

Radical Externalism

Amia Srinivasan

University College London

ABSTRACT: This paper presents a novel challenge to epistemic internalism, the view that epistemic justification supervenes on facts to which the believing agent has introspective access. The challenge rests on a new set of cases which feature subjects forming beliefs under conditions of ‘bad ideology’ – that is, conditions in which pervasively false beliefs sustain and are sustained by systems of social oppression. In such cases, I suggest, the externalistic view that justification is a matter of structural, worldly relations, rather than the internalistic view that justification is a matter of how things seem from the agent’s individual perspective, becomes the more intuitively attractive theory. But these ‘bad ideology’ cases do not merely yield intuitive verdicts that favour externalism over internalism. These cases are moreover *analogous* to precisely those canonical cases widely taken to be counterexamples to externalism: cases featuring brains-in-vats, clairvoyants, and dogmatists. That is, my ‘bad ideology’ cases are, in all relevant respects, just like cases that are thought to count against externalism – except that they intuitively *favour* externalism. This, I argue, is a serious worry for internalism, and bears interestingly on the debate over whether externalism is a genuinely ‘normative’ epistemology.

It is impossible by a mere individual...effort
to escape from the web of the social lie
TROTSKY¹

1. A new challenge for internalism

Consider the following case:

RACIST DINNER TABLE: Nour, a young British woman of Arab descent, is invited to dinner at the home of a white friend from university. The host, Nour’s friend’s father, is polite and welcoming

¹ As quoted and translated by Edmund Wilson (2003), p 438, from Trotsky’s *Biography of Lenin, Volume 1* (1936).

to Nour. He is generous with the food and wine, and asks Nour a series of questions about herself. Everyone laughs and talks amiably. As Nour comes away, however, she is unable to shake the feeling that her friend's father is racist against Arabs. But replaying the evening in her head she finds it impossible to recover just what actions on the host's part could be thought to be racist, or what would justify her belief in the host's racism. If pressed, Nour would say she just had a strong feeling that her host was racist: that she 'just knows'. In fact the host is racist – he thinks of Arabs as inherently fanatic, dangerous and backwards – and as a result did treat Nour in ways subtly different to how he would have treated a non-Arab guest. And in fact Nour is subconsciously sensitive to this subtly racist behaviour. It is this subconscious sensitivity that lead to her belief that her host is racist.

Here is my question: is Nour's belief that her host is racist (epistemically) justified? I think the intuitive answer is yes. Nour's belief, after all, is the product of a sensitivity to racism, a sensitivity that allows her to dependably track whether or not the people she encounters are racist. It would seem odd to say that Nour *ought not*, epistemically speaking, have formed the belief that her host is racist, or that she did something epistemically *impermissible* in forming this belief. And it would seem similarly odd to say that, having formed the belief, Nour ought to now, epistemically speaking, give it up. Indeed, it seems right to say that if Nour were to give up her belief in her host's racism, she would be losing an item of knowledge. Nour not only believes but *knows* that her host is racist. If so, it follows that Nour's belief must be justified.²

Of course, Nour has no awareness, introspective or otherwise, of how her subconscious racism-detection mechanism works – nor even *that* it works. (As she says, she 'just knows'.) We might well want to say that Nour would be better off, epistemically speaking, if she had such a higher-order awareness. Perhaps such an awareness would give Nour's cognitive economy a greater degree of overall coherence,³ or a greater robustness against

² I assume throughout that justification is a condition on knowledge.

³ Sosa 2009 on the virtues of 'reflective' over 'animal' knowledge. Cf Goldman 1988 on the distinction between 'weak' and 'strong' justification.

misleading counterevidence.⁴ And yet it seems counterintuitive to infer from the fact that Nour could be epistemically better off in these ways to the conclusion that her belief as it stands is unjustified. For Nour's belief that her host is racist is not only true, but *non-luckily* so. Nour's subconscious sensitivity to racism means that her belief gets on to the truth not as a mere matter of chance, not as a happy accident, but as a matter of predictable dependability. Surely then, her belief is justified.

Now consider a second case:

CLASSIST COLLEGE: Charles is a young man from a working-class background who has just become the newest fellow of an Oxford college. He is initially heartened by the Warden's explicit commitment to equality and diversity. The Warden assures him that, though the college is still dominated by wealthy fellows, Charles will be welcomed and made to feel included. Indeed, the Warden tells Charles, he too is from a working class background, and has experienced plenty of discrimination in his time. Charles is confident not only that the college will be a good community for him, but also that the Warden is a person of excellent judgment on these matters. However, a few incidents soon disrupt Charles' rosy view of things. At high table, when Charles explains that he went to a state school, a fellow responds with 'but you're so well-spoken!'. At a visit to the pub, a number of young fellows sing the Eton boating song while Charles sits uncomfortably silent. Finally, Charles hears that the other fellows call him 'Chavvy Charles'. Charles, who has a dependable sensitivity to classism, goes to the Warden to report that he has experienced a number of classist incidents in college. Shocked, the Warden asks him to explain what happened. But when Charles describes the incidents, the Warden laughs, assuring him that no harm was meant and none of these are genuinely classist incidents. He tells Charles that he is being overly sensitive – something to which (the Warden goes on) Charles is

⁴ One might also think that Nour would need such higher-order awareness in order for the fact of her host's racism to constitute what Grice (2001, ch 3) called a 'personal' reason for Nour to act – for example, by declining future supper invitations. (Thanks to Mikkel Gerken for this point.) My own view is that Nour's knowing that her host is racist is sufficient to make this a practical reason for Nour.

understandably prone to being, given his working class background. Charles leaves the conversation unmoved, continuing to believe that he has faced classist discrimination in the college, and dismissing the Warden's testimony. Charles meanwhile is unaware that some people from working class backgrounds (e.g. the Warden) suffer from false consciousness, distorting their ability to recognise class-based oppression.

Is Charles's ultimate belief that the college is classist justified?⁵ I think the intuitive answer is, again, yes. Like Nour, Charles forms a true belief on the basis of a dependable sensitivity to instances of classism. Of course, unlike Nour, Charles maintains his true belief in the face of seemingly credible, misleading evidence – namely, the Warden's testimony to the effect that Charles' belief is actually formed on an unreliable basis, viz. an *oversensitivity* to classist slights. And yet, this does not seem intuitively to affect the justification of Charles' belief. Charles is not only justified in forming the initial belief that the college is classist; he intuitively remains justified even after the Warden gives his misleading testimony to the contrary. Indeed, as with Nour, it seems right to say that Charles *knows* that the college is classist, both before and after the Warden's attempt to explain away Charles' belief.

Of course, in dismissing the Warden's testimony, we might think that Charles exhibits a mild form of epistemic dogmatism, a dogmatism that could, if indulged in other circumstances, lead him to recklessly dismiss non-misleading evidence and court ignorance. We might also think that Charles would be better off, epistemically speaking, if he had available to him the phenomenon of false consciousness as a debunking explanation of the Warden's testimony: if he were in a position to *explain away* the Warden's testimony, rather than simply dismiss it. Perhaps we even think that Charles is somewhat blameworthy, epistemically speaking, for this act of dogmatism. And yet, none of this intuitively precludes Charles from justifiably believing – indeed, I think, *knowing* – that the college is classist. Charles' belief that the

⁵ For the sake of simplicity, I am going to use the phrase 'the college is classist' as shorthand for the fact that Charles experienced a series of classist incidents in the college. Of course, just what it is for an institution to be classist, racist, sexist, etc., and just how this relates to the classism, racism, sexism, etc. of its constituent members, is a complicated issue.

college is classist, like Nour's belief that her host is racist, is true not as a matter of good luck, of happy accident, but as a function of her capacity to dependably get on to the truth. This intuitively seems sufficient to justify it.

Bear with me for a third, and final, case:

DOMESTIC VIOLENCE: Radha is a woman who lives in rural India. Her husband, Krishnan, regularly beats her. After the beatings, Krishnan often expresses regret for having had to beat her, but explains that it was Radha's fault for being insufficiently obedient or caring. Radha finds these beatings humiliating and guilt-inducing; she believes she has only herself to blame, and that she deserves to be beaten for her bad behaviour. After all her parents, elders and friends agree that if she is being beaten it must be her fault, and no one she knows has ever offered a contrary opinion. Moreover, Radha has thoroughly reflected on the issue and concluded that, given the natural social roles of men and women, women deserve to be beaten by their husbands when they misbehave.⁶

Is Radha's belief that she deserves to be beaten justified? I think the answer is: surely not. For Radha's belief is not merely false, but moreover the product of a convincing, and systematic, misogynistic illusion: that it is men's place to subordinate women. This illusion – one that infects not only the testimony of Radha's peers and respected elders, but her moral emotions (shame, remorse) and best attempts at rational reflection – ensures that Radha has no dependable access to the moral facts of her situation. Radha, despite her own best efforts, is tragically cut off from moral reality.

Of course, Radha's false belief is hardly her fault; it is not only explained but obviously *excused* by the misogynistic illusion of which she is a victim. Radha is doing, we might want to say, the best she can, given her own distorted epistemic connection to the world. Her belief is eminently understandable; indeed we would be naïve to expect anything better of ourselves in Radha's position. And yet none of this is the same as saying that

⁶ On contemporary attitudes towards domestic violence in India, see <http://data.unicef.org/child-protection/attitudes.html>

Radha's belief is justified.⁷ Intuitively, it is not. After all, if Radha's belief *were* justified, then it would only be its falsity that made the belief fall short of knowledge. We would say: if only it were *true* that Radha deserved to be beaten, she would then *know* it. But this, I think, is implausible. Radha doesn't know that she deserves to be beaten, not merely because her belief is false, but also because her belief betrays her systematically distorted connection to reality. Radha's belief doesn't amount to knowledge not only because it is false, but also because it is unjustified.

These three cases – RACIST DINNER TABLE, CLASSIST COLLEGE and DOMESTIC VIOLENCE – together present, I want to suggest, a serious challenge to a widely held view in epistemology. 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.^{8,9} 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

⁷ The commonsensical distinction between justification and excuse is one to which Austin (1956) famously exhorted philosophers to attend. If I run over your dog while carefully backing out of my drive, I might be excused for killing him, but I certainly wouldn't be justified for so doing.

⁸ According to the *mentalist* variety of internalism, justification supervenes on mental states, whether introspectively accessible or not (Conee and Feldman 2001). On my way of drawing the internalist/externalist distinction, mentalism counts as an externalist doctrine, and so is immune to my challenge.

⁹ If there are no non-trivial conditions that are *transparent* – i.e. such that one is always in a position to know whether one is in that condition (Williamson 2000, Srinivasan 2015) – then it follows from internalism that there are no non-trivial general truths about the supervenience base for justification (Srinivasan 2015; cf. Greco 2017).

an appropriate causal connection to its content,¹⁰ or is a product of a reliable or safe method.¹¹

It is generally thought, by internalists and externalists alike, that intuitive reflection on a range of well-known cases – brains-in-vats, unwitting clairvoyants, dogmatists, and so on – supports internalism over externalism. Meanwhile, the case for externalism is largely theoretical, resting primarily on externalism’s ability to neatly dispatch with sceptical threats, albeit at the cost of counterintuitive verdicts on cases. But the three cases I described above – RACIST DINNER TABLE, CLASSIST COLLEGE and DOMESTIC VIOLENCE – disrupt this tidy view of things. For externalism has a much easier time of vindicating what I take to be the intuitive verdicts on these cases than does internalism. In RACIST DINNER TABLE and CLASSIST COLLEGE, the subjects have a belief that is, *ex hypothesi*, reliably and safely connected to the truth. It is thus no mystery, from the externalist perspective, how such beliefs could be justified, since they straightforwardly satisfy the typical externalist conditions on justification.¹² In DOMESTIC VIOLENCE, meanwhile, the subject fails to exhibit such a connection between her belief and the truth, even while doing as best as she can by her own lights. Thus it is again no mystery, from the externalist perspective, how her belief could fail to be justified.

The internalist, meanwhile, has a more difficult time of it. How could Nour or Charles be *justified*, according to the internalist, given that neither has any introspective access to the reliable grounds of their beliefs¹³ – and indeed, in Charles’ case, has strong internalistic reason to think he is *not* so reliably grounded? And how could Radha be *unjustified*, according to the internalist, given that she is doing as well as she can, given how things seem

¹⁰ Goldman 1967.

¹¹ On reliability theories see Armstrong 1973 and Goldman 1975. On safety theories see Sosa 1996 and 2000, Williamson 2000 and Pritchard 2005.

¹² Externalistic treatment of cases like CLASSIST COLLEGE is complicated by the question of whether justification can be *defeated* by misleading evidence, a question on which externalists differ. I will discuss these complications shortly. Suffice it to say for now that at least some externalists would say that Charles’ belief retains its justification in the face of the Warden’s testimony.

¹³ A *permissive* internalist might want to protest that Nour *does* have introspective grounds for her belief that the host is racist: namely, the fact that she has a strong *inclination* to so believe. See fn. 19 for a discussion of why this is an unpromising internalistic response.

from her own limited view on things? The internalist appears to be faced with a choice between biting the bullet on these cases – conceding that they provide intuitive support for externalism – or trying to find a way to vindicate the ‘externalistic’ intuitions in a way consistent with internalism.

But the real challenge for internalism lies elsewhere. My three cases are not merely recalcitrant to internalistic treatment. They are moreover analogous to those very cases that internalists have canonically presented as counterexamples to *externalism*. Consider, for example, one of Laurence Bonjour’s famous such counterexamples:

CLAIRVOYANT: Norman, under certain conditions that usually obtain, is a completely reliable clairvoyant with respect to certain kinds of subject matter. He possesses no evidence or reasons of any kind for or against the general possibility of such a cognitive power, or for or against the thesis that he possesses it. One day Norman comes to believe that the President is in New York City, though he has no evidence either for or against this belief. In fact the belief is true and results from his clairvoyant power, under circumstances in which it is completely reliable¹⁴

Because Norman’s belief is based on a reliable method, the typical externalist will say that it is justified.¹⁵ But many find this externalist verdict absurd. Bonjour writes that Norman’s belief is ‘epistemically irrational and irresponsible, and thereby unjustified’.¹⁶ For it is part of one’s epistemic duty, he goes on, to ‘reflect critically upon one’s beliefs, and such critical reflection precludes believing things to which one has, to one’s knowledge, no reliable means of epistemic access’.¹⁷ Norman’s belief, Bonjour says, is from Norman’s own perspective nothing more than an ‘unfounded hunch’ (ibid). Thus Norman’s belief, *pace* the externalist, is unjustified.

And yet, CLAIRVOYANT is analogous to RACIST DINNER TABLE, in which, recall, Nour’s belief is intuitively *justified*. Both Norman and Nour

¹⁴ Bonjour 1980, 62.

¹⁵ Bonjour intends CLAIRVOYANT as a counterexample to reliabilist versions of epistemology, but it can be thought of as a potential counterexample, *mutatis mutandis*, to other varieties of externalism as well.

¹⁶ Ibid 63.

¹⁷ Ibid.

exhibit a sensitivity to the truth, a sensitivity of which they are unaware but that nonetheless produces reliably true beliefs. From her internal perspective, Nour's belief is no better, BonJour would presumably say, than a 'hunch', and must be therefore unjustified. And yet this seems intuitively false: Nour's belief that her host is racist seems eminently *justified*.^{18,19}

But how could it be that Norman's belief is any less justified than Nour's? Indeed it cannot be so. Insofar as these cases are analogous in the relevant respects, Nour and Norman's justification must stand or fall together. It is not enough, then, for the internalist simply to bite the bullet on RACIST DINNER TABLE, concluding that it intuitively supports externalism while cleaving nonetheless to internalism. If the internalist wants to continue to use CLAIRVOYANT as evidence against externalism and in favour of internalism, she needs to say something about RACIST DINNER TABLE. Either the internalist needs to tell us why RACIST DINNER TABLE is in fact relevantly disanalogous to CLAIRVOYANT, or she needs to offer us an

¹⁸ An externalist might protest that CLAIRVOYANT and RACIST DINNER TABLE are disanalogous in that while Norman's clairvoyant beliefs simply pop into his head, Nour (I've said) is responding to 'subtle cues'. But in so saying I don't mean to suggest that Nour enjoys 'my-host-is-racist' phenomenology, any more than Norman enjoys 'President-is-in-New-York' phenomenology. Nour's responsiveness to 'subtle cues' is entirely a matter of subconscious processing; the workings of Nour's racism-detection are as opaque to Nour's conscious mind as are the workings of Norman's clairvoyance is to Norman. Thanks to Ted Poston for raising this worry.

¹⁹ Not all internalists need agree with BonJour that Norman and Nour's beliefs are unjustified. BonJour endorses a strong form of internalism, according to which S's belief that p is justified iff S has (independent and undefeated) reason to believe that her belief was formed on a reliable basis. According to the weaker *permissive* internalism, S's belief that p is justified if S is inclined to believe p and S lacks any undefeated evidence that her belief that p is defective. Permissive internalism implies that both Norman and Nour's belief are, as the externalist thinks, justified. Unfortunately, permissive internalism is implausible. It implies, for example, that my hunch that my lottery ticket is a winner is sufficient to justify my belief that it is a winner. It also implies that anyone who has a 'hunch' that they are being racially discriminated against – for example, a white person who usually enjoys white privilege and suddenly finds herself being treated as a racial equal – is justified in her belief. Thus permissive internalism does not seem a likely refuge for the internalist.

error theory as to why the intuitions elicited by RACIST DINNER TABLE are not to be trusted.²⁰

I will discuss what I take to be the best prospects for such an internalist response below. For now let me return to CLASSIST COLLEGE and DOMESTIC VIOLENCE, which, I want to argue, are also analogous to cases traditionally thought to favour internalism over externalism. Recall that in CLASSIST COLLEGE Charles has an intuitively justified (and indeed intuitively knowledgeable) belief that his Oxford college is classist, despite the misleading testimony of the Warden to the contrary. Charles' belief is intuitively justified, I suggested, because it is based on Charles' dependable sensitivity to instances of classism. But CLASSIST COLLEGE is analogous to a kind of case that is canonically thought to elicit the opposite intuition, and thereby impugn externalism. Compare:²¹

DOGMATIST: At a time t_1 Mary walks into an art gallery and sees a red sculpture. There is nothing abnormal about Mary's perceptual faculties or the lighting conditions in the gallery. Thus she forms a true belief that the sculpture is red. At a slightly later time t_2 a gallery assistant tells Mary that the sculpture is illuminated by a hidden red light, such that any object it shines on would look red even if it weren't. Mary ignores the misleading testimony and continues to believe, on the basis of her reliable perceptual faculties, that the sculpture is red. What Mary does not know is that the exhibition – including the gallery assistant's misleading testimony – is being put on by a famous artists' collective dedicated to epistemic hoaxes.

²⁰ Of course, the externalist who wants to use RACIST DINNER TABLE as part of her case against internalism will have to say something similar about CLAIRVOYANT: that is, she is under pressure to explain why RACIST DINNER TABLE, but not CLAIRVOYANT, elicits unreliable, internalistic intuitions about justification. In the course of this paper I will offer such an error theory on behalf of the externalist.

²¹ There are many versions of this case, but mine follows most closely Lasonen Aarnio's "Trick on Suzy" (2010, 1). Cf. Chisholm (1966: 48), Bonjour (1980: 59–60), Pollock (1995: 41) and Pollock and Cruz (1999: 44). Lasonen Aarnio, unlike the other examples, embraces the (pure) externalist verdict that her protagonist continues to *know* in the face of the misleading testimony – a case of what she calls 'unreasonable knowledge'.

The standard intuitive verdict on DOGMATIST is that Mary's belief, while initially justified at t_1 , loses justification at t_2 , when the misleading testimony is delivered. After all, how could Mary's belief that the sculpture is red be justified after she is told by a seemingly reliable expert that her perceptual capacities are unreliable? In continuing to believe that the sculpture is red, doesn't Mary ignore evidence that, from her perspective, bears squarely and damningly on her belief? Isn't her belief, even if true, both irresponsible and blameworthy? And if so, how could it be justified?

Indeed, the intuition that Mary (and other similar dogmatists) are unjustified is generally thought so compelling that most externalists feel pressure to modify their externalism in order to vindicate it. What we might call a *pure* externalism says that the satisfaction of the externalist condition (e.g. reliability, safety) is both necessary and *sufficient* for justification. According to pure externalism, Mary's belief at t_2 – because it is still based on her reliable (safe, etc.) perceptual faculties – retains its justification.²² Cases like DOGMATIST prompt most externalists to reject pure externalism, instead favouring a modified, internalistically-inspired theory according to which S's belief is justified just in case S's belief satisfies the externalist justification-condition *and* S is not in possession of strong misleading evidence to the contrary.²³ This modification allows the moderate externalist to vindicate the intuition that Mary's belief is justified at t_1 but loses its justification at t_2 . Meanwhile, only a small minority of externalists are willing to bite the bullet on DOGMATIST-type cases in order to maintain pure

²² A pure externalist could argue that Mary loses justification at t_2 (and *mutatis mutandis* for other DOGMATIST-type cases) because the method on which Mary's belief at t_2 is based (which involves dismissing misleading evidence) is in fact unreliable or unsafe. Thus DOGMATIST turns out to be a case of Mary switching from a justification-conferring method to a justification-depriving method. I set aside this possibility in what follows, assuming that Mary's belief-forming method is stable from t_1 to t_2 .

²³ For some examples of moderate externalists, see Alston 1988, Bergmann 2006, Goldman 1986: 62-3 and 111-2, and Nozick 1981, 196. Some externalists (e.g. Bergmann and Goldman) endorse defeat conditions that are even more liberal than that endorsed by the moderate externalism I describe, allowing that even *unjustified* higher-order beliefs (e.g. those beliefs generated by baseless paranoia) are sufficient to destroy justification.

externalism.²⁴ DOGMATIST-type cases are widely thought to show that pure externalism must be false.

And yet, DOGMATIST is analogous to CLASSIST COLLEGE, in which Charles' belief, recall, intuitively *retains* its justification, despite the misleading evidence, just as the pure externalist predicts. How could Charles and Mary differ in justification? Both Charles and Mary use their properly functioning capacities – to detect classism, to detect colour – to arrive at their respectively true beliefs. They are then both met with testimony, from sources they have strong antecedent reason to believe to be trustworthy, to the effect that their belief-forming capacities are in fact unreliable. They both dismiss this misleading evidence, despite the fact that neither has available to them the proper explanation of why the evidence is misleading – that the Warden is suffering from false consciousness, that the 'gallery assistant' is part of the art piece – maintaining their original beliefs. So it seems that Mary's belief cannot be any less justified than Charles'. Either both beliefs are justified, or neither are. The internalist who wishes to use DOGMATIST as a counterexample to (pure) externalism will have to explain why it is that an analogous case appears to be a counterexample not to externalism, but internalism.

Third and finally, recall that in DOMESTIC VIOLENCE, Radha has an intuitively unjustified belief that she deserves to be beaten, a belief that is a symptom of a systematically distorted relationship to reality. Now consider the following case:

BRAIN-IN-A-VAT: Jane_{BIV} is a handless brain-in-a-vat, subjected to a compelling, electrochemical illusion to the effect that she is a normally embodied person

The external world sceptic asks how it is that Jane, a normally embodied person, can know that she has hands given that, for all Jane knows, she could be Jane_{BIV}. The externalist answers that because Jane is, *ex hypothesi*, a normally embodied person – i.e. because Jane is in fact not Jane_{BIV} – Jane's belief that she has hands enjoys a (reliable, safe, etc.) connection with the external world which in turn secures justification. While Jane and Jane_{BIV} are internal duplicates, the externalist insists, their beliefs enjoy different justificatory standing. In turn this explains, the externalist goes on, how it is

²⁴ e.g. Plantinga 1986, Lasonen-Aarnio 2010, Williamson 2014.

that Jane can know she has hands despite the fact that Jane_{BIV}'s situation is subjectively indiscriminable from Jane's.²⁵ This is what I meant when I earlier said that externalism is capable of neatly dispatching with sceptical threats.²⁶

But this capacity to fend off sceptical threats comes at a price, a price that strikes most internalists as extortionate. For typical externalist theories imply that Jane_{BIV}'s belief that she has hands is *unjustified*: for Jane_{BIV}'s belief, unlike Jane's belief, lacks the appropriate externalist connection to the world. But how could it be, the internalist asks, that Jane and Jane_{BIV} differ in justification, when (it is granted by all parties) things appear just the same for both Jane and Jane_{BIV}, when Jane_{BIV} is entirely blameless for her belief, when Jane_{BIV} seems to be acting no less responsibly, epistemically speaking, than Jane, and when it is just a matter of bad luck that Jane_{BIV} is envatted rather than embodied? Stewart Cohen famously called this the 'new evil demon problem' for externalism.²⁷ 'It strikes me as clearly false' Cohen writes, 'to deny that [the brain-in-a-vat's] beliefs could be justified. If we have every reason to believe e.g., perception is a reliable process, the mere

²⁵ Is Jane's situation also subjectively indiscriminable from Jane_{BIV}'s situation? The typical externalist will think not. Jane is able to know that she has hands, and from this (the externalist thinks) she can knowledgeably infer that she is not Jane_{BIV}. (This higher-order belief will also have to meet the relevant externalist condition for justification in order to constitute knowledge, but on the externalist's view this need not be particularly difficult.) Thus the externalist will object to the sceptic's initial claim that Jane, 'for all she knows', could be Jane_{BIV}.

²⁶ In what sense is the internalist less able than the externalist to dispatch with sceptical challenges? The internalist who thinks that Jane is not in a position to justifiably believe that she has hands obviously fails to counter scepticism. But what of the internalist who thinks that both Jane and Jane_{BIV} have justified beliefs, and that what makes the difference between their cases is that Jane's belief is true (and thus knowledgeable) while Jane_{BIV}'s belief is false? The problem for this sort of internalist is that she will find it difficult to say – unlike the externalist (see fn. 25) – that Jane can continue to know she has hands once the sceptical threat has been made salient to her. For even if Jane has a *prima facie* right to believe in accordance with appearances, and thus has an initially justified belief that she has hands, once she encounters the sceptic it would seem less (internalistically-speaking) responsible of her to maintain her belief. I don't take this to constitute an insurmountable challenge to internalism, but it does show why externalism has a relatively easier time dealing with scepticism.

²⁷ Cohen 1984.

fact that unbeknown to us it is not reliable should not affect [our] justification' (ibid, 281-2). Cohen explains that this judgment 'hinges' on his 'viewing justification as a normative notion' (ibid 282). So long as one's belief is 'appropriate to the available evidence', he says, one cannot be 'held responsible for circumstances beyond [one's] ken' (ibid). Justification, in other words, cannot be a matter of facts unavailable to one: the externalist's verdict on BRAIN-IN-A-VAT must be false.

And yet, BRAIN-IN-A-VAT is analogous to DOMESTIC VIOLENCE, where the intuitive verdict was, I suggested, that Radha's belief is *not* justified. Radha and Jane_{BIV} are both victims of systematic illusions; Jane_{BIV} is literally envatted, while Radha is, as it were, metaphorically envatted in misogynistic ideology. Both of their beliefs are internalistically impeccable: both Jane_{BIV} and Radha believe in accordance with how things seem to them, do not neglect any evidence, and do as well (epistemically speaking) as we would in their shoes (or vat). Why then should Radha's belief that she deserves to be beaten be any less justified than Jane_{BIV}'s belief that she has hands? The internalist who wants to use BRAIN-IN-A-VAT as part of her argument against externalism – who wants, that is, to insist that the 'new evil demon problem' really is a problem – needs to explain just why it is that our intuitions about DOMESTIC VIOLENCE appear to favour externalism over internalism.

Together, DOMESTIC VIOLENCE, RACIST DINNER TABLE and CLASSIST COLLEGE present a serious challenge for internalism. As I have said, these cases are more straightforwardly and intuitively handled by externalism than by internalism. More pressingly, these new cases are analogous to precisely those familiar cases that are generally thought to be counterexamples to externalism. The internalist who wants to continue to treat the traditional cases as evidence against externalism is thus under pressure to say something about my new cases. Here she has two options: to show that my new cases are in fact relevantly disanalogous to the traditional cases, or to show how the externalistic intuitions the new cases elicit can be explained away by a compelling error theory.

Having set out this new challenge to internalism, the remainder of this paper proceeds as follows. In §2 I offer a diagnosis as to why our intuitive verdicts diverge between these two sets of cases, old and new. In the new cases, subjects are operating under what we might call conditions of *bad ideology*: that is, conditions in which pervasively false beliefs sustain and are

sustained by systems of social oppression. When we consider subjects operating under such conditions, I want to suggest, the externalist verdict that justification is a matter of an agent's structural relationship to the world becomes much more intuitively appealing than the internalist verdict that justification is a matter of how things seem from the agent's individual perspective.

I then go on to canvass strategies for an internalist response to my challenge. In §3 I discuss a possible disanalogy between the traditional and new cases: that my cases, unlike the traditional cases, involve beliefs with moral content. This offers the epistemic internalist who is also a *moral rationalist* a potential strategy for vindicating the intuitive verdicts on my cases without jettisoning her internalism. In §4 I discuss the possibility that my 'bad ideology' cases are infected by political confounds: that our intuitions about them are untrustworthy precisely because the cases are politically charged. I argue that this is not so: my 'bad ideology' diagnosis, if anything, suggests that it is our intuitions in the more familiar, apolitical cases that are less trustworthy. In §5 I conclude with a discussion of how my challenge bears on the debate over externalism's status as a genuinely 'normative' epistemology.

2. A diagnosis: justification and bad ideology

What explains the divergence in intuitions across the two sets of cases I discussed in §1? The new cases 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: patriarchy, racism, classism. In such cases, I want to suggest, the salient epistemological question becomes not whether subjects are blameworthy or praiseworthy for their beliefs – whether their beliefs are reasonable by the subjects' own lights – but how these beliefs relate to systems whose function is to distort subjects' access to the truth. What intuitively matters most, in other words, is whether a subject's belief is a product of a distorting, ideological method, or the product of a truth-conducive method that is capable of piercing through systematic distortion.

Thus in DOMESTIC VIOLENCE, what intuitively matters most, epistemically speaking, is not that Radha's belief is perfectly reasonable by her own lights – a thought the externalist is happy to capture by saying her belief

is *excused*²⁸ – but that her belief is the product of a method (the internalisation of patriarchal ideology) that puts Radha systematically out of touch with the truth. This, we instinctively feel, is what matters for justification.²⁹

Meanwhile, in the case of Nour and Charles, we have subjects who are able to reliably get on to the truth despite their bad ideological circumstances. Again, this fact instinctively matters, epistemically speaking – indeed, it intuitively seems sufficient for their beliefs to be justified. It is perhaps true, as I have suggested, that Nour would be better off, epistemically speaking, if she were aware of the reliability of her belief-forming mechanism. And indeed it is perhaps true that Charles, in dismissing the Warden’s testimony, inculcates in himself a bad epistemic habit, one that might elsewhere undermine his capacity to get on to the truth.³⁰ But what feels counterintuitive is the internalist’s verdict that these deficits, such as they are, preclude Nour and Charles from having justified beliefs – indeed, *knowledge* – of their social realities. The internalist demand that Nour and Charles be aware of the grounds of their belief in order to be justified feels intuitively too high. Given their epistemically unfavourable circumstances, it is enough for justification, we instinctively feel, that they exhibit a reliable sensitivity to the truth.

Indeed, cases like RACIST DINNER TABLE and CLASSIST COLLEGE make salient the way in which internalism can pull us towards not only traditional external world scepticism,³¹ but also scepticism about the social world. Much emphasis has been put by analytic philosophers, in recent years, on the way in which oppression can deprive subjects of epistemic goods like

²⁸ On the externalist notion of excuse, see Williamson 2007.

²⁹ What should we say about *true* ideological beliefs, of the kind that arise because of self-fulfilling beliefs (see Haslanger 2007)? For example, what shall we say about a case in which a man’s belief that his wife is submissive is reliably true, but as a result of his treating her with the expectation that she will be submissive? (This sort of belief is usually accompanied by a further false belief that women are submissive by *nature*. But here I am concerned with the straightforwardly true belief.) A standard externalism would count such a belief as justified, suggesting that externalism can only explain the deficiencies of *false* ideological beliefs, not true ones. I am inclined to accept this implication, and indeed to say that part of what is troubling about self-fulfilling ideology cases is precisely that oppression can give rise to not only *true* but *justified* beliefs that the oppressed possess some undesirable characteristic. Thanks to Rachel Fraser for calling my attention to this issue.

³⁰ Cf. Hawthorne and Srinivasan 2013, p. 21ff.

³¹ See fn. 26.

justification and knowledge.³² But such a recognition should sit-by-side with Marx's thought (a thought he borrowed and refined from Hegel) that the position of the oppressed can afford a dispensation from some of the epistemic ills of the oppressors. A balance between these two thoughts might be achieved by recognising, as I have already suggested in passing, the ways in which the knowledge afforded by oppressed social positions might yet fail to be ideal forms of knowledge: by recognising, for example, that Nour would be better off if she were aware of her own reliable sensitivity to racism, or that Charles would be better off if he were able to explain away the Warden's misleading testimony as a product of false consciousness. The externalist verdict that Nour and Charles nonetheless know allows us to ward off scepticism about the social world while still duly noting the epistemic costs of oppression.

Instructive here is the Marxist literary theorist Fredric Jameson's description of conspiracy theory as 'the poor person's cognitive mapping in the postmodern age...a degraded figure of the total logic of late capital, a desperate attempt to represent the latter's system'.³³ Jameson's point is that conspiracy theorising is at best a second-best attempt on the part of the poor to come to terms, epistemically speaking, with an oppressive capitalist system: a 'degraded' and 'desperate' figure of their reality, but a representational one nonetheless. Conspiracy is neither an ideal form of representation – Jameson says it is 'marked by its slippage into sheer theme and content', variations on worn clichés that preclude genuine analysis – but nor is it merely a symptom of late capitalism. Take for example³⁴ a conspiracy theory that became popular amongst black Americans in the 1980s: that Church's Fried Chicken was secretly owned by the Ku Klux Klan, which in turn laced the chicken with a chemical that would sterilise black men.³⁵ In fact Church's Fried Chicken was not owned by the KKK, and their chicken was not laced with a sterilising drug. And yet black Americans in the 1980s (and, indeed, now) were quite right to believe that there was a coordinated plot against them – a plot orchestrated, in the post-Civil Rights era, through discriminatory government policy on education, housing and policing and the racist behaviour of private white citizens. This is a real plot, of which many black

³² e.g. Fricker 2007, Stanley 2015. Cf. Medina 2013.

³³ Jameson 1988, 356.

³⁴ The example is mine, not Jameson's.

³⁵ Turner 1987.

Americans are presumably well aware, but one significantly harder to describe and articulate than the confected Church's plot. The Church's conspiracy might be thought of as a false *post hoc* rationalisation for a genuine insight into the workings of structural racism in the U.S., a 'degraded figure' of American reality, but a figure nonetheless. Externalist theories of justification, I am suggesting, are well-suited to striking just the sort of balance that characterises Jameson's discussion of conspiracy theorising: a balance that vindicates the knowledge of the oppressed while making it clear the way in which such knowledge can be yet non-ideal. It is partly this anti-sceptical impetus, I want to suggest, that makes the externalist verdict in cases like RACIST DINNER TABLE and CLASSIST COLLEGE intuitively attractive.

I have argued that externalistic verdicts in bad ideology cases are attractive because in such cases what intuitively matters is not how things seem from the agent's own perspective – how well or poorly agents are doing by their own lights – but how the agent relates to the epistemically distorting systems in which they are embedded. Another way of putting this thought is that while internalism is an *individualistic* epistemology, externalism is a *structural* epistemology.

The internalist typically trades in individualistic normative notions – responsibility, blameworthiness, blamelessness, action-guidance – that naturally link up with how things look from the subject's own perspective. Justification, for the internalist, is a meritocratic good: it is available to all minimally competent agents, regardless of their circumstances, distorted or veridical. All that is required to be internalistically justified is individual conscientiousness. Moreover internalistic justification is not something that one can possess through the good luck of veridical circumstances: agents who are 'internally' equal will also be equal in justification. The externalist, meanwhile, does not think of justification in individualistic terms. Justification can come apart from questions of personal responsibility or blamelessness. What matters for externalistic justification, at least in part, is where the subject finds herself in the world – whether the subject, often through no fault of her own, has a veridical or distorted relationship to reality.

When we think of brains-in-vats, or clairvoyants, or dogmatic museum-goers, I'm suggesting, it's easy to think in individualistic terms: in terms of how well the subject is doing from the 'inside'. But when we turn to bad ideology cases, it feels natural to start caring more about how the individual's

belief structurally relates to the world. It becomes intuitive, in other words, to think of epistemic justification in externalistic rather than internalistic terms.

3. In search of a disanalogy

One way for the internalist to respond to the challenge I have set out is to show that my new cases are relevantly disanalogous to the traditional ones: that they differ in a way that makes a difference for justification. One such relevant disanalogy might be that, whereas the traditional cases involve subjects whose beliefs have paradigmatically non-moral content – that the President is in New York, that the sculpture is red, that one has hands – my cases involve beliefs with moral content. Clearly, Radha’s belief that she deserves to be beaten is a moral belief: it is a belief about moral desert. Somewhat more contentiously, so too are Nour’s belief that her host is racist and Charles’ belief that his college is classist.³⁶

Why might this be a relevant disanalogy? According to *moral rationalists*,³⁷ at least some moral truths are – unlike ordinary empirical truths – *a priori* knowable. That is, at least some moral claims can be known to be true or false simply in virtue of our understanding their content. Thus a moral rationalist might think³⁸ that a woman who is embedded in patriarchal ideology is nonetheless in a position to know that she doesn’t deserve to be beaten by her husband, simply through the rigorous exercise of her capacities for *a priori* moral reflection.³⁹ (It’s uncontroversial that such a woman could come to know that she doesn’t deserve to be beaten by joining a feminist consciousness-raising group or reading some bell hooks. But a moral

³⁶ ‘Racist’ and ‘classist’ are generally thought to be ‘thick’ terms, i.e. terms that have both descriptive and (morally) evaluative components. To say that someone is e.g. racist is not merely to say that they believe that certain people are, in virtue of their presumed ancestry, inferior, but that this belief is morally bad, or that the subject is bad in virtue of having this moral belief. But there are also some (albeit eccentric) uses of these terms that do not seem to involve the implicit moral claim, e.g. ‘there’s nothing morally wrong with being racist or classist’. For the sake of argument I assume that beliefs about classism and racism are at least partly moral in content.

³⁷ See e.g. Hare 1982 and Korsgaard 1996.

³⁸ Depending on the scope of her moral rationalism.

³⁹ And on the assumption that she has the requisite understanding of the relevant concepts, especially the concept of moral desert.

rationalist says something much stronger: that such a woman, without changing her actual epistemic situation – without gaining new experiences, new testimony, or new evidence – could come to know that she doesn't deserve to be beaten.)⁴⁰

An internalist who is also a moral rationalist could thus *agree* with the externalist that Radha's belief that she deserved to be beaten is unjustified, but for reasons quite different from those given by the externalist.⁴¹ Namely, Radha's belief is unjustified because it's *a priori* knowable that she doesn't deserve to be beaten. Had Radha simply reflected *properly*, the internalist-cum-moral realist could say, she would have come to the conclusion that she doesn't in fact deserve to be beaten. That Radha concludes otherwise only shows that she hasn't undertaken the requisite reflection; it follows that she lacks (internalistic) justification. In this way Radha is importantly unlike Jane_{BIV}, for whom no amount of assiduous reflection will reveal the truth that she doesn't have hands. Thus DOMESTIC VIOLENCE is, from the perspective of the moral rationalist, relevantly disanalogous to BRAIN-IN-A-VAT. In turn, DOMESTIC VIOLENCE is no threat to internalism, so long as the internalist is willing to embrace moral rationalism.

This is a fine response for the internalist to make, as far as it goes. But notice that it is in effect a denial of the case. I stipulated in DOMESTIC VIOLENCE that Radha thoroughly reflects on her situation and the moral issues at stake, and nonetheless – precisely because she has so thoroughly internalised the misogynistic ideology that surrounds her – is unable to see that she does not deserve to be beaten. The sort of internalist-cum-moral realist I am imagining simply denies this possibility: she thinks that the truth that one does not deserve to be beaten by one's husband is always yielded by proper reflection. This strikes me as false; it seems to me plainly possible for there to be someone like Radha who grasps the relevant concepts, reflects as thoroughly as she can, and still comes to the false belief that she deserves to be beaten. Indeed a case like DOMESTIC VIOLENCE seems to me quite

⁴⁰ By contrast, Marxist (and Marxist feminists) stress that the proletariat (or feminist) standpoint – a position of epistemic privilege from which the underlying political reality can be discerned through the distorting ideological superstructure – is something to be achieved not merely through reflection, but also through political struggle. See Hartsock 1983 for discussion.

⁴¹ Thanks to both Sharon Street and A.J. Julius for raising this objection.

pedestrian, and no less conceivable than, say, BRAIN-IN-A-VAT. (I don't think there are any real-world brains-in-vats, but I do think there are quite a few real-world Radhas.) Insisting on the impossibility of DOMESTIC VIOLENCE strikes me, thus, as a rather large bullet for the internalist to bite. That said, for the internalist willing to do so, the threat represented by DOMESTIC VIOLENCE can be neutralised. Just how many internalists will find this an attractive option remains an open question.

Can moral rationalism likewise rescue the internalist from the challenge posed by RACIST DINNER TABLE and CLASSIST COLLEGE? As I've set out RACIST DINNER TABLE, Nour's belief that her host is racist is based on subconscious processing of subtle behavioural cues, not any kind of rich 'my-host-is-racist' phenomenology, nor any rational intuition or inference stemming from her grasp of the concept *racist*.⁴² In other words, her belief is justified (insofar as it is justified) by her contingent, subconscious capacity to track racist behaviour, and not by any rational capacity to intuit or infer *a priori* knowable truths. Thus it seems unlikely that the internalist will be able to avail herself of moral rationalism to vindicate the intuition that Nour's belief is justified. For even if the moral rationalist thinks that Nour *could* come to have a justified, *a priori* belief that her host is racist, she will not think that Nour, having based her belief on a bit of subconscious processing, does in fact have such a justified belief. In other words, moral rationalism cannot explain why it is that Nour's belief is justified.

What about Charles' belief that the college is classist? In my initial description of CLASSIST COLLEGE I did not say much about how Charles forms his belief, except that it is based on his experience of a series of classist incidents. There is a way of filling out the case that is congenial to the moral rationalist. We could say that Charles' competence with the concept *classist*, together with the relevant empirical knowledge, yields a rational insight that the incidents he experiences are indeed incidents of classism – a rational insight, moreover, that is immune to defeat by misleading testimony. Thus the moral rationalist could explain how it is that Charles is able to continue knowing that the college is classist, while Mary's belief, in DOGMATIST, loses its justification in the face of misleading testimony.

But this is not the only way of filling in CLASSIST COLLEGE. We could alternatively say, for example, that Charles' belief is, like Nour's belief,

⁴² As I'm using the notion, moral rationalism is consistent with moral judgments being justified either non-inferentially (intuitionism) or inferentially.

a product of some subconscious processing. Or we could say that it is a product of a quasi-perceptual moral sense: it is as if Charles *sees* the classism of the incidents. In either case, the intuition that Charles' belief is justified, even in the face of the Warden's misleading testimony, appears to stay firm. And yet, in neither case can the moral rationalist explain why it is that Charles' belief is justified. If Charles' belief is based on subconscious processing, it is obviously not a product of rational intuition or inference. If it is based instead on a kind of quasi-perceptual moral sense, it is again not based on an *a priori* method – one available to all those who possess the relevant moral concepts – but rather a method, like sense perception, that is contingently found in some subjects but not others. In neither case can the moral rationalist explain why it is that Charles' belief is intuitively justified, for in neither case is his belief formed on the basis of an *a priori* method. Thus while moral rationalism can explain why, on some suitable precisification of CLASSIST COLLEGE, Charles' belief is unjustified, it cannot explain why Charles' belief is intuitively justified on other precisifications. Our intuition that Charles' belief is justified does not seem to rest on our explicitly or tacitly assuming that Charles arrives at his belief through some *a priori* method.

Thus the internalist-cum-moral rationalist will, at best, be able to use the moral disanalogy between my cases and the traditional cases to deflect the challenge from DOMESTIC VIOLENCE. At the least, the challenge from RACIST DINNER TABLE and CLASSIST COLLEGE stand.

4. In search of an error theory

The diagnosis I offered in §2 as to why our intuitions differ across the traditional and new sets of cases suggests a promising avenue for an internalist error theory of the intuitions elicited by my cases: political confounds. Specifically, the internalist might argue that our intuitions in the 'bad ideology' cases are unreliable *precisely* because they are cases about bad ideology. Perhaps it just feels too politically unsavoury to say that one can justifiably believe that one deserves to be beaten (as with Radha), or that victims of racism or classism can be unjustified in believing themselves to be such victims (as with Nour and Charles). For political reasons we want to resist the internalist verdicts, and side with the externalist. But this is to

allow – the objection goes – our political judgments to contaminate our epistemological one.

The first thing I want to say by way of response is that it's not at all clear that we can do our epistemological theorising free of ethical or political confounds. After all, it is precisely the thought that it would be *unfair* to withhold justification from the brain-in-a-vat, or *unfair* to confer justification on Norman the clairvoyant, that motivates many to endorse internalism. Indeed Jennifer Nagel et al found that people's willingness to attribute justification to protagonists in sceptical cases correlated positively with measures of empathy.⁴³ Of course, two wrongs don't make a right. Just because the internalist might depend on ethical confounds for her case doesn't mean the externalist should follow suit.⁴⁴ So let us consider a case in which the protagonist *truly* believes something that those who share my leftist political sensibilities would think false:⁴⁵

ABORTION: Thomas is a young man who has grown up in a community gripped by bad leftist ideology. All his life he has been surrounded by pro-choice peers and adults, and exposed only to pro-choice arguments. But Thomas cannot shake the feeling that abortion is, despite what everyone says, wrong. His elders and peers tell him that this feeling must be a manifestation of his deep-seated hatred of women. But in fact it's the result of his genuine sensitivity to the moral truth, namely the truth that abortion really is categorically wrong. Despite all the

⁴³ Nagel et al 2013s

⁴⁴ Thanks to Harvey Lederman for pushing me on this point.

⁴⁵ Since moral truths are presumably necessary truths, if (as I think) abortion is morally permissible, then it is necessarily the case that abortion is morally permissible. Thus to entertain a case in which abortion is morally impermissible, as in ABORTION below, is to imagine a counter-possible case. I take it that such a world is at least conceivable. Alternatively, one could refashion the case to involve a contingent but nonetheless politically-loaded belief, e.g. about man-made climate change. Thus one could imagine a world (non-actual but certainly possible) in which man-made climate change was nothing but a liberal conspiracy, and in which a protagonist – through careful first-hand evaluation of the evidence and in the face of overwhelming testimony and gaslighting – formed a reliable true belief that man-made climate change was a myth. For my part I have the intuition that such a belief would be justified, though I am strongly inclined to believe in the reality of man-made climate change.

misleading evidence from his peers and elders, Thomas cleaves to his belief that abortion is wrong.

Is Thomas' belief that abortion is wrong justified? I think the answer is surely yes. If abortion really were wrong, and it really were just a matter of bad ideology that made leftists think otherwise, then an individual who was genuinely sensitive to the badness of abortion, and formed a belief in its badness on such a basis, would – analogously with Charles in CLASSIST COLLEGE – be justified. And yet the confound error theory would predict that my intuitions wouldn't go this way, because saying that a young man would be justified in his belief that abortion is wrong grates against my political sensibilities. (It does grate somewhat. But I feel the intuition powerfully nonetheless.) ABORTION suggests that what is doing the work in the bad ideology cases isn't leftist political sympathies – after all, there is nothing in the notion of bad ideology that is inherently leftist – but the politically-neutral fact that these cases involve subjects forming beliefs under bad ideological conditions. Whether a subject is embedded in right- or left-wing ideology, we intuitively feel that what matters for justification is their ability to dependably track the world as it really is.

Moreover, there is something to be said for the thought that the political nature of my cases underscores rather than impugns the trustworthiness of our intuitions about them. For my cases are far less *recherché* than their traditional analogues. Subjects operating under conditions of misogyny, racism and classism are a commonplace (or so I think) in the actual world, while brains-in-vats, clairvoyants and epistemology-hoaxing artists are presumably more modally remote. This in turn suggests that if any intuitions should be thrown out as less reliable, it should be those we have in response to the traditional cases, rather than the new, bad ideology cases. The diagnosis I offered in §2, in other words, not only does not impugn the reliability of our intuitions about 'bad ideology' cases. It also casts into doubt the trustworthiness of our intuitions about the traditional, humdrum cases we know and love. The challenge to internalism stands.⁴⁶

⁴⁶ I have been here presuming here a realist metaepistemology, according to which there is one correct account of epistemic justification. According to the epistemic *pluralist*, by contrast, there is more than one kind of epistemic justification, each appropriate to different contexts. An internalistic notion of justification, for example,

5. Conclusion: externalism as a normative epistemology

It's often said that externalism is not a 'normative' epistemology. This is said by both internalists and externalists. Internalists mean it as a criticism: any epistemology worth having must be a normative epistemology, so externalism isn't worth having. Externalists mean it as an explanation: they simply aren't in the game of normative theorising, which is why their theory implies that blameless people (like Jane_{BIV}) can be unjustified and irresponsible people (like Norman and Mary) can be justified. That externalism is not a 'normative' theory is also taken to explain why externalist theories are not 'action-guiding', in the sense of being operationalizable by an agents under conditions of uncertainty or ignorance. By contrast, an internalist epistemology, which articulates its justification-norms in terms of how things seem from the subject's own perspective, can always be used to guide one's own epistemic actions, and neatly tracks whether the subject is blameworthy or blameless.⁴⁷ Internalism treats believers as normative agents, we say, while externalism treats them like mere thermometers.⁴⁸

All this 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 an overly restrictive notion of a normative theory. Marxism, for example, is arguably a normative theory, in the sense that it is responsive to the 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

might be relevant when thinking about first-personal management of beliefs, while an externalistic notion might be appropriate when evaluating agents third-personally as knowers. If one adopts such a pluralist metaepistemology, one might think that my cases show that the externalistic notion of justification is of particular relevance when thinking about how to evaluate beliefs under non-ideal political conditions. One might think, in other words, that externalism gives us the right theory of justification for our political epistemology.

⁴⁷ Though see Srinivasan 2015 for an argument that no theory can satisfy this demand.

⁴⁸ A metaphor that originates with Armstrong (1973).

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. For Aristotle this dependency of the normative goods on ‘external’ facts appears to be a natural consequence of humans’ social existence: the things most worth having are those we can only have through our relationships with others. So too, the externalist thinks, with epistemic goods: the epistemic goods really worth having are not ours to have alone. It is in this sense that externalism is, or can be, a normative epistemology.

One way of reading this paper is as I’ve presented it – as a straightforward argument against epistemic internalism and, *ipso facto*, for externalism. On this way of reading it, my cases present a challenge to the internalist, one that demands that she say something about why these cases do not threaten her view of things. I would be happy to hear from internalists how this challenge might be met. But I would be just as happy to have this paper read as an explanation of why, for at least some of us, externalism is attractive as a genuinely normative theory of justification and knowledge. Insofar as one thinks, as I do, that we live in a world suffused with bad ideology – insofar, we might say, that one’s view of the world is a *radical* one – an epistemology that operates in terms of structural notions becomes more attractive than one that trades in merely individualistic ones. One might be an externalist, in other words, not in *spite* of externalism’s detachment from the individualistic normative notions we hold dear, but precisely because of it.

Acknowledgments

For reading and commenting on multiple drafts, I am deeply indebted to Jane Friedman. For discussion of these issues, my thanks also to Clare Birchall, Paul Boghossian, Daniela Dover, Rachel Fraser, Mikkel Gerken, Jeremy Goodman, A.J. Julius, Harvey Lederman, Nakul Krishna, Lucy O'Brien, Ted Poston, Daniel Rothschild, Robert Simpson, Sophie Smith, and audiences at UCL, KCL, Edinburgh, Southampton, NYU, UCLA and USC.

Works Cited

Alston, W. (1988). "An Internalist Externalism". *Synthese*, 74: 265–283.

Austin, J.L. (1956). "A Plea for Excuses". *Proceedings of the Aristotelian Society*, 57: 1-30.

Armstrong, D.M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

Bergmann, M. (2006). *Justification Without Awareness*. Oxford: Clarendon Press.

BonJour, L (1980). "Externalist Theories of Empirical Knowledge" *Midwest Studies in Philosophy*, 5: 53–73.

Cohen, S. (1984). 'Justification and Truth,' *Philosophical Studies* 46: 279-296.

Conee, E. and Feldman, R. (2001). "Internalism Defended". *American Philosophical Quarterly*, 38(1): 1-18.

Chisholm, R. (1966). *Theory of Knowledge*. Englewood Cliffs, N.J.: Prentice-Hall.

Fricker, M. (2007). *Epistemic Injustice*. Oxford: OUP.

Goldman, A. (1967). 'A Causal Theory of Knowing'. *The Journal of Philosophy* 64(12): 357-372.

Goldman, A. (1979). 'What is Justified Belief?' in G.S. Pappas (ed.), *Justification and Knowledge*. Dordrecht: Reidel, 1-25.

Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

Goldman, A. (1988). 'Strong and Weak Justification'. *Philosophical Perspectives* 2: 51-69.

Greco, D. (2017). 'Cognitive Mobile Homes'. *Mind* 126(501); 93-121.

Grice, P. (2001). *Aspects of Reason*. Oxford: Clarendon Press.

Hare, R.M. (1982). *Moral Thinking*. Oxford: OUP.

Hartsock, N. (1983). "The Feminist Standpoint: Developing the Ground for a Specifically Feminist Historical Materialism", in Sandra Harding and Merrill B. Hintikka (eds.), *Discovering Reality*, D. Reidel Publishing Company.

Huemer, M. 2001. *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.

Hawthorne, J. and Srinivasan, A. (2013). 'Disagreement Without Transparency: Some Bleak Thoughts', in J. Lackey and D. Christensen (eds.), *The Epistemology of Disagreement: New Essays*. Oxford: OUP.

Jameson, F. 1988. "Cognitive Mapping", in Cary Nelson and Lawrence Grossberg, eds. *Marxism and the Interpretation of Culture*. London: Macmillan.

Korsgaard, C. 1996. *The Sources of Normativity*. Cambridge: CUP.

Lasonen Aarnio, M. (2010). "Unreasonable Knowledge". *Philosophical Perspectives* 24: 1-21.

Medina, J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression,*

Epistemic Injustice, and Resistant Imaginations. Oxford: OUP.

Nagel, J., San Juan, V. & Mar, R.A. (2013). "Lay Denial of Knowledge for Justified True Beliefs" *Cognition* 129: 652-661.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Plantinga, A. (1986). "The Foundations of Theism: A Reply," *Faith and Philosophy*, 3(3): 310-312.

Pollock, J. L. & Cruz, J. (1999). *Contemporary Theories of Knowledge*, 2nd edition, Lanham, MD: Rowman & Littlefield.

Pollock, J. L. 1995. *Cognitive Carpentry: A Blueprint for How to Build a Person*, Cambridge, MA: MIT Press.

Pritchard, D. 2005. *Epistemic Luck*. Oxford: OUP.

Pryor, J. 2000. "The Skeptic and the Dogmatist". *Nous* 34(4); 517-549

Sosa, E. 1996. 'Postscript to 'Proper Functionalism and Virtue and Epistemology'', in J.L. Kvanvig (ed.), *Warrant in Contemporary Epistemology*. Lanham, MD: Rowman and Littlefield, 271-280.

Sosa, E. 2000. 'Skepticism and Contextualism', *Philosophical Issues* 10: 1-18.

Sosa, E. (2009). *Apt Belief and Reflective Knowledge*, Volume I and II, Oxford: OUP.

Stanley, J. (2015). *How Propaganda Works*. Princeton: PUP.

Srinivasan, A. (2015). "Are We Luminous?", *Philosophy and Phenomenological Research* 90(2): 294 – 319.

Srinivasan, A. (2015). "Normativity Without Cartesian Privilege", *Philosophical Issues* 25(1), 273-299.

Pryor, J. 2000. "The Skeptic and the Dogmatist". *Nous* 34(4); 517-549

Turner, P.A. (1987). "Church's Fried Chicken and The Klan: A Rhetorical Analysis of Rumor in the Black Community", *Western Folklore*, 46(4): 294-306

Williamson, T. (2000). *Knowledge and Its Limits*. Oxford: OUP.

Williamson, T. (2007). 'On Being Justified in One's Head', in Mark Timmons, John Greco and Alfred Mele (eds.), *Rationality and the Good: Critical Essays on the Ethics and Epistemology of Robert Audi*. Oxford: OUP, 106-122.

Williamson, T. (2014). "Very Improbable Knowing". *Erkenntnis* 79: 971.

Wilson, E. (2003). *To the Finland Station*. New York: The New York Review of Books.