irresponsible to refuse to subject the way things seem to you to critical scrutiny, especially when you have information indicating that your seemings have been warped by the distorting influence of bad ideology. And it is hubristic to reject all testimony that suggests that you might be mistaken as itself mistaken (or malevolent, or both), especially if this testimony concerns the distorting influence of bad ideology on your beliefs. Intuitively, an attitude like Alice’s belies the wrong sort of approach to the whole project of forming beliefs for her hubristic beliefs to be justified. Moreover, this intuition remains compelling even if we stipulate that, as a matter of fact — but unbeknownst to everyone involved — Alice was in fact getting things right all along, and it was Liam who was wrong. Suppose that Liam does in fact treat Alice in a sexist way, and that her perception of him as dismissive and aggressive is not the effect of racial stereotypes but of a reliable sensitivity to sexism on her part. This does not help Alice’s case. Intuitively, even if someone’s gut feeling results from a reliable sensitivity to the facts in a certain domain, her belief is not justified if the thought-process that takes her from gut feeling to belief is just epistemic hubris wrapped in a cloak of righteousness.

⁴ Thanks to Gabrielle Johnson and Sukaina Hirji for pressing me to clarify this construal of the case.

Similar remarks apply in *Prejudice*. If each of the women on the committee retains her narrative about what happened in the meeting and her view about the merits of the university policy, dismissing her colleague's arguments to the contrary as just thinly veiled Islamophobia/transphobia, then none of these beliefs are epistemically justified. In general, if someone learns information indicating that their responses to evidence have been distorted by pernicious stereotypes about social groups to which their colleagues belong, then it is intuitively epistemically irresponsible for them not to pause and reconsider those responses. And this verdict remains intuitively compelling even if we stipulate that they are in fact reliably detecting a form of oppression to which they are regularly subject. That does not make it okay to simply ignore evidence that their beliefs might still be warped by insidious prejudices of their own.

This is bad for Srinivasan. For *Microaggressions* and *Prejudice* are structurally analogous to *Classist College* and to the classic case of the dogmatist: the four cases are all examples of misleading higher-order defeat. Such cases have three main features. First, someone is told by a credible source that distorting factors are disrupting their ability to reliably detect the facts in some domain. Second, as a matter of fact (though unbeknownst to everyone involved), these distorting factors have not affected them and they are reliably detecting the facts. Third, they choose to dismiss the credible source and continue to believe in accordance with the way things seem to them. The question in all cases is whether the protagonist's belief is justified at the end of this little story. In cases involving no ideology, like the traditional dogmatist case, the intuitive answer is that it is not justified. In cases involving a single axis of oppression, like *Classist College*, Srinivasan argues that the intuitive answer is that it is justified. But in cases involving multiple axes of oppression, like *Microaggressions* and *Prejudice*, the intuitive answer is again that it is not justified. So it is simply untrue that bad ideology cases support externalism while only the *recherché* cases of philosophical imagination support internalism. Counterexample: *Microaggressions* and *Prejudice* are bad ideology cases, but they support internalism.

Moreover, multiple-axis cases like mine are more realistic than single-axis cases like Srinivasan's, since in real life there is never only a single axis of oppression operating in a context. So, if we should assume that our intuitions are more reliable in more realistic cases — as Srinivasan argues we should — then we should assume that our intuitions are more reliable in *Microaggressions* and *Prejudice* than in *Classist College*.

If the reader is not convinced by my cases, then she is welcome to devise her own. Here's how. Consider someone who is a member of a marginalized group, which makes her unusually reliable when it comes to detecting facts pertaining to her own oppression. Imagine that she disagrees about something with a member of another marginalized group — perhaps one that she has not thought about a great deal and perhaps one whose marginalization she does not take very seriously. Stipulate that, in fact, this other person does treat our protagonist in a subtly prejudiced way, which she subtly detects. But stipulate further that she acquires a lot of higher-order evidence indicating that her subtle sense of ill-treatment results from bad ideology surrounding the marginalized group to which her interlocutor belongs and is thus not a reliable indicator of ill-treatment in this case. Now imagine that she digs in her heels as the evidence mounts. Imagine that she refuses to question the way things seem to her, despite learning more and more information indicating that her seemings are thoroughly permeated with distorting ideological influence. If you can imagine any cases like this in which the protagonist's dogmatic beliefs start to seem poorly-justified, then you can imagine a case that illustrates the point I am making here.

What we need is a way of thinking about epistemic justification that accommodates the fact that almost all of us are socially privileged along some dimensions and marginalized along others, which affects our epistemic position in complex ways. In section 4 I will argue that a suitably "radical" internalism is well-

placed to accommodate this fact, and will propose a first-pass version of such a view. Before that, though, I briefly take up Srinivasan’s challenge to explain away our externalist-friendly intuitions in her cases, since I have already said most of what it takes to meet this challenge.

3. Taking up the gauntlet

Let’s get one thing straight. Srinivasan says that “the” internalist verdict is that Radha is justified but Nour and Charles are unjustified, while “the” externalist verdict is that Nour and Charles are justified but Radha is unjustified. But these definite descriptions are a stretch, given how minimally Srinivasan characterizes internalism and externalism. As I noted above, Srinivasan casts internalism as a simple supervenience claim and externalism as the simple denial of that claim: internalism says that internal duplicates cannot differ in justification, and externalism denies this. So, internalists are not forced to accept the verdicts that Srinivasan finds counterintuitive. All the internalist must say is that any internal duplicate of Radha, Nour, or Charles has the same degree of justification for the belief that they deserve to be beaten, their host is racist, or their college is classist as these agents do. This may be any degree of justification, as far as the supervenience claim goes.⁵ Likewise, it is not clear that externalism can reap the intuitive benefits that Srinivasan takes her cases to provide. For the simple denial of a supervenience claim obviously does not entail the specific verdicts about cases that she finds compelling. Indeed, some externalists would agree with my verdicts on *Microaggressions* and *Prejudice* — and, by analogy, *Classist College* — on the grounds that gut-feeling-plus-hubris is an unreliable belief-forming process.⁶

Nonetheless, Srinivasan is right to suggest that there are versions of *Racist Dinner Table* and *Classist College* that are radical analogues of the clairvoyant and dogmatist cases, thus directly challenging the views about justification that these cases are supposed to support. The classic cases are supposed to elicit the intuition that someone who in fact exercises a reliable sensitivity to certain facts is nonetheless not justified in believing the hunches that result from this reliable sensitivity, if she has no idea that her hunches result from it, no idea that she has it, and indeed no idea that such a thing is even possible (clairvoyant), or if she has acquired higher-order evidence indicating that it is unreliable in the case at hand (dogmatist). If this is true then there are versions of *Racist Dinner Table* and *Classist College* in which Nour and Charles are unjustified. In *Racist Dinner Table*, we would have to stipulate that Nour has no idea how her sensitivity to racism operates, no idea that she has such a sensitivity, and indeed no idea that such a thing is even possible. Her sense that her friend’s father is racist against Arabs seems to just pop into her head for no

⁵ Srinivasan herself notes that some forms of internalism do not entail that Nour is unjustified; in a footnote she says that forms of “permissive internalism”, according to which one is justified in believing P if one has an undefeated hunch that P (p.9, fn.19), allow that Nour is justified. Srinivasan argues that permissive internalism is implausible. I agree. I am not a permissive internalist. But the mere possibility of this view is enough to show that internalists are not forced, by the very nature of internalism, to accept the verdicts that Srinivasan foists upon them.

⁶ Srinivasan also concedes this point, observing that “most externalists” modify their view in response to cases like that of the dogmatist. She distinguishes her “pure” externalism from a “moderate” externalism that allows that misleading higher-order evidence can defeat justification (pp.11-12), citing Alston (1988), Bergmann (2006), Goldman (1986), and Nozick (1981) as moderate externalists. Srinivasan takes her intuition about *Classist College* to support pure externalism. Nonetheless, the fact that most externalists are moderate externalists shows that externalism does not all by itself entail the verdicts that she finds plausible in her cases. This also means that my cases do not challenge all forms of externalism. Rather, they challenge the particular version of externalism that Srinivasan takes her cases to support. And they support the verdict that Srinivasan casts as “the” internalist verdict on *Classist College*: that someone who in fact exercises a reliable sensitivity to the facts is nonetheless not justified in accepting the resultant hunches if she has compelling higher-order evidence indicating that her hunches are unreliable. Thanks to Alex Worsnip, Daniel Fogal, and Stephanie Leary for helpful discussion of this point.

reason. In *Classist College*, by contrast, we would have to stipulate that Charles is aware that he is good at detecting classism under normal conditions, but that the warden points to credible reasons for Charles to think that his sensitivity to classism is unreliable in the particular cases at hand, and yet Charles ignores his testimony. These would be radical cases that are structural analogues of the traditional cases.

But, in these versions of the cases, Nour and Charles don't seem very well-justified. Moreover, it is difficult even to understand what is happening in these versions of the cases. If Nour really has no idea that it is possible for people to be reliably sensitive to racism, and if the hunch that her friend's father is racist seems to just pop into her head for no reason, then why does she start believing it? Believing the content of what seems to her to be a random intrusive thought for no reason at all, or just for fun, seems clearly unjustified. And if Nour believes it because she baselessly assumes that the way things seem to her must be the way they are, then she exhibits precisely the sort of hubris that seems objectionable in *Microaggressions*, *Prejudice*, and the traditional clairvoyant and dogmatist cases. Likewise, if Charles has no positive grounds on which to dismiss the Warden's credible testimony about his unreliability in the case at hand, then what is he doing in dismissing the Warden's testimony? If his dismissal is groundless, or just for fun, then it again seems patently unjustified. And if it is based on surety that he must be getting things right and thus that the Warden must be getting things wrong, then it again seems objectionably hubristic.

Other, more natural ways of filling out the details of *Racist Dinner Table* and *Classist College* make Nour and Charles' beliefs seem better. People in marginal social positions often have some grasp, however nascent, of the fact that we are better than others at detecting subtle biases against members of our group, having experienced such things a lot. We can trust our instincts on this basis. And people in marginal social positions often become familiar with the various ways in which others in the same social position learn to capitulate, or not to capitulate (as the case may be), to those in dominant positions in order to safely navigate our social world. We can trust or distrust others' testimony on this basis. Moreover, we can do all of this even if we lack the theoretical concepts to describe what we are doing. So, although Srinivasan stipulates that Charles lacks the concept of false consciousness, he might still understand that the Warden's instincts have been corrupted by overexposure to their classist environment and ignore the Warden on this basis. If Nour and Charles' beliefs are like this, then they seem much better-justified. But internalism easily accommodates this. If Nour and Charles are not just reliable but aware of their reliability (however they conceptualize it), then the cases are no counterexamples to internalism.

Thus, anybody who wants to use Srinivasan's cases to challenge internalism faces a dilemma. Either Nour and Charles have some grasp of their reliability, or they don't. If they don't, then their beliefs do not seem very well-justified after all. And, if they do, then internalism has these cases covered.

What about *Domestic Violence*? I return to some important versions of this case below. But the short answer is that this is just a bad analogy. Srinivasan suggests that, while the brain-in-a-vat is literally in a vat, Radha is "metaphorically envatted in misogynistic ideology" (p.14). But this vague imagery masks a crucial disanalogy. In the brain-in-a-vat case, what drives the intuition that the brain is justified in believing that it has hands is the fact that its experiences are indiscriminable from those of a veridical perceiver. There seems something off about saying that the perceiver is justified in believing that it has hands but the brain is unjustified, given that neither believer can even tell whether she is in the good case or the bad. This motivates the thought that justification supervenes on introspectively accessible states. But no analogous argument can be made in *Domestic Violence*, because there is no veridical believer whose experiences are indiscriminable from Radha's. For moral facts supervene on descriptive facts. So, given that Radha does not deserve to be beaten, no-one indiscriminable from her does deserve to be beaten. There are no distant possible worlds that are just like Radha's except that women deserve to be beaten if they are insufficiently

obedient or caring. That is just not how moral facts work. So, the very feature that drives the internalist-friendly intuition in the brain-in-a-vat case is absent in *Domestic Violence*. It's a bad analogy.

But couldn't there be a the case in which Radha is trying her best to figure out the truth about her situation, and meets all of the internalist criteria for justification, yet is still duped by ideological social forces into believing that she deserves to be beaten? Of course there could. I discuss some such cases in the next section. My take will be that — *pace* Srinivasan — Radha's belief in these cases may be somewhat justified. To see why, we must now take some initial steps toward articulating a genuinely "radical" internalism.

4. Radical Internalism

Recall what seemed wrong with the protagonists in *Microaggressions* and *Prejudice*: they seemed to display an objectionable kind of epistemic hubris. This kind of hubris is particularly egregious — and particularly dangerous — in bad ideology cases. But it is the kind of thing to which epistemologists have been objecting all along. For instance, Laurence Bonjour, the author of the clairvoyant case, said back in 1980 that the clairvoyant's belief is "irresponsible, and therefore unjustified" (p.63). In spelling out this intuition, Bonjour suggested that we have epistemic duties to "reflect critically upon [our] beliefs, and such critical reflection precludes believing things to which one has, to one's knowledge, no reliable means of epistemic access" (p.63). Likewise, Stewart Cohen says of the dogmatist that he "has been epistemically irresponsible" and "ought not to have proceeded in the way he did", and that, as a result, his belief is unjustified (1983, p.284). And Hilary Kornblith says of a similar character that "had [his] actions been guided by a desire to have true beliefs, he would have listened carefully to the objection. Since his [continued belief] is due, in part, to this epistemically irresponsible act, his continued belief is unjustified" (1983, p.36). In short: my aversion to epistemic hubris is not new. On the contrary, there is a distinguished tradition of aversion to this attitude among theorists of epistemic justification, and especially among internalists.

To make progress here, I suggest that we move beyond characterizing internalism and externalism as a supervenience claim and its denial. Nobody in this literature is kept up at night by the thought that a certain pattern of facts might hold or fail to hold across worlds. Rather, we have deep theoretical intuitions that drive some of us toward externalism and others toward internalism. Srinivasan has given us a compelling articulation of the deep intuitions that drive her toward internalism. I will now attempt to follow suit.

Here's an attempt to articulate the deep intuition underlying the internalist aversion to epistemic hubris (at least in my own case): I want to reward people who are *trying*. Some people invest a lot of effort into trying to do something good — in ethics, to act well and be well-motivated, and in epistemology, to shape their beliefs into as full and accurate a picture of the world as possible. Of these people, some try but fail because the external world does not cooperate; for example, because their belief-forming method is unreliable in a way that they could not possibly detect. My inclination is to praise these people for their efforts and to emphasize that, though they failed, they did their best. Meanwhile, some of the people who put a lot of effort into their attempts to do good succeed in doing so. And sometimes their success is brought about by their effort, in that "right sort of way" that anyone involved in a literature dealing with deviant causal chains recognizes even if we cannot spell it out. This, in my book, is what an achievement is. I want to reward people who try to do something good and succeed (in the right sort of way) for their achievements.⁷

This deep intuition about who to reward goes hand-in-hand with a deep intuition about who *not* to reward. Lots of people do well, morally or epistemically as well as prudentially, as a result of generous endowments

⁷ For more on the notion of an achievement see Bradford (2015), and for more on rewarding trying see Mason (2019).

rather than as a result of careful effort on their part. My deep sense is that it is inappropriate to reward such people. They don't deserve our veneration because they are only doing well because their surrounding circumstances set them up to succeed without trying. They're just lucky. To illustrate: I have encountered a lot of people who put little effort into their studies because they were convinced that their charisma and wit would carry them to success, but whose families could afford such nice schools and had such powerful connections that this plan actually worked and they ended up in high-powered, high-paying positions. I have also heard tell, in externalists' thought-experiments, of people who put little effort into forming and revising their beliefs – that is, into questioning the reliability of their methods, seeking out a range of sources of evidence, carefully considering others' testimony, and so on – because they were convinced that their intuitive sense of what the world is like must be correct, but whose undetectable skills or favorable external circumstances ensured that this plan actually worked and they ended up with a lot of true beliefs. I am deeply unimpressed by both of these kinds of person. The same anti-elitist sense of who *not* to reward rebels inside me when I consider the thought that these hubristic successes might be praiseworthy.

What does this have to do with epistemic justification? The internalist tradition with which I identify holds that justification is a kind of *entitlement* – the entitlement to believe something – and that, as with other forms of entitlement, having a strong sense of it is not sufficient for having it. Nor is a generous endowment. Rather, we can earn the entitlement to believe something by putting in epistemic work. On this view, being epistemically responsible is not just a matter of sitting around with one's eyes open, waiting for evidence to come in and then making the most minimal revisions to one's beliefs that render them consistent with the new evidence. Instead, it includes proactively considering ways in which one's evidence is *incomplete*; that is, in which one is only learning about part of the world rather than gaining a fuller picture. And it includes proactively considering ways in which one's evidence, or one's way of responding to evidence, is *distorted*; that is, in which it misrepresents the phenomena under investigation, rather than presenting an accurate picture of them. In short, justification is earned by doing one's epistemic due diligence.⁸

I think that the fact that internalism stands opposed to epistemic hubris means that it could be particularly useful in thinking through bad ideology cases. For the problem in a lot of bad ideology cases – including *Microaggression* and *Prejudice* – is that members of dominant groups fail to reflect on ways in which their evidence is incomplete or distorted. Moreover, as we have seen, ideology can render people hubristic in this way even while being a member of some marginalized group(s) makes them unusually good at seeing through parts of its charade. The problem with unreflective members of dominant groups is often that they fail to put in enough of the kind of work that the internalist says is necessary for being well-justified.

Thus, “radical” internalists can put Srinivasan's Marxist insight into practice. We can do this by developing principles enjoining people to reflect critically on their beliefs in a way that takes seriously the fact that we are almost all socially privileged in some respects and socially marginalized in others, and that this affects our epistemic positions in complex ways.

Here is a first-pass attempt at some such principles:

1. Seek out evidence as to the dimensions along which you occupy a dominant social position and those along which you occupy a marginalized social position — and, accordingly, as

⁸ I owe the phrase “epistemic due diligence” to Elise Woodard (who uses it in a somewhat different way), to whom I am grateful for many productive conversations about radical epistemology over the years. I should also note that I do think that one's overall degree of justification is affected by the amount and quality of one's evidence, as well as by the degree to which one has been diligent in gathering this evidence. Thanks to Daniel Fogal for pressing me to clarify this.

to the topics on which your assumptions are likely to be more accurate and your inferences better than others', though ideology will incline both you and the others toward believing that your assumptions and inferences are worse than theirs (*mutatis mutandis* for dominant positions). As you acquire evidence about your social position along various dimensions, begin to apply principles (2) and (3), while inquiring further into how this position benefits and/or limits you epistemically.

2. If your best evidence suggests that you occupy a dominant social position along some dimension, then *be humble*: pay attention to what those in marginal positions are saying on topics relevant to the dimension along which you are dominantly situated; lend their testimony more weight than you would other testimony on the topics; question the truth of your assumptions and the validity of your inferences on these topics. Encourage other dominantly-situated people in your conversational circles to do the same. And if your evidence suggests that you are dominantly-situated along some dimension relevant to a conversation, then bring that up in the conversation, and note the epistemic limitations that it likely entails.
3. If your best evidence suggests that you occupy a marginalized social position along some dimension, then *be bold*: respond to skepticism from those in dominant positions on topics relevant to the dimension along which you are marginally situated by reminding yourself of the ways in which you are more likely to see these topics accurately and to draw good inferences from available information than they are. Point this out to them if it is safe to do so, and offer similar defenses of other marginally-situated people in your conversational circles if it is safe to do so.

These principles are a first pass. I'm sure that there are problems with them, and they will need refining. But I am confident that the basic idea behind them is correct. This idea is that, if we are genuinely concerned with shaping our beliefs into as full and accurate a picture of the world as possible, we should proactively seek evidence and pursue lines of inquiry that go beyond those suggested to us by our immediate epistemic community. We should find the stones left unturned by the forms of evidence and ways of reasoning easily accessible to us, and then we should turn those stones. If we have reason to think that our social position improves our access to information or our ability to reason about certain matters, then we should engage in further inquiry as to how that works and which matters it covers, and should trust our instincts on such matters in the face of disagreement with those less well-positioned, trying to share what we believe with them in an accessible way insofar as it is safe to do so. And if we have reason to think that we are the ones who are less well-positioned in certain respects, then we should trust the instincts of those who we have reason to think are better-positioned, and should work to try to understand what they are trying to tell us.

One might worry that someone needs the concept of a social position, and quite an extensive education in standpoint epistemology, in order to follow these principles. But that is not how I intend for the principles to be read. We need not have the concept of a marginalized social position in order to be familiar with the ways in which social power can render someone's evidence incomplete or distorted and in which being on the outside can make someone better at identifying and reasoning about what is really going on. Quite the reverse: everyone is familiar with these phenomena, since they happen all the time. As we grow up, we all see versions of them happening at home and at school. We are thus all in a position to recognize the central tenets of standpoint epistemology *de re* if not *de dicto*, and thereby in a position to ask how these tenets can be applied to our own case. (Indeed, for present purposes we can see these tenets as plausible applications of general epistemological principles about higher-order evidence — on which see Toole *fc.*) And that is enough for principles like mine to find purchase in everyone's lives.

Notice that these principles do not use the words “justified” or “justification”. They are simply instructions for how to manage one’s epistemic states. Nonetheless, my preferred brand of internalism sees justification as an entitlement to believe something that we can earn by putting in epistemic work, which is a matter of following principles like these. On this picture, justification is a complex matter. It comes in degrees. And someone’s overall degree of justification for a belief depends on several factors; not only the amount and quality of her evidence for the truth of the relevant proposition, but also the amount and quality of the work she has done to consider ways in which this evidence is incomplete or distorted and act accordingly – to turn the stones. One notable upshot is that it is not clear that it makes sense to speak of someone as being justified or unjustified *simpliciter*. Nor is it clear that it makes sense to speak of a belief’s being “fully justified”. We can always be more justified. For we can always seek more evidence, or reflect further on the quality of our current evidence and on its possible limitations, taking steps to correct these limitations if any seem necessary. This is particularly clear in bad ideology cases: we can talk to more people whose social position is likely to give them insight into the situation, investigate our biases and parochialities and the ways they might be leading us astray, and so on. We can always put in more work.⁹

This picture gets the right result about all the cases we have discussed so far. *Microaggressions* and *Prejudice* involve people who do a great job of following principle 3, but a lousy job of following principles 1 and 2; they overlook the ways in which their dominance limits their epistemic position, being too focused on the ways in which their marginalization improves it. As a result, their beliefs are not very well-justified. They are not totally unjustified; each has some justification, given her evidence and given that her social position can be expected to make her good at detecting facts pertaining to her own marginalization. But they are not very well-justified, as this is not the whole story, and they obtusely ignore the rest of the story. Likewise, in the versions of *Classist College* and *Racist Dinner Table* in which Nour and Charles’s beliefs are based on nothing but unfounded hubris, they are not justified because they are making no attempt whatsoever to put in epistemic work. But in the versions of these cases in which Nour and Charles have some grasp on the ways in which their social position affects their epistemic position, they are somewhat justified, as they are then following principle 3. These are the right results.

We might worry about cases in which someone *is* putting in epistemic work, but is subject to such badly distorting ideological influence that she still gets things wrong, no matter how hard she tries and how much she checks and double-checks. Some versions of *Domestic Violence* are like this. Recall: this case is not about someone with a reliable sensitivity to facts about a form of oppression to which she is subject, which leads to gut feelings that she wonders whether to trust. Rather, this is a case in which a victim of oppression has no recalcitrant gut feelings and no understanding of her own marginalization, having been convinced that everything is hunky-dory by sophisticated but ultimately false and dangerous arguments. In this case, bad ideology works so insidiously that it shields itself from view. So, we can imagine Radha actively trying to understand her social world, and repeatedly engaging in sustained and careful reflection on the matter, at the end of which she always ends up reaffirming that she sometimes deserves to be beaten. In this version of the case, my view entails that her belief that she sometimes deserves to be beaten is fairly well-justified – one of the verdicts at which Srinivasan balks.

But I don’t think that this is such a terrible verdict. For an internalist, saying that someone’s belief is justified in no way constitutes an endorsement of their belief-forming circumstances. So, saying that Radha’s belief

⁹ Some epistemologists think that there can be no *epistemic* reasons to do things like talking to people and investigating our biases, and that our reasons for doing such things must be practical. I don’t think that. I think there can be epistemic reasons for action. For a defense of this view see Singer and Aronowitz (forthcoming). For a more general exploration of the relationship between norms of inquiry and other epistemic norms see Friedman (forthcoming).

is justified is not an approval of patriarchal ideology — no more than saying that the brain-in-a-vat’s belief is justified is an approval of the practice of putting brains in vats. Rather, for the internalist, saying that someone’s belief is justified is a way of giving her credit for doing the best she can, epistemically speaking, under whatever circumstances she is in. And that is true of Radha in this version of the case. By contrast, to say that her belief is unjustified because it is suffused with patriarchal ideology is to act as if Radha had no agency in her belief-formation and belief-revision process. But that is not true, and it is disrespectful to Radha. She is not a helpless dupe. Rather, she has amassed a set of mutually supporting metaphysical and moral arguments that collectively support the conclusion that she sometimes deserves to be beaten. And, when she has subjected these arguments to critical scrutiny, she has come back to them every time. Therein lies Radha’s moderate degree of justification. By acknowledging it, we acknowledge the epistemic agency that she still has within her patriarchal environment. This is preferable to talking about oppressed people as if they have been brainwashed, which literally adds insult to injury.¹⁰

Moreover, by allowing that Radha’s belief is fairly well-justified we allow ourselves to distinguish between the many different ways in which ideological distortion occurs. Sometimes bad ideology limits or obscures evidence. Sometimes it leads people to overlook evidence or weigh their total evidence poorly. Sometimes it teaches people spurious patterns of reasoning. And so on. It is important to distinguish these different forms of ideological distortion, because some lead people to be epistemically irresponsible while others lead responsible agents astray. This means that it will take different strategies for us to resist them. When we recognize that bad ideology can lead people to form dangerously false beliefs *and be fairly well-justified in so doing*, this reminds us that ideology operates subtly, such that even those who are on the lookout for it can be taken in by its charms. But if we say that any belief subject to ideological influence is thereby unjustified, we rob ourselves of the conceptual tools with which to talk carefully not only about how bad ideology works but also about how most effectively to combat it in different cases.

More on that in a minute. First, let me discuss two last kinds of case. Take a version of *Domestic Violence* in which Radha is not following principles 1–3, but this is only because of the effects of patriarchal oppression on her self-esteem. Or take a version of *Racist Dinner Table* in which Nour is unaware that she is reliably sensitive to racism against Arabs only because racist ideology has led her to underestimate her reasoning abilities and assume that she is bad at thinking about complex epistemic and political matters. My view suggests that, since these agents’ failures to recognize the epistemic benefits associated with their marginal social positions result from a failure to put in epistemic work, their beliefs are not very well-justified. And that seems unduly harsh, especially given that it is because they are victims of oppression that they are disinclined to do their epistemic due diligence.

This is an uncomfortable implication of my view. But I don’t see a good way out of it. I think we are saddled with it as long as we accept something like principle 1. And I am firmly committed to principle 1. We need this principle in order to say what is wrong with dominantly-situated people who fail to gather or reflect on evidence of their own limited epistemic positions. I hold that this failure limits their ultimate degree of justification. And I see no way to say this about people who fail to put in epistemic work and thus end up unduly bold without saying something parallel about people who fail to put in epistemic work and thus end up unduly humble. I leave it to other scholars to develop views that can accommodate the lenience we

¹⁰ One might think that Radha is a victim of *hermeneutical injustice* (à la Fricker 2007); she lacks the concept of domestic violence and as such is unable to understand her own experiences. But, if Radha lacks the concept of domestic violence, presumably she could not be doxastically justified in believing that she is a victim of domestic violence. So Srinivasan’s suggestion that Radha is also unjustified in believing that she deserves to be beaten is difficult to swallow, as it suggests that there is no proposition that Radha would be justified in believing. Thanks to Grace Helton for this point.

might intuitively want to show the latter agents. (Notice that Srinivasan’s externalism is not such a view; these agents’ belief-forming processes are thoroughly contaminated with the pernicious influence of bad ideology, so Srinivasan’s view also implies that they are unjustified.)

Now consider someone who is in fact in a dominant social position, but is surrounded by catastrophically misleading evidence indicating that this social position is marginalized. To be clear: I think that few, if any, real cases are like this. Most real people who occupy dominant social positions have ample evidence of this dominance available to them, such that a good-faith effort to follow principle 1 would uncover it, and their failure to recognize their dominance is largely the product of willful ignorance and wishful thinking.¹¹ But it remains possible for this not to be the case. It is possible that the best evidence available to someone indicates that they are marginalized when in fact they are not, even though they are putting in every effort to manage their beliefs well. (For instance, it could be part of their ideology that they must constrain their inquiry to trusted sources, with the result that they avoid acquiring evidence that would be dissonant with the ideology in a sincere attempt to not be led astray.¹²) On the view I have defended, these people might be warranted in being bold, and their beliefs might end up fairly well-justified.

This is another uncomfortable result. But I still think that it is the right result. To repeat: it is important to distinguish ways in which ideology leads people to be epistemically irresponsible from ways in which it leads responsible agents astray. To repeat again: I doubt that anyone really is so mired in misleading evidence that a well-intentioned effort to manage her epistemic states only sends her spiraling ever-deeper into an echo chamber. But if anyone is in this position, then it is important to distinguish her from ill-intentioned internet trolls who don’t care about truth — and from the more common case of a dominantly-situated person in a homogenous and unreflective epistemic community who just doesn’t try hard enough to follow principles 1 and 3. In the right kind of case, in which the fault lies solely with the evidence rather than the agent, I am comfortable saying that false and pernicious beliefs can be fairly well-justified. To repeat once more: these distinctions help us to see the different strategies that we will need to take in order to resist the different ways bad ideology operates. They also help us to avoid alienating our more conscientious interlocutors by lumping them in with the trolls. So, as well as getting the normative facts right, these distinctions are politically useful.

5. Coda: “Individualistic” epistemologies, “normative” epistemologies

Srinivasan argues that those of us with radical worldviews should prefer externalism to internalism on the grounds that internalism is an “individualistic” epistemology and externalism a “structural” epistemology. But I am unmoved by the allegation that internalism is “individualistic”. The fact that a view of justification includes principles addressed to individuals does not mean that it is blind to structural concerns. It may instead mean that the view’s proponents understand that these structural concerns can only be remedied by people doing things. That’s what I think. And I think that this remains the case if there are important facts about how groups and institutions should change in order to remedy structural problems. For groups and institutions are comprised of individuals. When they act, individuals do. And it is the individual acts, appropriately related, that constitute the group or institutional act. So, to implement principles about how groups or institutions should change, we ultimately need individuals to do things. And, when individuals try to change groups and institutions, what they do should be informed by their understanding of the

¹¹ Indeed, in many cases the agent already possesses sufficient evidence to recognize the ways in which their dominance limits their epistemic position, and they simply fail to draw well-supported conclusions from their total evidence. What such agents lack is not information, but *discernment*. Thanks to Ram Neta for helpful discussion of this point.

¹² Thanks to Mike Rea, Christina Van Dyke, and Kate Ritchie for pushing me to discuss this kind of case.

structure in which they are embedded and of their position within it. The principles that I offered in the previous section are intended as some first-pass suggestions for how this might go in the epistemic realm.

Those of us with radical worldviews should feel the need for principles like these, because we should want to talk not only about bad ideology's distorting influence but also about what to do about it. It's no good sitting around quoting Trotsky and lamenting the pervasive impact of oppressive social forces on our evidence and reasoning. By itself, that won't change anything. And it's no good simply stating that a belief is unjustified if it is distorted by bad ideology and justified if it is not. Again, by itself, that won't change anything. Nor is it much good to discuss vignettes in which we stipulate information about the reliability of subjects' belief-forming mechanisms that ordinary people don't have about themselves or each other, calling the subjects justified or unjustified according to this god's-eye-view stipulated information. These are unhelpful responses to the realities of epistemic oppression. As with other forms of oppression, the only acceptable response to epistemic oppression is to fight it as best we can. So, those of us with radical worldviews should be interested in what we — actual people, with the information that we have — can do to resist bad ideology. We should be figuring out what it amounts to to try to shape our beliefs, and those of others in our epistemic communities, into as full and accurate a picture of the world as possible, given that this world is permeated by ideological influence. For that project, we will need to develop some of the more action-guiding principles that are the internalist's stock-in-trade.

At the very end of her paper, Srinivasan responds to the challenge that externalism is not really a normative view since it fails to offer action-guiding principles. Here is what she says in response:

[This challenge] assumes that what it is to be a normative theory is to talk in terms that are familiar to us from ethics: blameworthiness and blamelessness, responsibility, action-guidance. But we might think this is an overly restrictive notion of a normative theory. Marxism, for example, is arguably a normative theory, in the sense that it is responsive to a gap between how things are and how things should be. But Marx was uninterested in the questions of what any given individual ought to do or who is to be blamed — concerns that he dismissed as typically bourgeois. Aristotle meanwhile thought the question of whether something is a *good* version of its kind — whether a citizen is a good citizen, or whether a thermometer is a good thermometer — was a paradigmatically normative question. And yet Aristotle thought that being a good version of one's kind is not something that lies solely within the will of that thing. Virtue requires the cooperation of the external world... So, too, the externalist thinks, with epistemic goods: the epistemic goods really worth having are not ours to have alone.

This response is inadequate. Clearly, externalism is “responsive to a gap between how things are and how they should be” in that it proposes criteria for the evaluation of people's beliefs and the processes that lead to them. Nobody questions this. The problem is that when a criterion's application turns on information that is inaccessible to the subjects of evaluation, it cannot be used by those subjects to determine what to do, even if they want to be positively evaluable according to the criterion.¹³ This is what theorists are worried about when they say that externalism is not normative. Of course, we can use the word “normative” to mean “evaluative”. But that will not make externalism normative, any more than using

¹³ One often hears the complaint that no principle is perfectly action-guiding, since we cannot always tell whether its conditions of application obtain. This is true. But it just means that action-guidingness comes in degrees, not that action-guidingness is not a desideratum of epistemic principles that are for use by actual people with our actual information. Thanks to Nick Hughes for helpful discussion of this issue.

“tail” to mean what “leg” means will make horses have five tails. A verbal maneuver is obviously not enough to tell us how to start working to try to remedy the underlying structural injustices that lead to discrepancies in the ease with which different people can believe well on different topics.

The ultimate question at issue between radical internalists and radical externalists is a version of the “What is a theory of justification *for?*” question traditionally at issue between internalists and externalists.¹⁴ Radical epistemologists answer this question with our radical worldviews in mind. We want, as Srinivasan puts it, “the right theory of justification for our political epistemology” (p.25, fn.46). I am wholeheartedly on board with this project. But I think that, *pace* Srinivasan, externalism about justification is just not very useful for radical politics. Of course we should pay attention to the ways in which bad ideology distorts our beliefs. But we can’t stop there — that’s just the first step. If that’s as far as externalism gets us, then, it doesn’t get us very far. Rather than taking information about ideological distortion to settle the facts about epistemic justification, we should see it as a call to arms. We should be thinking about how to combat this distortion and thus earn the right to believe things in spite of it. We can do that by adopting a genuinely radical internalism.

¹⁴ Thanks to Fatima Amijee for this way of putting the question.

References

- Alston, William (1988). "An Internalist Externalism". *Synthese* 74: 265-283.
- Armstrong, D.M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Bergmann, Michael (2006). *Justification Without Awareness*. Oxford: Clarendon Press.
- BonJour, Laurence (1980). "Externalist Theories of Empirical Knowledge", *Midwest Studies in Philosophy* 5: 53–73.
- Bradford, Gwen (2015). *Achievement*. Oxford: OUP.
- Celikates, Robin, Haslanger, Sally, and Stanley, Jason (eds.), *Analyzing Ideology*. Forthcoming from OUP.
- Cohen, Stewart (1984). "Justification and Truth". *Philosophical Studies* 46: 279-296.
- Fricker, Miranda (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: OUP.
- Friedman, Jane (fc). "The Epistemic and the Zetetic". Forthcoming in *The Philosophical Review*.
- Goldman, Alvin (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Kornblith, Hilary (1983). "Justified Belief and Epistemically Responsible Action". *The Philosophical Review* 92(1), 33-48.
- Mason, Elinor (2019). *Ways to be Blameworthy: Rightness, Wrongness, and Responsibility*. Oxford: OUP.
- Moody-Adams (1994). "Culture, Responsibility, and Affected Ignorance". *Ethics* 104(2): 291-309.
- Nozick, Robert (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Singer, Daniel J. and Aronowitz, Sara (fc). "What Epistemic Reasons are For: Against the Belief-Sandwich Distinction". Forthcoming in Dunaway and Plunkett (eds.), *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*.
- Srinivasan, Amia (2020). "Radical Externalism". Forthcoming in *The Philosophical Review*.
- Toole, Briana (fc). "From Standpoint Epistemology to Epistemic Oppression". Forthcoming in *Hypatia*.