

A Plea for Epistemic Excuses

Clayton Littlejohn
King's College London
cmlittlejohn@gmail.com

[For F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.]

I. Introduction

A typical epistemology course begins with a discussion of the distinction between justification and knowledge and ends without any discussion of the distinction between justification and excuse.¹ This is unfortunate. If we had a better understanding of the difference between justificatory and excusatory defenses, we would have a better understanding of the significance of the intuitions that shape the internalism-externalism debate. My aim in this paper is to fill a lacuna. There has been little discussion of the justification-excuse distinction in epistemology, so I shall offer some suggestions about how that distinction should be drawn.² I shall explain how the kinds of excuses that should interest epistemologists exculpate and what distinguishes excuses from other kinds of defenses.

The topic might not seem very sexy, but I hope that the discussion will be valuable insofar as it should shed some light on the intuitions that underwrite Cohen's (1984) new evil demon argument. Once we have a better understanding of the difference between justificatory and excusatory defenses, we shall see that externalists have nothing to fear from the demon and the intuitions that talk of demons can scare up. The intuitions that underwrite the new evil demon objection cannot be used to determine whether belief is governed by some particular norm (e.g., the truth norm, the knowledge norm, or some evidentialist norm).³ Since the justificatory status of our beliefs depends upon whether they conform to epistemic norms, the intuitions Cohen hopes to elicit cannot tell us whether justification depends upon any of the external conditions that aren't just in the head. Most of what I say here in response to Cohen's argument appeals to distinctions that internalists and externalists recognize. In

¹ I want to thank Maria Alvarez, John Callanan, Christina Dietz, Julien Dutant, Pascal

² Various authors discuss excuses in passing, but there has been little explicit discussion as to how the distinction should be drawn. In Littlejohn (2009, 2012) I tried my hand at explaining how the distinction should be drawn and why we should think of reasonableness/rationality as a crucial part of some standard excuses. See also Sylvan (MS).

³ The intuitions that underwrite Cohen's (1984) argument against externalist accounts of justification live a double life and figure prominently in the literature on the norms that determine when our assertions are warranted and norms that determine whether it's proper to treat something as a reason for action. See Lackey (2007) and MacKinnon (2013), for example. Authors that defend externalist accounts of warranted assertions (e.g., DeRose, Weiner, and Williamson) or externalist accounts of the norms governing practical reason (e.g., Hawthorne and Stanley, Littlejohn) often defend their views from critics by appeal to a distinction between primary and secondary propriety, a distinction that appears for all intents and purposes to be the distinction between justification and excuse. This maneuver has been criticized (e.g., by Gerken (2011)) and I address these criticisms throughout the paper.

arguing that the intuitions he appeals to cannot have the significance he takes them to, I'm not arguing from any externalist assumptions. Regardless of whether you're an internalist or an externalist, you shouldn't think of your positive responses to the relevant error cases as any sort of indication that the relevant beliefs conform to epistemic norms or are justifiably held. In the final section of this paper, I'll introduce a kind of argument against internalism. We'll see that the newest evil demon problem is a problem for internalists about epistemic justification.

2. Cohen's Demon

Cohen's (1984) new evil demon objection is, to my mind, the single most important objection to externalism in the literature. If you bracket the intuitions that underwrite the objection, the remaining objections to externalism don't have any force. If, however, Cohen's intuitions are correct and they have the significance he takes them to, these other objections to externalism would be otiose since Cohen's objection would settle the issue.

While the objection has tremendous dialectical significance, I don't think it has nearly the power that many epistemologists seem to accord to it.⁴ Intuitions about error cases won't help us determine which norms govern belief, so they can't help us determine whether these norms should be understood along internalist or externalist lines. The test to determine whether a norm is a norm isn't whether it can be demon-proofed.

I'll present Cohen's objection as an argument for a non-skeptical version of internalism about doxastic justification. For my purposes, internalism about justification should be understood as the view that internal conditions wholly determine whether your *beliefs* are justified.⁵ There are different ways of drawing the distinction between the internal and the external, but for our purposes, let's say that C is an internal condition iff C supervenes upon your non-factive mental states. C is an external condition iff C isn't internal. Since *being sufficiently reliable*, *being accurate*, *being accepted by the experts*, and *constituting knowledge* aren't things that supervene upon your non-factive mental states, they won't be necessary for justification if the internalists are right. Are they right? Many epistemologists think that they are.

Cohen's attack on externalism begins with an objection to factive accounts of justified belief:

⁴ It is a mistake to try to modify reliabilism to vindicate the intuitions as Comesana (2002) and Goldman (1998) have tried to do.

⁵ I stress belief here because there's a tendency in this literature to move freely from claims about a *person* being justified in believing something and a person's *belief* being justified. The new evil demon argument is supposed to be an argument against externalism about doxastic justification (i.e., the justification of belief) and it's not at all clear how intuitions about whether a *believer* is justified bear on claims about the justification of this believer's *beliefs*. For discussion of the distinction between personal and doxastic justification, see Bach (1985), Engel (1992), Littlejohn (2009), and Lowy (1978). This kind of distinction between two forms of assessment is found in other literatures. See Bennett (1995) for a discussion of the distinction in ethics and Moore (1997) for an argument that the distinction is required to make sense of our intuitions about imperfect defense.

The legacy of the Cartesian view is scepticism. Descartes demonstrated this in the first meditation that no such connection is forthcoming ... Given any plausible specification of C for any S, it will always be logically consistent to suppose that not B. That is what the evil demon argument shows. Where, e.g., C comprises facts about sensory data, and where B is a belief about the truth of some empirical proposition, it is always logically possible that the evil demon has arranged for C to obtain where B is false (1984: 281).

There aren't many externalists who accept the Cartesian view.⁶ Most externalists think it's a non-starter. Let's hope that their reasons for rejecting such views aren't Cohen's for Cohen's objection applies to the remaining externalist views with equal force:

I think that the evil demon hypothesis ... uncovers a defect in the Reliabilist position. We can see this by supposing the hypothesis to be true. Imagine that unbeknown to us, our cognitive processes (e.g., perception, memory, inference) are not reliable owing to the machinations of the malevolent demon. It follows on a Reliabilist view that the beliefs generated by those processes are never justified.

Is this a tenable result? I maintain that it is not ... [P]art of what the hypothesis entails is that our experience is just as it would be if our cognitive processes were reliable. Thus, on the demon hypothesis, we would have every reason for holding our beliefs that we have in the actual world. It strikes me as clearly false to deny that under these circumstances our beliefs could be justified. If we have every reason to believe e.g., perception is a reliable process, the mere fact that unbeknown to us it is not reliable should not affect its justification-conferring status (1984: 281).⁷

If Cohen is right, none of the external conditions that figure in the Cartesian view or your orthodox externalist views could be required for justification. Is he right?

We can state the core objection as follows:

The New Evil Demon Argument

⁶ The proud few include Sutton (2007), Littlejohn (2012), McDowell (1998), Unger (1975), and Williamson (this volume). Unger is an interesting case because he seemed to think that you could not rationally believe what you do not know and thought that virtually nothing could be known. On his view, justification is both factive and something that supervenes upon your non-factive mental states!

⁷ It's odd that so many externalists regard the first objection as decisive against the view that justified beliefs have to be true since they don't seem to think that the underlying objection has much force against their views. If there's a point that Cohen and I probably agree on, it's that externalists are fooling themselves if they think that their views are immune to the difficulties that arise for factive accounts of justification of the sort defended by Littlejohn (2012), McDowell (1998), Sutton (2007), Unger (1975), and Williamson (this volume).

- P1. In the bad case, the subject's beliefs don't have the properties that figure in externalist accounts of justification [Observation].
- P2. The beliefs are equally justified in the good and bad case [Equality].
- P3. The beliefs are justified in the bad case [Sufficiency].
- C. The internal properties are sufficient for justification in the good case and bad. The external properties aren't necessary for justification.

As stated, the argument appeals to two intuitive verdicts. According to the first, Equality, you and your counterparts' beliefs are equally justified in the good case and the bad. This doesn't tell us whether they're justified, mind you, only that they won't differ in justificatory status. According to the second, Sufficiency, the conditions present in the bad case are sufficient for justification. Taken together, they tell us that justification doesn't depend upon any of the external conditions that distinguish the good case from the bad and that the internal conditions common to the two cases ensure that your beliefs are justified in the relevant pairs of cases.⁸

There are good ways to resist this line of reasoning, but let me say straight off that there is *something* going for the subject and the subject's beliefs in the bad case. I have, as I'm sure most of you have, a positive reaction to the subject in the bad case who forms the mistaken beliefs that she does. The issue isn't whether we have this reaction. The issue is what to make of this reaction. Is this a good indication that the subject's beliefs in the bad case are perfectly well justified? No. To determine whether the subject's beliefs are justified, we'd have to determine whether they conform to the norms that govern them. This positive reaction is an unreliable guide when it comes to determining whether some act or attitude conforms to the relevant norms. So, while I think we all have a similar sort of positive reaction when we consider the case, I think the evidential value of this response is nugatory. We shouldn't take it that the positive response is a way of registering that the subject has conformed to the standards that determine whether her attitudes are justified.

The right way to approach questions about justifications and norms is by appeal to *substantive* considerations, considerations that don't seem to figure in Cohen's argument at all. Compare Cohen to Kant. Much in the way that Cohen thinks that everything that bears on the justification of belief is in the head, Kant thought that everything that bears on the justification of action is in the head.⁹ Kant's argumentative strategy was to show that all moral assessment is concerned exclusively with the quality of the agent's will by seeking to show that the agent's will was the only

⁸ I shall argue later that Equality is in tension with Sufficiency. We can add this to the list of reasons *not* to take this argument to be a good argument for a non-skeptical internalist view.

⁹ We shall see below that Cohen rejects this. As he sees things it is *irrational* not to shoot, say, a terrorist disguised as an electrician even if the disguise is perfect and all the evidence indicates (misleadingly, perhaps) that this guy is an electrician. As he sees things, even if you're rationally compelled to believe the guy is an electrician and you're rationally compelled to believe that you're required not to shoot and that it's crazy or irrational to shoot, he thinks that it's irrational not to shoot.

thing of unconditional value.¹⁰ Kant's internalist supervenience thesis which states that the justification of action depends wholly upon what's in the agent's head was supposed to be established on substantive grounds, not epistemic grounds. It explains why there should be epistemic constraints (if it works), but isn't motivated by the thought that there must be such constraints. In contrast, Cohen's argument rests on no substantive assumptions about the norms of belief, the aim of belief, epistemic value, etc. and I think that his methodology must be flawed because it cannot be a test of a genuine norm that it's demon proof. We'll see this below. It's a mug's tame trying to find norms that can both deliver the right verdicts and guide all rational agents so that they can always respond with justification no matter how ignorant or mistaken they might be.¹¹

In §3, I'll explain why we shouldn't be moved by arguments from error. My aim isn't to show that the bad case is a case of excusable belief; rather, it's to show that the positive response elicited by the error cases is just the kind of response we'd have to have to think that an excuse is in order. If there's a case for thinking that the belief is justified, not merely excused, that case has to be made on other grounds. In §4, I address an ancillary argument that Cohen offers in support of his internalist view. In §5, I shall offer some reasons for thinking that the best we could hope for in the bad case is an excuse. My criticisms in §3 don't rest on any externalist assumptions. Even if we knew that internalism was true, we should know better than to rely on arguments like the new evil demon argument. My arguments in §5 are arguments for externalism. They're designed to show that views motivated by the new evil demon argument are mistaken, not merely unmotivated.

3. Pleas

Let's start with Austin:

In general, the situation is one where someone is accused of having done something, or (if that will keep it any cleaner) where someone is said to have done something which is bad, wrong, inept, unwelcome, or in some other of the numerous possible ways untoward. Thereupon he, or someone on his behalf, will try to defend his conduct or to get him out of it. One way of going about this is to admit flatly that he, X, did do that very thing, A, but to argue that it was a good thing, or

¹⁰ Also important were Kant's claims about the proper mode of response to this value. For useful discussion, see Herman (1993) and Gardner (2004).

¹¹ If you want a historical analogy, someone who revises their views about norms and justification in light of the considerations Cohen offers is like someone who revises their views about the moral significance of promises in the way that Ross did when intuitions about error cases led him from the view that there's a reason to keep promises to the view that there's really only reason to try to do the things that would fulfill the promises you think you have. Ross didn't appreciate that there's a trade-off here. If you start trying to revise your norms so that rational judgments about what to do (where 'rational' means supported by the evidence at hand), you'll quickly find that your putative norms deliver the wrong verdicts. See Littlejohn (2014) for discussion of cases of normative ignorance. Cohen agrees that there aren't epistemic constraints on practical norms, but thinks that there are on epistemic norms.

the right or sensible thing, or a permissible thing to do, either in general or at least in the special circumstances of the occasion. To take this line is to justify the action, to give reasons for doing it: not to say, to brazen it out, to glory in it, or the like. A different way of going about it is to admit that it wasn't a good thing to have done, but to argue that it is not quite fair or correct to say baldly 'X did A'. We may say it isn't fair just to say X did it; perhaps he was under somebody's influence, or was nudged. Or, it isn't fair to say baldly he did A; it may have been partly accidental, or an unintentional slip. Or, it isn't fair to say he did simply A—he was really doing something quite different and A was only incidental, or he was looking at the whole thing quite differently. Naturally these arguments can be combined or overlap or run into each other. In the one defence, briefly, we accept responsibility but deny that it was bad: in the other, we admit that it was bad but don't accept full, or even any, responsibility.

Let's note two things about this passage. First, Austin's examples of excuses are a motley bunch. Among the things that Austin thinks can excuse are accidents or slips, nudges, and epistemic defects. Second, Austin works with a *dichotomous scheme*:

- A justification shows that it's not appropriate to criticize the subject for ϕ -ing because there was sufficient reason to ϕ .
- An excuse shows that it's not appropriate to criticize the subject for ϕ -ing even though there wasn't sufficient reason to ϕ .

On this scheme, there are just two kinds of defenses. A justification shows that criticism isn't appropriate because there was sufficient reason for the agent to ϕ . Even the subject committed some wrong in ϕ -ing, the justificatory defense is supposed to show that there was sufficient reason for ϕ -ing in spite of the fact that it was wrongful.¹² In offering an excuse, we concede that the subject didn't have sufficient reason for ϕ -ing but maintain that it wouldn't be appropriate to blame the subject. As Austin understands things, excuses cover every unjustified but blameless wrong. If we apply this to belief, the idea would be that you justifiably believe p iff you believe for sufficient reason.¹³ You should be excused for believing p iff you lack justification but

¹² Bear in mind that wrongs needn't be wrong all things considered.

¹³ I fear that most epistemologists will read 'sufficient reason' as something like 'sufficiently strong evidence'. That's not how I intend it. In the case of action we cannot say whether your reason for ϕ -ing was sufficient until we know what it was *and* what reasons there were that spoke against ϕ -ing. Thus two subjects could ϕ for the very same reasons and only one of them ϕ for a sufficient reason. In the case of belief, I want to allow for the same sort of possibility. It's possible that two subjects believe p for the very same reasons where only one subject's reasons are sufficient. The first subject, but not the second, might have had defeating evidence available to her. It's an interesting question whether the reasons that apply to a subject and prevent her reasons for

are somehow blameless for believing what you do. To embed this within a larger framework, we could start with some norm(s) that govern belief with this form: you shouldn't believe p unless C obtains. These norms alert us to the kinds of conditions that would prevent a belief from being justified and a belief counts as justified iff it doesn't violate any such norm without sufficient reason (where the reasons that might justify such violations would be specified by further norms). A belief is excusably held iff it's held in violation of such a norm but the believer isn't to be blamed for this.

You might have noticed that this tells us nothing about how excuses excuse. Austin hasn't offered any *positive* account of what excuses are or how they get their work done. All he's told us about excuses is that they provide a defense of something there wasn't sufficient reason to do. There's no positive account of the mechanisms that remove responsibility and no indication that excuses share anything in common apart from their exculpatory power.

If we work with Austin's dichotomous scheme and with his open-ended conception of excuse we should be skeptical of *contrast arguments*:

1. This defense, D , differs from some paradigmatic case of an excusatory defense, D' , in some salient way (i.e., the underlying mechanisms by which they remove responsibility differ).
2. Thus, D is not an excusatory defense.
3. All successful defenses are either justifications or excuses.
4. D , if successful, must be a justification.

Anyone who accepts Austin's dichotomous scheme and his gloss on what excuses are will accept (3), but they shouldn't think that there's any good inference from (1) to (2). You wouldn't reason in that way if you thought that we couldn't isolate some single mechanism that's operative in successful excusatory defenses. Once you're alive to the possibility that two successful excusatory defenses might operate in different ways, you'd quickly see that contrast arguments quickly lead to trouble. Case in point. Any sensible person working with this scheme would recognize infancy as an excusing condition. What about cases in which women kill or attempt to kill an abusive partner? Should these women say that their actions were justified or excused? It seems that a justification should only be available in those cases where they had sufficient reason to harm or kill in self-defense. If so, what about the cases where the grounds for self-defense were absent? Should we say that the killing or attempt was excused only insofar as the woman could plausibly claim that she was like an infant in that she lacked the rational capacities necessary for assuming responsibility for her deeds? That would be degrading. The woman who judges on the available evidence that she's in mortal danger and uses force to defend herself shouldn't have to argue that she was somehow like an infant or a mentally ill person to defend herself from criticism if it turns out that the threat was merely apparent or that she could have defended herself with less force. For this sort of reason, it's generally recognized that contrast arguments are just bad arguments. We should be able to offer an exculpatory defense that doesn't amount to a justification that doesn't purport to show that the person

believing from being sufficient to justify belief are all available to the subject or part of her evidence. At this stage, we should leave the issue open. In the paper's final section, I shall point to cases that I think show that the reasons that bear on whether to believe p needn't be accessible to the subject or part of her evidence.

being defended lacks the rational capacities needed to assume responsibility for her deeds, but contrast arguments close off this possibility.¹⁴

It's not hard to find contrast arguments in the literature on epistemic justification.¹⁵ Externalists shouldn't be afraid to say that the subject in the bad case should be excused for believing what she does. To combat this suggestion that Cohen's deceived subject is merely excused, people often argue as follows:

1. The subject's beliefs aren't *just* blamelessly held. These beliefs differ in significant and salient ways from the merely blameless beliefs (e.g., beliefs held by people who are brainwashed).
2. Thus, it would be a mistake to say that the reason the subject shouldn't be blamed for believing what she does is that she should be excused.
3. If they aren't just excusably held, they must be justifiably held.
4. Thus, if we can defend her from criticism, our defense must be justificatory.

In offering examples of excusable epistemic failures, standard examples involve drugs that interfere with your rational faculties and capacities (Steup), incapacitation (Steup), brainwashing (Madison, Smithies), and delusion or insanity (Gerken, Wedgwood). It's obvious that there's some difference between erring because you've been drugged or brainwashed and erring because you've reasoned impeccably from some beliefs formed in response to hallucination, so we're supposed to conclude from this that the difference here is the difference between justification and excuse. As Madison succinctly puts it:

Mere blamelessness is insufficient for the positive epistemic status of the beliefs in cases of subjects and their evil genius-deceived counterparts. It is therefore implausible to suggest that the traditional internalist's judgments of sameness of justification in these cases could arise merely by confusing these two notions (2014: 69).

There are two problems with this kind of argument. First, it seems to assume that all successful non-justificatory defenses are somehow alike (e.g., someone who has been brainwashed, say, should be excused for believing what she does, so other

¹⁴ See Gardner (2007) and Ferzan (2011) for a helpful discussion of the history of the battered woman defense.

¹⁵ See, for example, Bird (2007: 97), BonJour (2002: 248), Cohen and Comesana (2013: 18), Coffman (2010: 3), Fantl and McGrath (2009: 125, fn. 23), Lackey (2007: 606), Locke (forthcoming: 23), Pryor (2001: 117), Rizzieri (2013: 103), Russell (2001: 39), and Wedgwood (2002: 349). Other writers (e.g., Douven (2006)) contrast what's excusable with what's reasonable without explaining why such a contrast is important. It should be noted that while Wedgwood (2002) thinks there's an important contrast to be drawn between cases of, say, insanity and cases of mistaken belief based on hallucination, he focuses on rationality rather than justification. My critical remarks do not apply to his discussion. We can also find contrast arguments in other literatures. See Rosen's (2014) arguments concerning duress, justification, and excuse.

excuses should be like that). We're invited to think that since Cohen's subject isn't anything like a subject who has been brainwashed or drugged, the subject's beliefs couldn't merely be excused. Since they cannot be blamed and cannot be merely excused (or merely blameless), they must be justified. It would, of course, be daft to deny that there's some difference between the cases in which someone follows the evidence and arguments where they lead and forms a false belief and the case in which someone is brainwashed, but it's still not clear why we'd account for this difference by saying that brainwashing excuses and systematic deception justifies. Why shouldn't we say that this is just another case of excuse, albeit one in which the grounds for excusing differ from those that are present in cases of brainwashing? If the category of excuse is sufficiently heterogenous, we could say that insanity and ignorance both function as excuses, albeit excuses that do their work in different ways. Maybe it's right to say that *mere blamelessness* doesn't account for the positive reaction to the case, but maybe it's a mistake to equate excuses with mere blamelessness. All we know about excuses thus far is that they aren't justifications and if the category isn't homogenous, maybe some excuses involve some sort of positive element, something that we might admire, praise, commend, etc.

The second problem with contrast arguments has to do with the kinds of pleas offered to create the contrast. Consider Audi's brief for internalism. In this passage he explains why we should think of false perceptual beliefs based on hallucination as being justified:

[G]iven the vivid hallucination, I am in no way at fault for believing what I do, nor do I deserve any criticism (at least on the non-skeptical assumption that we may generally trust our senses in this way). Far from it. I am like a surgeon who skillfully does all that can be expected but loses the patient. There I should feel regret, but not guilt; I should explain, but need not apologize; and when we know what my evidence was, we approve of what I did.

We consider it reasonable (2001: 23).

Audi's plea says very little directly about whether there was sufficient reason for holding the perceptual belief. Instead, the idea is that once we know about the subject's evidence and how the belief was formed, we don't find fault, we praise rather than blame, we note that the performance was skilled, and so on. It's interesting to note that Cohen's remarks concerning his case and his argument that the deceived subject's beliefs are justified are similar:

My argument hinges on viewing justification as a normative notion. Intuitively, if S's belief is appropriate to the available evidence, he is not to be held responsible for circumstances beyond his ken (Cohen 1984: 282).

If someone offers a justification, you'd expect considerations that would purport to show that the subject had sufficient reason for her actions or attitudes would be front and center. What we find instead is a focus on things like praiseworthiness, character, responsibility, whether the subject could feel guilt, and how reasonable or rational the subject was. Austin wouldn't have thought that these were the marks of justification.¹⁶

¹⁶ In fairness, both Audi and Cohen mention *something* about the subject's grounds, reasons, or evidence. What's not clear from these passages is what role considerations

One reason to think that Austin would have been cool to some of the standard things internalists say about Cohen's case is that Austin's list of excusatory conditions included ignorance and factual mistake (e.g., the subject was unaware that by doing what she was doing she was doing A; the subject saw things differently and so didn't see that she was doing A, this untoward thing). While Austin focused on actions, not attitudes, it seems that if we were to follow Austin's lead the natural thing to say is that humdrum factual mistakes excuse beliefs. If there's some good reason to think of Cohen's victim as someone who has justified beliefs rather than merely excusably held beliefs, I don't think that we'll find support for this in anything that Austin says. We'll have to find support for saying this in spite of what Austin says about mistakes and excusable wrongdoing.

Because of an unfortunate tendency to think that contrast arguments help us identify defenses as justificatory, it might be helpful to consider an alternative scheme to Austin's, Strawson's (1962) *trichotomous scheme*:

- A justification shows that it's not appropriate to criticize the subject for ϕ -ing because there was sufficient reason to ϕ .
- An excuse shows that it's not appropriate to criticize the subject for ϕ -ing even though there wasn't sufficient reason to ϕ and even though she has the intact rational capacities necessary for being held accountable for having ϕ 'd.
- An exemption shows that it's not appropriate to criticize the subject for ϕ -ing on the grounds that she either lacked the rational capacities needed for assuming responsibility for her actions and attitudes or because these rational capacities were compromised.

On this view, there three kinds of defenses: justifications, excuses, and exemptions.¹⁷ To show that something is justified, we have to show that no wrong was committed or that it was done for adequate reason. To show that a subject should be exempted, we'd identify global exculpatory conditions that would preclude us from holding the subject accountable generally. To show that something should be excused (rather than exempted) we'd have to show that there's some exculpatory factor that removes responsibility even though the subject has the rational capacities that would allow us to hold her responsible for her actions and attitudes. If there really are three kinds of exculpatory defenses, it's clear that contrast arguments are bound to fail. A salient difference between the way in which D and D' remove responsibility might show that D and D' belong to different categories, but classifying D as an exemption, say, wouldn't give us any reason to classify D' as a justification rather than an excuse.

about fault, blame, and responsibility have to do with grounds, reasons, or evidence. If they're using considerations about whether the subject should be held responsible to determine whether the subject's beliefs or actions are based on sufficiently strong undefeated reasons, they're making a mistake. The mistake is either conflating two kinds of assessment or in accepting a mistaken philosophical view about how these different sorts of considerations are related to each other.

¹⁷ Some authors refer to exemptions as 'denials of responsibility'. For helpful discussions of the trichotomous scheme, see Duff (2007), Ferzan (2011), Gardner (2007), Horder (2004), Tadros (2005) and Wallace (1998).

We know this more fine-grained scheme has something going for it. People often say that forgetfulness or accidental slips have some power to exculpate, but there's no temptation to say that these must be justifications simply because they differ in salient ways from temporary insanity or infancy. While nothing of substance seems to turn on whether we lump all the excuses together or distinguish excuses from exemptions, there is a substantive question as to whether a defense that cites the subject's ignorance or mistaken beliefs justifies or merely excuses. I'll work with the trichotomous scheme and explain why the internalists haven't (yet) offered any good reason to think that the beliefs formed in the bad case are justified rather than excused.¹⁸

3.1 Epistemic Defects and Exculpation

If we're using the language of the trichotomous scheme, we wouldn't say that someone should be excused for having ϕ 'd unless we were prepared to defend the subject from criticism while making two concessions. The first is that there wasn't sufficient reason for them to have ϕ 'd. (Excuses aren't justifications.) The second is that the subject had the intact rational capacities necessary for being held accountable for having ϕ 'd. (Excuses aren't exemptions.) If an excuse involves both concessions, how do they work?

In general, we're blamed for manifesting *de re* unresponsiveness.¹⁹ Norms identify the relevant *res* and we manifest this *de re* unresponsiveness when our responses manifest the failure to respond to the reasons that determine what an appropriate response would be. What excuses are supposed to do is show that while an agent's response might well have involved a failure to respond to some reason, the subject shouldn't be blamed. It has to do this even though we're acknowledging that her rational capacities were intact and operative. This last part is crucial. The excuse highlights features of the case that show that this subject's response doesn't show an insufficient concern for the reasons that apply to her but it presupposes that her response was something she could be held accountable for. When factual ignorance and mistake do, they do this because someone who isn't aware of some reason wouldn't manifest insufficient concern for that kind of reason by failing to be guided by it. So, there's at least some reason to think that ignorance and mistake can excuse.

This only gets us part of the way. It explains why we should believe ignorance and mistake are potentially exculpatory, but it doesn't explain why we should think of defenses that cite ignorance or mistake as excusatory rather than justificatory. That's the difficult question. In this passage, Fantl and McGrath discuss Bernard Williams' petrol and tonic case and explain why they think that epistemic defects like ignorance and mistake figure in justificatory defenses rather than excusatory defenses:

¹⁸ Although ignorance and mistake can excuse, they don't invariably excuse. One reason that ignorance or mistake might not excuse is that the ignorance or mistake is not non-culpable. It should also be noted that normative mistakes might not have the same significance as factual mistakes. It's one thing to say that someone shouldn't be blamed for hitting someone with their car because they mistakenly believed that there wasn't a cyclist ahead and another to say that someone shouldn't be blamed because they mistakenly believed that cyclists don't matter. I'll have more to say on this below.

¹⁹ See Arpaly (2002) for discussion.

Notice that if we asked the unlucky fellow why he did such a thing, he might reply with indignation: 'Well, it was the perfectly rational thing to do; I had every reason to think the glass contained gin; why in the world should I think that someone would be going around putting petrol in cocktail glasses!?' Here the unlucky subject ... is not providing an excuse for his action or treating what he did as a reason; he is defending it as the action that made the most sense for him to do ... He is providing a justification, not an excuse (2009: 125).²⁰

Their argument seems to be something like this:

1. If the agent's action was the thing that made sense to do and was rational/reasonable, the agent's defense constitutes a justification, not merely an excuse.
2. The agent who acts in ignorance or on a mistaken belief often does the thing that made sense to do and was rational/reasonable.
- C. Epistemic defects like ignorance and mistake exculpate by subverting obligations to do otherwise and so justify rather than excuse.

Their example concerns the justification of action, not belief, but they seem to want to treat the cases in parallel. I'm happy to follow their lead. I'll offer two points in response.

The first is that the features Fantl and McGrath point to in their defense (i.e., that the agent's response 'made the most sense' and was 'rational') would have to be in place for an excuse to do its work in many of the standard excusatory defenses. This is important, so I'll say it again. An excuse wouldn't be appropriate if the agent's response didn't make sense and wasn't a rational thing to do in the situation.

Look at Austin's examples of excusing conditions:

- The agent was nudged.
- It was partly an accident or an unintentional slip.

²⁰ In their most recent discussion of pragmatic encroachment, Fantl and McGrath affirm something they call 'Clifford's Link', which says that if your belief that *p* is justified, *p* is warranted enough to justify actions suggested by *p* (2014: 87). The passage above suggests that they think that facts about the justification of belief determine the justification of our actions, but in other places they seem to back off this and say that all the justification of belief will give you is something weaker (e.g., that the premise is justified enough to use in reasoning that leads to action but doesn't justify the actions that bring that reasoning to its conclusion). It seems to me that this weaker claim is incompatible with the intuitions they offer to support Clifford's Link (e.g., their remarks at 2014: 92). In Littlejohn (2012) and (2014), I've argued that Clifford's Link causes trouble for the sort of internalist approach to justification they seem to endorse. (If you combine any of the standard internalist views of epistemic justification with Clifford's Link, you'll face skeptical worries and the view generates all sorts of morally problematic consequences.)

- It wasn't fair to simply say that he was ϕ -ing because he was doing something else and his ϕ -ing was merely incidental.
- He was looking at the matter quite differently and didn't realize that he was ϕ -ing.

If Austin sorted his cases into the right categories, it's hard to see how these conditions could excuse if the agent's response wasn't rational or one that 'made the most sense'. Suppose that someone in the agent's position could have expected to be nudged if they tried things one way and could have avoided it by trying it a different way. The nudge wouldn't excuse, would it? Or perhaps there was no way to avoid a potential nudge but the agent didn't take as much care as we'd normally expect to avoid it. A nudge wouldn't excuse here, but would if the agent took reasonable steps to avoid such interference. Or, suppose that someone in the agent's position could have expected that something would make her slip but took no precaution and plunged ahead anyway. A slip would be an indication of carelessness and would inculcate rather than excuse. The slip excuses only when we're operating under the assumption that care was taken and that the agent slipped in spite of having taken *reasonable* care. It's striking that the last two examples that Austin gives are epistemic in character. Someone might pump the water and poison the inhabitants without realizing what they were doing. Does their ignorance get them off the hook? It depends. If they had no knowledge that the water was poisoned, they might get off the hook. If they weren't the same 'on the inside' as someone who acted rightly in pumping the water, though, it's hard to see how doing something else or seeing things differently could get the agent off of the hook.²¹

Remember that when we're offering an excuse in the trichotomous scheme, we're assuming that the agent's rational capacities are intact and operative so that the agent can be held accountable for what she does in general. If someone's rational capacities are intact and operative, it's completely obscure why we'd let them off of the hook if we didn't judge that they used those faculties excellently. What could their excuse be if they didn't?! *If*, however, they use their rational capacities excellently and show all the care and concern that could be expected of someone in their situations, an

²¹ Some writers say that if the subject pumping the water reasonably believed that they were doing the right thing, this might get them off the hook. It might, but Gerken's (2011) point that beliefs about the normative status of a belief or action cannot be required for an excuse is well taken. Still, it's plausible that if you're the same on the inside as someone who knows that it's right to ϕ , that's a good sign that your ϕ -ing would either be excused or justified. Being the same on the inside as someone who knowingly does what's right might be *sufficient* for justification or for excuse, but it's not necessary. I think the crucial test is whether the agent is the same on the inside as someone who conforms to the norms that apply to her. To justifiably act or believe, you needn't have any beliefs at all about the justificatory status of your acts or beliefs. If you're the same on the inside as someone who acts or believes with justification, I'd suggest that your acts or beliefs are likely to be either justified or excused. I have some reservations about Boyd's (forthcoming) separabilism because it links a kind of epistemic assessment to facts about the agent's epistemic position with respect to the norms that apply to her.

excuse might be in order. If they did exercise those capacities excellently, however, they'd have to do things that could be reasonably taken to be the right thing to do in the circumstance. Anyone who buys into the trichotomous scheme has to think that the marks that Fantl and McGrath fix on are the marks of excuse.

In short, here's my argument against their take on the case:

1. Because we're offering an excuse rather than an exemption, we have to show that the agent cannot be blamed for ϕ -ing even though she has the intact rational capacities required for assuming responsibility for having ϕ -d.

2. If the agent has the intact rational capacities required for assuming responsibility for having ϕ -d she deserves an excuse only if she shows excellence in the exercise of these rational capacities (i.e., by doing things that could reasonably be taken to be right, by doing the thing that it made sense for a conscientious person to do, etc.).

C. Thus, the agent can be excused for having ϕ -d only if she shows excellence in the exercise of her rational capacities (i.e., by doing things that could reasonably be taken to be right, by doing the thing that it made sense for a conscientious person to do, etc.).

Now, if you think that non-factive mental duplicates whose mental states stand in the same causal relations exercise their rational capacities in the same way and so exercise them equally well, we can go in for a kind of internalism about excusability. We also might be able to go in for a kind of internalism about justification, but *that* depends upon some substantive matters about the kinds of norms that govern our actions and attitudes (e.g., whether these norms all have internal application conditions so that mental duplicates invariably do equally well in conforming to them).²²

²² In the literature, there are a number of strategies for arguing that reasonable but mistaken factual beliefs and non-culpable ignorance excuse actions that violate norms rather than justify them, subvert obligation, or show that the putative norms aren't genuine norms. The subject who poisons others acting in the mistaken belief that they are pumping potable water might have reparative duties, duties to right past wrongs. See Littlejohn (2012) for discussion of this point. Subjects who use force against an apparent aggressor will act with justification if the apparent aggressor is an aggressor and can only be excused if the apparent aggressor is merely apparent. As evidence for this, people often appeal to intuitions about the justified use of force to intervene and the agent's rights of non-interference. See Duff, Ferzan, Gardner, Moore, and Robinson on this point. Darley and Robinson (1996) have also done empirical work into folk intuition and found that folk intuitions support a deeds account of justification according to which the availability of a justificatory defense depends upon the objective nature of the agent's deeds, not just the motivations of the agent at the time of action. Such work is significant in two ways. First, it suggests that intuitions like those that Cohen elicits don't have the significance he takes them to when we're considering the justification of action. Second, it suggests that insofar as there's a principled link between the justificatory status of actions and beliefs that guide these actions we

I've suggested that maintaining the trichotomous scheme and maintaining the distinction between exemptions and excuses forces us to recognize that excusing conditions really only excuse when they are present alongside a display of virtue. An excuse affirms the excellent use of the subject's capacities in spite of the wrong and shows that the wrong isn't itself any indication that the subject is anything less than fully virtuous.²³ If the agent's handling of the situation doesn't manifest the excellent use of her rational capacities but those capacities were intact and operative, it would be totally obscure why these so-called excusing conditions would exculpate. Thus, rationality and reasonableness are the mark of the excusable, not the justified. This, I'd argue, is to be expected because the crucial test to determine whether a subject's ϕ -ing was justified is whether it conformed to the norms governing it. This depends upon what reasons applied to the subject and whether she managed to meet their demands. If we knew in advance that someone would display this sort of excellence in how they used their rational capacities, we wouldn't necessarily know in advance whether their actions and attitudes would be justified.²⁴ I think the answer to that is obvious. We couldn't.

3.3 Rationality as the Mark of Justification

When excuses are distinguished from exemptions, an excuse typically works by affirming the excellence of the subject's exercise of her rational capacities. A good test to determine whether the subject has exercised these capacities well is by asking whether the subject could be the same on the inside as someone who acts rightly or believes with justification. If it turns out that pairs of internal duplicates don't do equally well in conforming to the norms that apply to them, then it would turn out that their actions or attitudes wouldn't be equally justified. However, while their actions and attitudes wouldn't be justified we would still be disposed to think that there's something good to be said about them in light of how they responded to the situation. Even if the subject doesn't conform to the norms governing her conduct, we're still often disposed to evaluate her and her response positively in a way we wouldn't if we

should expect ignorance and mistake to exculpate by excusing rather than justifying in both cases. I don't see any evidence that the epistemologists who take the opposing view about how ignorance and mistake exculpate have considered any of the arguments in this literature. They have missed some very valuable discussions that address points that have clear parallels to the points at issue in the discussion of the new evil demon case.

²³ Mind you, you needn't *be* virtuous to be excused. One merely needs to have responded as a virtuous person would have done.

²⁴ Thus, it's useful to distinguish, as writers like Broome (2013) and Weatherston (2008) have, between identifying reasons and processing them correctly and processing apparent reasons correctly. Justification requires the identification and the proper processing, but rationality is really just a matter of the proper processing. The failure to respond to actual reasons isn't itself a sign that there's some sort of malfunction in the subject's rational capacities. As Graham (2012) notes, something that doesn't fulfill its function needn't be thought of as malfunctioning (e.g., a properly functioning toaster's coils might heat up as they should and for as long as they should even if you haven't put in the proper toastables).

were dealing with an exemption, because we see the failure as occurring in spite of the excellent way she's exercised her capacities.

To determine whether justification is an internalist notion or an externalist one is a *substantive* issue, one that really isn't settled by the fact that internal duplicates will show the same excellence in their use of their rational capacities. It *could* turn out that, say, all moral norms are concerned with internal conditions in which case justified action would turn out to be an internalist notion. It *could* turn out that, say, all epistemic norms are concerned with internal conditions like consistency and coherence. On the other hand, it *could* be that rightness and justification depends upon things like the consequences, the objective nature of the deeds we perform, or the accuracy of our attitudes. Such questions *should* be settled on *substantive* grounds, however, not simply by noting that internal duplicates will display the same excellence in the use of rational capacities.

Or so I say. Should you agree? I think that this is required for maintaining the trichotomous scheme, so if you think that there's something to this scheme, maybe you're already on board. If you're not already on board and you want some further argument, I shall argue that the excellent use of rational capacities is not a guarantee that a subject will meet the standards that apply to her. It doesn't really matter whether the standards that apply to her have internal application conditions or external ones, the subject's excellent use of her rational capacities simply cannot guarantee norm conformity. In the cases where the capacities are used excellently, it would be perverse of us to blame, so we need to offer some sort of defense. Because a norm we recognize as a norm has been violated, we'd have to be muddled to offer a justification. It seems we'll need some non-justificatory defense that affirms the subject's rational capacities and that just is an excuse. This is why I think the new evil demon argument is no good. It is no good because it tries to settle a substantive question about what's required for justification without invoking any substantive claims.

Maybe *this* is where the internalists want to register their disagreement. They might agree that *if* we wish to maintain the trichotomous scheme we would have to think of excuses as involving a kind of excellent use of a subject's rational capacities but they might insist that this just shows that the trichotomous scheme is mistaken.²⁵ Insisting that the reasonable and the rational is the mark of justification, they might insist that the way that excuses work is the way that I've claimed exemptions work, which is to say that they do their work by showing that a subject cannot be held accountable for her actions or attitudes generally because she doesn't have intact, functioning, developed rational capacities.

This might be the kind of response that Cohen would offer.²⁶ In his initial discussion, he dismisses the idea that we might distinguish rationality from

²⁵ Steup is one of the few authors explicit about this. He insists that incapacitation is the *only* thing that can epistemically excuse (1988: 78).

²⁶ Huemer's (2006) case for internalism rests on a similar point. As he sees it, it's important that normative differences are rationally recognizable. He criticizes reliabilists for saying that there can be pairs of cases, C1 and C2, where there's justification for believing that things are as they appear in C1 but not C2 where this difference is not rationally recognizable. In Littlejohn (2011), I tried to address his argument.

justification, saying, "'reasonable' and 'rational' are virtual synonyms for 'justified'" (1984: 283). In his most recent discussion, he reiterates this point, saying:

... distinguishing between rational beliefs and justified beliefs does not make any sense. It makes no more sense than distinguishing between moral acts and justified acts. Justification is domain relative. A justification for ϕ -ing can be moral, rational, legal, etc. It is moral to ϕ just in case ϕ -ing is morally justified. In the same way, it is rational to F just in case one's ϕ -ing is rationally justified. So it is rational to believe just in case believing is rationally justified. When one distinguishes between rational belief and justified belief, one can only be distinguishing between rationally justified belief, and some other kind of justified belief. What is the other kind of justification (MS: 5)?

Cohen's remarks suggest that he thinks that normative notions as such are connected to things like whether the subject can be held responsible. As he puts it, his objection to externalism turns on viewing justification as a 'normative notion' and thus one that doesn't depend upon any circumstances beyond the subject's ken (1984: 282). I think he *has* to take this line if he's going to identify rationality with justification within a domain. If he didn't, then he'd have to concede (as I think he should) that the distinction between rationality and justification makes perfectly good sense because we could say that justification can depend upon factors beyond those that bear on justification by depending (in part) upon matters beyond the subject's ken. To determine whether S's ϕ -ing was justified, we have to ask whether S conformed to the norms that applied to her in ϕ -ing. To determine whether S's ϕ -ing was rational, we have to ask whether S's ϕ -ing results from responding excellently to the apparent reasons. When rational subjects are guilty of failures of rationality, they are guilty of some failure in the way they reason. When rational subjects are guilty of failures of justification, they are guilty of some failure to conform to a norm that applies to them. These are clearly different heuristics and if it happens to be the case that all and only the rational beliefs are justifiably held that has to be because this non-trivial claim is true: excellence in rationality guarantees norm conformity. If this non-trivial claim is false, we have more than a distinction in heuristics. If it's false, rational beliefs needn't be justified and internalism about the former doesn't support internalism about the latter.

We can use cases of factual mistake to cause trouble for R=J:

1. N is a norm that says that you must X if C obtains. As it happens, you must not X if C doesn't obtain [Assume].
2. C obtains [Assume].
3. A rationally believes C doesn't obtain [Assume].
4. It is rational for A to believe she must not X. [Assume A rationally believes N is a norm and rationally believes that C doesn't obtain].
5. If it is rational for A to believe she must not X, A is justified in refraining from X-ing [On the assumption that non-culpable factual mistakes justify].

6. A is justified in refraining from X-ing [From (4) and (5) because we're assuming that mistake/ignorance justifies].
7. A must X [From (1), (2)].
8. A is not justified in refraining from X-ing [(7)].

Since refraining cannot be justified if it is refraining from what you must do, (6) and (8) cannot both be correct. If C is the application condition for a norm, it's being an application condition for a norm doesn't depend upon the attitude we take towards its being an application condition for a norm, so we wouldn't want to say that *knowing* or *rationaly believing* C to be an application condition for a norm precludes its being one. So, it seems that there are two ways to block the argument for (6) and (8). The first is to reject R=J. That's what I'd recommend. Once we do that, we can see that justification requires more than rationality. Being rational is exercising your rational capacities as you would in responding correctly to reasons but being justified is doing that and responding correctly to reasons. Rationality would provide a kind of excuse, but nothing more.

If you wanted instead to retain R=J, you'd have to accept this access thesis, an epistemic constraint on genuine norms:

The Access Thesis: If C is the application condition for a genuine norm, it's not possible to have a rational but mistaken belief about C's obtaining.

The trouble with the Access Thesis is that it seems that for any condition (internal or external) it should be possible to have arbitrarily strong evidence that C obtains even if C fails to obtain. Thus, if strong evidential support is sufficient for rational belief, it seems that the Access Thesis leads to a *highly* revisionary account of norms. Not only couldn't there be norms that, say, required you to keep promises or try to keep promises made a nanosecond earlier, there couldn't be norms that, say, required you to respond to your own pains or beliefs, not if it's possible to have strong evidence that you believe things you don't actually believe or are in pain when you aren't in pain.²⁷

²⁷ An example of Arthur Prior's shows that it's possible for two subjects perfectly alike physically and very similar mentally who differ in terms of whether they believe *p* because of contingent facts about their situation. Here's the case. A monk enters a room with four other monks. Not thinking that much good will come from the thinking in this room, he mutters to himself, 'Most of the things thought in this room will be false'. As it happens, the four other monks each think precisely one thing. In the first version of the case, all four monks think true things, in which case the muttering monk believes something false. Let's call this proposition, the proposition that the monk believes, *m*. In the second version of the case, two monks think true propositions and two monks think false propositions. In this second version of the case, the muttering monk *cannot* believe *m*. In this case, the muttering monk neither believes nor asserts anything at all. He could not assert or believe anything at all. A propositional attitude like belief has to take a proposition as its object and all such propositions must either be true or false. This is what bivalence tells us. There *is* no proposition that the monk could believe in the second version of the case that is true or that is false, so the monk doesn't have the propositional attitude he does in the first case in the second case. (Thanks to John Hawthorne for introducing me to the case. He discusses them at greater length in his manuscript on the surprise exam paradox.) Now, it is consistent with all of this that the monk entered the room with arbitrarily

The Access Thesis is deeply implausible. The best way to block the derivation of (8) is to reject R=J. Without too much muss and fuss, we can generate a similar argument against R=J using factual ignorance instead of mistake. This argument would force defenders of R=J into accepting further implausible access theses.

I said earlier that the intuitions that underwrite Cohen's objection are not the sort of thing that we should rely on to determine what sort of norms govern our actions or attitudes. In light of the argument just sketched, I can now explain why. If the Access Thesis is false, then for just about any genuine norm we can think of, N, say, it should be possible to imagine a situation in which (a) the subject knows that N is a genuine norm, one that says that she's required to ϕ if C obtains, (b) she's actually required not to ϕ , and (c) she mistakenly believes that C obtains. If she ϕ 's in the belief that C does obtain, we will respond positively to her ϕ -ing. We will see her ϕ -ing as a manifestation of her competence at handling reasons. While we will have the same kind of positive reaction to her ϕ -ing that we have when considering Cohen's subjects in the bad case, the ϕ -ing is not justified. That has been stipulated at the outset. It's done in violation of N. Thus, the positive response cannot be understood as the recognition that the subject conforms to the norms that determine whether it's appropriate for her to ϕ . Thus, the positive response cannot be understood as the recognition that the subject's ϕ -ing was justified. That's been stipulated from the outset.

While I think that the right response here is to reject R=J and the Access Thesis, some epistemologists will want to hang tough. They face a related worry about how *normative* mistakes will generate obligations if justification amounts to nothing more than rationality (understood as belief based on strong evidence). We start by imagining cases with similar structure as the problematic cases of factual error or ignorance:

strong evidence for believing that 'Most of the things thought in this room will be false' would express a true proposition if sincerely uttered. You can add in your auditory hallucinations if you like so that the monk in the first case *knows* that *m* is true and the second monk is, apart from having an attitude towards *m*, as close to the monk in the first case as any monk could be without being a complete mental duplicate. At this point, someone could say that this just goes to show that the norms that govern belief will have application conditions that refer to internal mental states that aren't belief. Might these norms' application conditions refer to seemings instead of beliefs? They might, but that surely won't make the problem go away. You can generate the very same problem with seemings if you introduce some monks that have seemings about the number of false seemings formed in a room, by a group, etc. One of the monks might have arbitrarily strong evidence for believing that it seems to him that *p* when it doesn't (for just the sorts of reasons discussed), in which case you could get a case with the requisite structure (i.e., a norm that has an application condition that refers to the subject's seemings, evidence that misleadingly suggests that things seem to the subject a certain way, etc.). You can't go internal enough to avoid the structural problem.

1. A rationally believes that N is a norm that says that one must X if C obtains. A also rationally believes that A must not X otherwise.²⁸
2. A knows C obtains [Assume].
3. N isn't a genuine norm. In fact, there's a genuine norm, N', that says that one must not X if C obtains [Assume].
4. A is rationally required to believe she must X. [Plausible in light of (1)-(3).]
5. If A is rationally required to believe she must X and rationality just is justification, A is not justified in refraining from X-ing.
6. A is not justified in refraining from X-ing (i.e., A must X).
7. A must not X.

Since (6) and (7) could not both be correct, we have to block the reasoning at some stage. A defense of the Access Thesis won't help because we can generate the problem even if the subject has full information. We could introduce a different sort of access thesis:

The Second Access Thesis: It is not possible to have the rational but mistaken belief that N is a genuine norm.

You can have arbitrarily strong evidence for believing that N is a genuine norm even if it's not one. It seems that on every standard view of rationality it should be possible to rationally believe N is a norm when you have sufficiently strong evidence in support of that belief. If the evidence really is misleading, shouldn't the putative norm be spurious? If so, mustn't the Second Access thesis be false?

If proponents of R=J want to stick to their guns, they'd have to deny (7). The evidence that I'd describe as misleading isn't *really* misleading on this view if the evidence is sufficiently strong.²⁹ Because of this evidence, they'd have to say that there's a norm that applies to A that requires A to X. One of the strange features of this view is that it seems that if a subject has strong evidential support for believing that N is a genuine norm, this kind of reasoning wouldn't just lead to the correct conclusion that N is a norm, it would seem to generate *knowledge* of its conclusion:

There is a norm, N, that says that one must X if C obtains. C obtains. So, I must X.

Suppose comes to believe this on strong evidence and, in keeping with R=J, this norm is (now) genuine. Suppose A doesn't X in line with her belief. Suppose that she's akratic. If a subject doesn't X when she knows that she must X (and this isn't a case of duress), it's tempting to say that the subject should be blamed for X-ing, but it seems completely perverse to blame someone for performing actions that don't manifest *de re* unresponsiveness. Remember that by *our* lights, A wouldn't manifest *de re* unresponsiveness by refraining from X-ing. By our lights, this is a case of inverse akrasia, a case in which a subject does right by refraining from acting on her own

²⁸ We're assuming that this requires nothing more than believing these propositions on the basis of evidence that provides sufficiently strong support.

²⁹ I suppose that if the evidence were weak, it would be misleading. As the evidence gained in strength, however, it would cease to be misleading, as it would help to determine which normative propositions were true.

judgment about what to do.³⁰ Our subject could have been Huck Finn, say. If, like me, you wouldn't blame Huck for failing to turn in a runaway slave, you'd probably want to resist R=J because we can modify the story very slightly by insisting that all of Huck's evidence indicated that runaway slaves ought to be turned over to the authorities. We could build in features to the story that make it rational (according to standard theories of rationality) for Huck to believe that he's required to turn Jim in, but you're a despicable person if you think that it follows from this that Huck should be blamed for helping Jim escape or that what Huck ought to have done was turn Jim over to the authorities. Denying R=J lets you retain your orthodox conception of rational belief without turning into a baddy yourself because it lets you say that Huck's actions were right even though they didn't mesh with his evidentially well-supported beliefs.

The general point to take from this is that you don't want to build into your conception of a norm some epistemic requirements that say that it's not possible to rationally believe the norm to apply when it doesn't, to rationally fail to see that it applies when it does, or to have rational but mistaken beliefs about which norms are genuine. Once we see that, we can see that there's a potential gap between rationally believing that something is true or that something is required and the belief or action in question conforming to the relevant norms. If a rationale is needed for distinguishing justification from rationality, this is precisely the rationale. Something is justified when there's sufficient reason for it and we determine whether there's sufficient reason for something by considering whether it conforms to the relevant norms. The general philosophical lesson to draw from this is that the excellent use of our rational capacities does not ensure that we'll do what the reasons demand from us. You can make this point by pointing to examples in which the relevant reasons are features of the situation that the subject can be ignorant of or mistaken about, but you don't need to think of reasons that way. Even if you think reasons are just in the head, that doesn't take care of the problem. So, the general point is one that every internalist and externalist should appreciate. In turn, every internalist and externalist should appreciate that the positive response elicited when we judge that someone rationally X'd is not the recognition that the subject's X-ing was justified. Whatever the norms are, we can imagine someone responding to misleading evidence and will respond positively to their X-ing even if done in violation of a norm.

This suggests, in turn, that rationality is quite often a sign of *excuse*, not *justification*. It is a sign of excuse because an excuse would often be out of order if the agent's actions and attitudes weren't rational responses to the apparent reasons, responses that showed the subject to be an excellent *processor* of reasons, albeit an imperfect *detector* of them.³¹ If a subject's rational capacities are intact, we couldn't

³⁰ The expression 'inverse akrasia' is due to Arpaly (2002). For further discussion of this line of argument, directed at views like Zimmerman's (2008) on which a subject's normative evidence can generate norms, see Littlejohn (2014). Like Harman (2011), I'm quite skeptical of the claim that normative ignorance/mistake has the same kind of exculpatory significance as factual ignorance/mistake, but I'm particularly put off by the idea that a subject's normative ignorance/mistake could generate obligations and inculcate in cases where outside parties can clearly see that there's nothing that calls for the subjects to act in line with their mistaken moral judgments.

³¹ An imperfect detector of reasons might not manifest any sort of worrisome *de re* unresponsiveness, but that's what we'd have with an imperfect processor of (possessed)

defend the subject's failure to conform to a norm without showing that the subject's actions or attitudes manifested the care, consideration, and careful reasoning that typically results in right action or knowledge.

4. The Reasons Rationale(s)

This should be enough to show that the intuitions that underwrite the new evil demon objection don't have any evidential significance. We are disposed to respond positively when someone X's if they're moved by misleading evidence that suggests that X-ing has the properties that would mean that it conformed to the relevant norms. The positive response cannot, then, be understood as the recognition that the X-ing was justified. So, the intuitive positive response that underwrites the new evil demon objection really doesn't tell us anything interesting about the justificatory status of the beliefs formed in the bad case.

Cohen (1984) offers further arguments against externalism. It seems that they're supposed to stand on their own feet. This further argument has to do with reasons. In explaining why we should accept Equality, he remarks:

Now part of what the hypothesis entails is that our experience is just as it would be if our cognitive processes were reliable. Thus, on the demon hypothesis, we would have every reason for holding our beliefs, that we have in the actual world. Moreover, since we actually have reason to believe that our cognitive processes are reliable, it follows that in the demon world we would have every reason to believe that our cognitive processes were in fact reliable ... It strikes me as clearly false to deny that under these circumstances our beliefs could be justified. If we have every reason to believe e.g., perception, is a reliable process, the mere fact that unbeknown to us is not reliable should not affect its justification conferring status (1984: 282).

The argument seems to be something like this:

The Reasons Rationale

1. Subjects in the good case and the correlative bad cases have the same reasons.
2. If subjects have the very same reasons, they will have the same justification to believe.
- C. Subjects in the good case and the correlative bad cases have the same justification to believe.

It seems that one of the central claims of the reasons rationale can be taken in one of two ways. The claim that pairs of subjects have the same reasons might be understood as the claim that they possess the same evidence or it might be taken to mean that the same reasons apply to them. Consider two claims:

reasons. For arguments that rationality cannot consist of responding to reasons on the grounds that the failure to identify the things that constitute reasons needn't indicate any sort of irrationality, see Broome (2013).

Same Reasons Apply: If two subjects are internally the same, the same reasons bear on whether they should believe p .

Same Reasons Possessed: If two subjects are internally the same, the subjects have the very same reasons/evidence for their beliefs.

These claims are potentially distinct. It's not obvious that the set of reasons that bear on whether to believe p is simply the set of evidential reasons that support a subject's beliefs. For some propositions it seems that you shouldn't believe p if you don't have sufficient evidence for that belief. If you have *some* evidence for p but the evidence doesn't provide sufficiently strong support, you shouldn't believe p , so there's a reason that bears on whether to believe p where *this* reason isn't itself a piece of evidence. For this reason, we should distinguish Same Reasons Apply from Same Reasons Possessed. Thus, there are *two* reasons rationales to consider, not just one.

Let's look at the argument from Same Reasons Possessed:

The Possessed Reasons Rationale:

I. Subjects in the good case and the correlative bad cases have the very same reasons/evidence.³²

³² Cohen (1984), Dougherty and Rysiew (2013), Gibbons (2010), and Turri (2009, 2013) all defend views on which this should come out as true because possessed reasons are identified with a subject's states of mind (e.g., experiences, beliefs, intuitions, apparent memories). Fantl and McGrath (2009) defend a view on which this should come out as true because possessed reasons are identified with the propositions a subject has in mind, including falsehoods. To my mind, the linguistic evidence cited by Hornsby and Unger tells strongly against all these views. Agnes' reason for being upset isn't her experience. If Agnes is upset because the neighbors are making so much racket, her reason for being upset is that the neighbors are making so much racket. That is to say, her reason for being upset is the fact that the neighbors are making so much racket. Reasons of all kind are facts and a subject's reasons are facts known to her. The mistake that people make in accepting the first premise of the argument is probably due to the mistaken assumption that the subject in the bad case believes things on the basis of reasons. There's nothing that could be the deceived subject's reasons for ϕ -ing. If the deceived subject is angry because it sounds as if the neighbors upstairs are making too much noise, we wouldn't say (a) her reason for being angry is that the neighbors are making too much noise and wouldn't say (b) her reason for being angry is that it seems to her that they are making too much noise. The trouble with (a) is that propositionally specified reasons have to be true. The trouble with (b) is that Agnes isn't crazy. She doesn't take the fact that it sounds to her as if someone is making noise to be something that merits an angry response on her part directed towards the neighbors. For a careful discussion of reasons and error cases, see Alvarez (2010). For an argument that the subject couldn't ϕ for the reasons that she'd ϕ for in the good case, see also Hornsby (2007), Hyman (1999), Littlejohn (2012), and Unger (1975). Whereas Alvarez, Hyman, myself, and Unger would probably all agree that the subject in the bad case doesn't ϕ for reasons at all, Hornsby offers a disjunctive account of ϕ -ing for reasons.

2. If two subjects have the very same reasons/evidence, they have the same justification for ϕ -ing.
- C. Subjects in the good case and the correlative bad cases have the same justification for ϕ -ing.

We should reject both premises of this argument.

If we think of possessed reasons as something like the subject's reasons, it's clear that (1) wouldn't be true if taken as a claim about a subject's reasons for doing or feeling something. Your reason for being angry with your neighbors is something that you desire not to be the case. Your reason for being angry might be that they are drilling so early in the morning. Your reason for being angry with them is something you know and something that's the case, so it's *not* something that could be a deceived subject's reason for being angry. I don't know how weird you are, but I'd guess that you aren't *so* weird that you're the kind of person who'd be angry that it seems the neighbors are drilling so early, that it sounds as if they're drilling, or that you believe that they're drilling. Since your reason for acting in anger will often be your reason for being angry, it's possible for you to act for reasons that couldn't be a subject's reason for acting in the bad case. Since your reason for acting in anger will often be your reason for believing that your actions were appropriate to the situation, your reason for believing things will often be things that couldn't be a subject's reason for believing something in the bad case.

People often miss this because they conflate motivating reasons with explanatory reasons (i.e., a subject's reasons for ϕ -ing vs. reasons why a subject ϕ 'd). If you think I'm being uncharitable, you can check for yourself to see if the people who think of a subject's reasons or evidence as something that she shares in common with her systematically deceived counterparts recognize the distinction and apply it correctly. The extant arguments that purport to show that a subject's reasons for believing things is limited to the reasons that could have been your reasons in the bad case would show that your reasons for feeling things or doing things would be similarly constrained. Since *they* aren't constrained, it's not plausible that your reasons for believing things is so constrained since your reasons for doing things and feeling things are often your reasons for believing things.

A further worry about Same Reasons Possessed is that it leads to skepticism. It's obvious that it does if E=K is correct.³³ If your evidence consists of all and only what you know, you wouldn't know more than your deceived counterpart if your evidence and their evidence were the same. We can generate similar skeptical consequences from Same Reasons Possessed if we weaken E=K. If propositional evidence has to be true and your evidence will include anything you know non-inferentially, it follows from the fact (and it is a fact) that we have non-inferential perceptual knowledge that our evidence in the good case will be better than the evidence we'd have in the bad case. This is a trivial consequence of being in the good case.

Let's consider the argument's second premise, the proposition that the same things will be justified for two subjects if they have the same evidence or possessed reasons. This is also mistaken. The facts that determine whether a subject's response to her situation is appropriate, fitting, right, permitted, or justified *doesn't* supervene upon facts about that subject's evidence. If you're angry with your neighbor because

³³ See Hyman (2006) and Williamson (2000).

you believe that they're drilling so early, this anger isn't fitting or appropriate if they aren't drilling. If you lash out in anger at your neighbor because you believe they're drilling so early, this action isn't appropriate, permitted, or justified if they aren't the ones who are drilling. This is why apologies and reparations are called for when you act in ways that harm someone (e.g., screaming at them, banging on the ceiling, filing a complaint about them, etc.) when you act on a mistaken belief. Similarly, if some subject takes steps to defend themselves in the belief that some person is an aggressor, we can't say whether this person's use of force is justified until we know whether this person actually *is* an aggressor. The person who tries to use mace (or something worse) to fend off a perceived aggression doesn't have the same rights to non-interference or assistance that they might otherwise have if they act in the mistaken belief that they're facing an aggressor. If they are facing an aggressor, we might be duty bound to assist them and duty bound to refrain from interfering with them. If they aren't, we might be duty bound to interfere with their action and to refrain from assisting them. If these points are correct when it comes to emotion and action, similar points should hold for belief. The justification of believing that you ought to ϕ depends upon whether it would be right to ϕ .

Once we see that both premises of the Possessed Reasons Rationale are mistaken, we'll see that this second version of the reasons rationale is also unsound:

The Applicable Reasons Rationale:

1. If two subjects are internally the same, the same reasons bear on whether they should ϕ .³⁴
2. If the same reasons bear on whether they should ϕ , they have the same justification for ϕ -ing.
- C. Subjects in the good case and the correlative bad cases have the same justification for ϕ -ing.

The argument's second premise is fine, but the argument's first is mistaken. Justified emotional response and action typically require a proper fit between appearance and reality. If there's a link that connects the justificatory status of a belief to the emotions and actions that the belief rationalizes, externalism about the justification of emotion and action supports externalism about the justification of belief.

5. The Newest Evil Demon Problem

I've argued that none of the considerations Cohen offers us give us any reason to reject externalist accounts of justification. My aim has been to show that externalists can draw distinctions that we have good independent reason to draw and that once we draw these distinctions externalists can accommodate some of the intuitions that he appeals to and can rightly reject some of the important assumptions that feed into his attack on externalism. In this concluding section, I shall gesture at a problem for internalist views motivated by Cohen's argument.

³⁴ Authors like Cohen (1984), Gibbons (2010), and Huemer (2006) might argue that this premise is supported by the idea that the relevant sort of internal sameness shows that the same reasons apply to these subjects because it shows that the same things would be rational for them. Since I don't think rationality is the mark of justification, I don't think we need to rehash this again.

It's important to remember that while there are intuitions that seem to cause trouble for externalists, there are also considerations that point in a different direction:

The Newest Evil Demon Argument

P1. In the bad case, the subject's beliefs don't have the properties that figure in externalist accounts of justification [Observation].

P2. The beliefs are equally justified in the good case and the bad case [Equality].

P3. For some beliefs, your beliefs wouldn't be justified unless you were in a better epistemic position than a mental duplicate of yours in a bad case [External Constraint].

C. The relevant class of beliefs couldn't be justified however strong your epistemic position might appear.

This objection is directed against *every* non-skeptical internalist view supported by the New Evil Demon Argument. To avoid the skeptical conclusion, we'll either have to reject Equality or External Constraint. Since Equality is the crucial assumption in the new evil demon argument, I don't think that Cohen would want to reject that. Thus, he'll have to reject External Constraint.

I've argued for External Constraint elsewhere on intuitive and theoretical grounds. If External Constraint is correct, externalism is the only coherent alternative to skepticism. If External Constraint is correct, there's no sound argument from Sufficiency and Equality to some non-skeptical internalist view because Sufficiency and Equality lead to skepticism.

To get a feel for how the Newest Evil Demon Argument is supposed to work, consider the role that beliefs play in rationalizing feelings, beliefs, and actions. Consider this quartet of claims:

1. The Enkratic Requirement: Rationality requires you to be enkratic (i.e., to avoid failing to ϕ when believing that you're required to ϕ).
2. $R = J$: Rationality is the mark of justification.
3. $J \rightarrow NC$: Justification requires norm conformity.
4. $E \rightarrow R$: If you have sufficiently strong evidence for p and believe on the basis of it, you rationally believe p .

Once you have these assumptions in place, consider cases in which you have strong evidence for the following propositions:

- (a) I ought to be angry with Bob.
- (b) Bob threatened someone I love.
- (c) I ought to punch Bob in the nose.
- (d) I ought to believe that Bob threatened someone I love.

In having strong evidence for such propositions, (4) tells us that you can rationally believe them. From this and (2), it follows that you can justifiably believe them. If you justifiably believe that you ought to be angry with Bob but aren't angry with him or you justifiably believe that Bob threatened someone you love, it wouldn't be rational for you to maintain this belief and fail to be angry with him. If so, it seems that it would be rational for you to be angry with him. If it is, this anger is justified. If it is, there's no norm that governs your emotional responses that you'd violate by being angry with him. And yet, there could be, couldn't there? Isn't this a plausible norm for anger?

FA: You shouldn't be angry with someone unless that anger is fitting.

The anger wouldn't be fitting if, say, your evidence is misleading. Maybe it wasn't Bob who threatened someone you loved. If it wasn't Bob, the anger wouldn't be fitting. Since FA seems like a pretty good norm, I'd think that there's a perfectly good sense in which you shouldn't be angry with Bob. If so, (3) tells us that the anger isn't justified. At best, it's excused. If so, (2) tells us that it's not rational. Now, this is where I'd get off, but Cohen assures us that the distinction between the anger it would be rational to feel and justified anger doesn't make sense. If it's not rational, what is it? Are you supposed to be Enkratic? Violating the Enkratic requirement is the clearest possible case of irrationality! If R=J is right and you're rationally required to be Enkratic, you cannot justifiably have the attitudes that make you akratic. So, are we supposed to say that the arbitrarily strong evidence that supports the belief that Bob threatened someone you love doesn't make your belief rational?

I don't see a good way forward here apart from (a) admitting that rationality and justification fracture in just the way I've suggested and (b) admitting that the justificatory status of your beliefs depends upon whether they'd justify the affective responses they rationalize. I don't see any good way to make sense of (b) without embracing a factive account of justification, so that's how I'd sort this mess out.

We get similar problems if we think about (c) and (d). The evidence, (t) tells us, makes it rational to believe (c) and (d). This, in turn, tells us that these beliefs are justified. This, in turn, tells us that no norm is violated in forming such beliefs. This, in turn, tells us that we ought to believe that Bob threatened someone I love. This in turn tells us that you ought to punch Bob. Again, the problem here is that the evidence can be misleading. Because it's misleading, we don't want to detach the conclusion that you ought to hit someone and we don't want to detach the conclusion that you ought to believe that Bob threatened someone you love. To my mind, the obvious thing to do is distinguish rationality from justification and let justification depend upon the justificatory status of the stuff that the belief rationalizes.³⁵

6. Conclusion

³⁵ Interestingly, Cohen's response to this sort of argument is to insist that R=J is correct and simply say that it's not *rational* to do things like punch someone you rationally believe threatened someone you love. His example is slightly different, but he thinks that the rationality of *action* depends upon facts that you couldn't know. This is like saying that it's not rational to put everything on red and spin the wheel if the ball will come up red. Since he's an internalist about the rationality/justification of belief and externalist about the rationality/justification of action, he rejects the Enkratic Requirement. I see that as a significant cost. Feldman (1988) once embraced this combination, too, but he dropped it after I pressed him to answer what he thought someone should do if that person had justifiably judged that they're required to ϕ . He went internalist about the justification of action and belief, a coherent but terrifying view when you think about how easy it is for the baddest of the baddies to meet the requirements for forming rational beliefs about obligation on all your standard accounts of rational belief.

In closing, I've explained how externalists can account for some of the intuitions that Cohen draws on in his argument for internalism and argued that there are good reasons for thinking that his strategy for motivating internalism is flawed. The distinction between justification and excuse is perfectly coherent and when properly understood should be sufficient for undermining his argument for internalism. A proper argument for internalism requires some sort of independent argument that shows *on substantive grounds* that all epistemic norms have internal application conditions. Since he has offered no substantive argument for thinking that epistemic norms are that way, I don't think he's offered us any reason to go in for internalism. A proper argument for internalism would be one that purported to show that nothing of epistemic significance could be found outside of the head, one that followed Kant's lead in arguing that everything of moral significance has to do with the quality of the agent's will. In the closing section, I've outlined a strategy for showing that there are epistemically significant factors that aren't in the agent's head. On the view that I'd defend, the crucial question to ask in determining whether you believe what you should about *p* is whether you'd know *p* if you were to believe it. If not, your belief could not be justifiably held. If you would know, it could be justifiably held.

In the opening we saw that Cohen thought that this factive account of justification leads to skepticism. It doesn't face any *untoward* skeptical consequences. His argument that this Cartesian view of justification leads to skepticism rests on two claims. The first is that the view classifies all false beliefs as unjustified. The second is that the conditions that determine whether our beliefs are justified are just those conditions common to ourselves and our deceived counterparts. His argument *assumes* that the supervenience base for justification is just in the head. Once we reject that assumption, there's nothing that's left of that objection. My view says that your belief in *p* is justified iff you know *p*. *You* can't believe that this view leads to skepticism unless you're a skeptic. If you are, that's *your* problem, not mine.

References

- Alvarez, M. 2010. *Kinds of Reasons*. Oxford University Press.
- Arpaly, N. 2002. *Unprincipled Virtue*. Oxford University Press.
- Audi, R. 2001. An Internalist Theory of Normative Grounds. *Philosophical Topics* 29: 31–45.
- Austin, J. L. 1956. A Plea for Excuses. *Proceedings of the Aristotelian Society* 57: 1–30.
- Bach, K. 1985. A Rationale for Reliabilism. *The Monist* 68: 246–63.
- Baron, M. 2005. Justifications and Excuses. *Ohio State Journal of Criminal Law* 2: 387–406.
- Bennett, J. 1995. *The Act Itself*. Oxford University Press.
- Booth, A. and R. Peels. 2010. Why Responsible Belief is Blameless Belief. *The Journal of Philosophy* 107: 257–65.
- Boyd, K. Forthcoming. Assertion, Practical Reasoning, and Epistemic Separabilism. *Philosophical Studies*.
- Broome, J. 2013. *Rationality Through Reasoning*. Wiley.
- Cohen, S. 1984. Justification and Truth. *Philosophical Studies* 46: 279–95.
- Cohen, S. MS. Theorizing about the Epistemic.
- Comesana, J. 2002. The Diagonal and the Demon. *Philosophical Studies* 110: 249–66.
- Darley, J. and P. Robinson. 1998. Testing Competing Theories of Justification. *North Carolina Law Review* 76: 1095–1143.

- DeRose, K. 2002. Assertion, Knowledge, and Context. *Philosophical Review* III: 167-203.
- Dougherty, T. and P. Rysiew. Experience First. In J. Turri, M. Steup, and E. Sosa (ed.), *Contemporary Debates in Epistemology*, 2nd edition. Wiley.
- Duff, R.A. 2007. Excuses Moral and Legal: A comment on Marcia Baron's 'Excuses, Excuses'. *Criminal Law and Philosophy* 1: 49-55.
- Engel, M. 1992. Personal and Doxastic Justification. *Philosophical Studies* 67: 133-51
- Fantl, J. and M. McGrath. 2009. *Knowledge in an Uncertain World*. Oxford University Press.
- Fantl, J. and M. McGrath. 2014. Practical Matters Affect Whether You Know. In M. Steup, J. Turri, and E. Sosa (ed.), *Contemporary Debates in Epistemology*. Wiley-Blackwell, pp. 84-95.
- Feldman, R. 1988. Subjective and Objective Justification in Ethics and Epistemology. *the Monist* 71 405-19.
- Ferzan, K. 2011. Justifications and Excuses. In J. Deigh and D. Dolinko (ed.), *The Oxford Handbook of the Philosophy of Criminal Law*. Oxford University Press.
- Gardner, J. 2004. The Wrongdoing that Gets Results. *Philosophical Perspectives* 18: 53-88.
- Gardner, J. 2007. *Offenses and Defenses*. Oxford University Press.
- Gerken, M. 2011. Warrant and Action. *Synthese* 178: 529-47.
- Gibbons, J. 2010. Things That Make Things Reasonable. *Philosophy and Phenomenological Research* 81: 335-61
- Goldman, A. 1988. Strong and Weak Justification. *Philosophical Perspectives* 2: 51-69.
- Graham, P. 2012. Epistemic Entitlement. *Nous* 46: 449-82.
- Harman, E. 2011. Does Moral Ignorance Exculpate? *Ratio* 24: 443-68.
- Hawthorne, J. MS. The Surprise Examination Paradox.
- Herman, B. 1993. *The Practice of Moral Judgment*. Harvard University Press.
- Horder, J. 2004. *Excusing Crime*. Oxford University Press.
- Hornsby, J. 2007. Knowledge in Action. In A. Leist (ed.), *Action in Context*. De Gruyter, pp. 285-302.
- Hornsby, J. 2008. A Disjunctive Theory of Acting for Reasons. In F. Macpherson and A. Haddock (ed.), *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press.
- Huemer, M. 2001. *Skepticism and the Veil of Perception*. Rowman and Littlefield.
- Huemer, M. 2006. Phenomenal Conservatism and the Internalist Intuition. *American Philosophical Quarterly* 43: 147-58.
- Hyman, J. 1999. How Knowledge Works. *Philosophical Quarterly* 49: 433-51.
- Hyman, J. 2006. Knowledge and Evidence. *Mind* 115: 891-916.
- Lackey, J. 2007. Norms of Assertion. *Nous* 41: 594-626.
- Littlejohn, C. 2009. The Externalist's Demon. *Canadian Journal of Philosophy* 39: 399-434.
- Littlejohn, C. 2011. Defeating Phenomenal Conservatism. *Analytic Philosophy* 52: 35-48.
- Littlejohn, C. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- Littlejohn, C. 2013. The Russellian Retreat. *Proceedings of the Aristotelian Society* 113: 293-320.
- Littlejohn, C. 2014. The Unity of Reason. In C. Littlejohn and J. Turri (ed.), *Epistemic Norms*. Oxford University Press.
- Lowy, C. 1978. Gettier's Notion of Justification. *Mind* 87: 105-8.

- Lyons, J. 2013. Should Reliabilists be Worried about Demon Worlds? *Philosophy and Phenomenological Research* 86: 1-40.
- McKinnon, R. 2013. The Supportive Reasons Norm of Assertion. *American Philosophical Quarterly* 50: 121-35.
- Madison, B. 2014. Epistemological Disjunctivism and the New Evil Demon. *Acta Analytica* 29: 61-70.
- Madison, B. Forthcoming. Epistemic Value and the New Evil Demon. *Pacific Philosophical Quarterly*.
- Mitova, V. Forthcoming. Truthy Psychologism about Evidence. *Philosophical Studies*.
- Moon, A. 2012. Three Forms of Internalism and the New Evil Demon Problem.
- Moore, Michael. 1997. *Placing Blame*. Oxford University Press.
- Nottelman, N. 2007. *Blameworthy Belief*. Springer.
- Peels, R. 2011. Tracing Culpable Ignorance. *Logos & Episteme* 4: 575-82.
- Prior, A. 1958. Epimenides and the Cretan. *Journal of Symbolic Logic* 23: 261-66.
- Robinson, P. 1996. Competing Theories of Justification: Deeds vs. Reasons. In A. Smith and A. Simester (ed.), *Harm and Culpability*. Oxford University Press.
- Rosen, G. 2014. Culpability and Duress: A Case Study. *Proceedings of the Aristotelian Society*, Supplementary Volume 88: 69-90.
- Silins, N. 2005. Deception and Evidence. *Philosophical Perspectives* 19: 375-404.
- Smithies, D. 2013. Review of Duncan Pritchard's *Epistemological Disjunctivism*. *Notre Dame Philosophical Reviews*.
- Steup, M. 1988. The Deontic Conception of Epistemic Justification. *Philosophical Studies* 53: 65-84.
- Steup, M. 2012. Justification, Deontology, and Voluntary Control. In S. Tolksdorf (ed.), *Conceptions of Knowledge*. DeGruyter.
- Sylvan, K. MS. Rationality and Justification: Reasons to Divorce?
- Tadros, V. 2005. *Criminal Responsibility*. Oxford University Press.
- Turri, J. 2009. The Ontology of Epistemic Reasons. *Noûs* 43: 490-512.
- Turri, J. 2013. Bi-Level Virtue Epistemology. In J. Turri (ed.), *Virtuous Thoughts: The Philosophy of Ernest Sosa*. Springer.
- Unger, P. 1975. *Ignorance*. Oxford University Press.
- Wallace, R.J. 1998. *Responsibility and the Moral Sentiments*. Harvard University Press.
- Weatherston, B. 2008. Deontology and Descartes' Demon. *The Journal of Philosophy* 105: 540-69.
- Wedgwood, R. 2002. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349-69.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. Forthcoming. Justifications, Excuses, and Skeptical Scenarios. In F. Dorsch and J. Dutant (ed.), *The New Evil Demon*. Oxford University Press.
- Zimmerman, M. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.