

The Externalist's Demon

Clayton Littlejohn

(Forthcoming in the *Canadian Journal of Philosophy*)

Abstract

Internalists often claim that theirs is the view that best accommodates our epistemic intuitions. To make this vivid, they ask us to imagine subjects whose mental lives match our own but are systematically deceived about their surroundings and claim that these subjects are no less justified in their beliefs than we are in ours. Because it seems that externalist accounts of epistemic justification cannot accommodate this intuition, it seems we have good reason for preferring some sort of internalist account. In this paper, I want to show that arguments that appeal to intuition to motivate internalist accounts of justified belief are highly problematic. I also want to show that if we draw a distinction between two types of justification ascription, the externalist has no difficulty at all accommodating the intuitions commonly taken to motivate rival internalist views. We have to distinguish between a *believer* who is justified in her beliefs and *beliefs* that are justified. While the internalist account of *personal* justification is quite plausible, intuition gives us little reason if any for adopting an internalist account of *doxastic* justification. Critics will say that the personal/doxastic justification distinction is incoherent or unmotivated, so I shall explain what this distinction amounts to and why we need it to make sense of certain features of our ordinary practice or normative appraisal.

Introduction

Let's say that two individuals are epistemic counterparts iff they happen to be in precisely the same (non-factive) mental states. If one has a veridical experience, her counterparts will undergo a subjectively indistinguishable conscious experience. If she remembers something, her counterparts will seem to recall the same event or fact. If she knows something to be true, her counterparts will believe it to be true. Counterparts always find the same things intuitive. Any difference between those who know a great deal about the external world and their systematically deceived counterparts is a difference the deceived counterparts could never appreciate.¹

Suppose that when it comes to her mundane beliefs about her immediate surroundings, Audrey gets things right as a rule. Audrey's counterpart Cooper is systematically deceived. Audrey sees that she is holding a book of matches, knows that she wants to light her stove, and strikes a match to light it. Coop seems to see a book of matches, thinks he has struck a match, and believes he has lit his stove. Because of the demon's efforts, Coop undergoes these experiences in the dark. He strikes non-existent matches with phantom limbs. Intuition suggests that while Coop might be mistaken in nearly everything he believes about the external world, he is no less justified in believing what he does than Audrey is justified in believing what she does.

¹ Brewer (1999) suggests that if you have perceptual knowledge of the external world, you might know that you are not in the unfortunate situation of your deceived counterparts. He might say that there is a sense in which you can appreciate the difference between your situation and the situation of someone deceived while allowing that your counterparts who undergo subjective indistinguishable hallucinatory experiences are in no position to appreciate the difference between their situation and yours.

To bolster the intuition and explain why we ought to regard mistaken beliefs backed by hallucinatory experiences as justified and not merely excused, Audi offers these remarks:

... given the vivid hallucination, I am in no way at fault for believing what I do, nor do I deserve criticism. 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*.²

In general, it seems that what we ought to say about epistemic counterparts is this:

Parity: Necessarily, if S and S' are epistemic counterparts, S is justified in believing *p* iff S' is justified in believing *p*.

Once we accept Parity as true, it is hard to see how we could reject Internalism or refuse to reject Externalism:

Internalism: Necessarily, the facts about the justification of a subject's beliefs are fixed completely by the facts about the subject's non-factive mental states.

Externalism: It is not the case that the facts about the justification of a subject's beliefs are fixed completely by the facts about the subject's non-factive mental states.³

To see why, consider two versions of Externalism. According to process reliabilism, you cannot have a justified belief about *p* unless the processes that led to that belief reliably lead to the truth.⁴ According to the knowledge account, you cannot have a justified belief about *p* unless you know *p*.⁵ While Audrey arrives at her beliefs by means of reliable processes and those beliefs constitute knowledge, Coop arrives at his beliefs by the most unreliable of means and knows nothing of the external world. Assuming Parity is true, we have to say that he is no less justified than she in spite

² Audi (2001: 23).

³ Externalism, so understood, is broad enough to encompass Goldman's (1986) reliabilism, Sutton (2005) knowledge account of justified belief, and Williamson's (2000) view that a belief is justified by the evidence which consists of propositions known.

⁴ The new evil demon problem first surfaced in Stewart Cohen's (1984) discussion of the reliabilist views defended in Goldman (1979). The intuition figures prominently in more recent criticism of externalist views. See Audi (2002), Bird (2007), Conee and Feldman (2004), Graham (forthcoming), Nelson (2002), Pollock and Cruz (forthcoming), and Wedgwood (2002).

⁵ Defended by Unger (1975) and more recently by Sutton (2005). Such a view seems to be a consequence of Hawthorne and Stanley's (Forthcoming) account of justified use of belief in theoretical deliberation.

of these external differences. Thus, it seems justification must be fixed by the non-factive mental states that Audrey and Cooper share in common. The point seems perfectly general. If you take the justification of belief to involve some condition that does not strongly supervene on a subject's non-factive mental states, it seems you've adopted a view of justification incompatible with Parity. That your view is incompatible with Parity strongly suggests that your view is mistaken.

Most think that Parity and Externalism are incompatible.⁶ Many think that their incompatibility gives us good reason to reject Externalism.⁷ I shall argue that there is no good argument from Parity to Internalism. In §1, we will see that there are two problems with arguing from Parity to Internalism. First, some versions of Externalism are consistent with Parity. As these views are incompatible with Internalism, Parity alone cannot give us conclusive reason to accept Internalism. Second, anyone who accepts Internalism on the grounds that Parity is true has to explain away the intuitions and theoretical considerations that motivate Externalism. It seems, *ceteris paribus*, a view that accommodates the widest range of intuitions and theoretical observations is preferable to a view that requires you to explain away the considerations taken to motivate Internalism or Externalism. In §2, we shall see that Externalism can accommodate the intuitions taken to motivate Internalism if we revise some commonly held views about the logic of justification ascriptions. In that section, I shall also argue that we have good reason to revise these views about the logic of justification ascription. The aim of this paper is not to settle the debate between internalists and externalists, but to show that Externalism is compatible with the intuitions taken to motivate Internalism and that a strategy for settling this debate is far less promising than internalists have taken it to be.

To preview a bit, there are really two thoughts that lead people to say that Externalism cannot be squared with our intuitions. First, there is the observation that considerations wholly inaccessible to a subject cannot be relevant to whether this subject is justified in her attitudes. I shall not take issue with this observation. Second, there is the assumption that only those considerations relevant to whether a subject is justified in her beliefs are relevant to the justification of her beliefs. There is, I shall argue, good reason to distinguish between the considerations that determine whether a *believer* is justified (i.e., considerations that have to do with personal justification) and the considerations that determine whether a *belief* is justified (i.e., considerations that have to do with doxastic justification). It is the second assumption that ought to be rejected. If it is possible for a believer to be justified in holding beliefs that are not justified,

⁶ Notable exceptions to this include Bach (1985) and Engel (1992).

⁷ A notable exception to this is found in Brewer (1997) who insists that hallucinatory experiences that lead to false beliefs do not give the believer the right to believe, but only appear to.

externalists can say that personal justification is an internalist notion. It would not follow that doxastic justification is also an internalist notion. It may well be an externalist notion. Some will say that the distinction between a believer who is justified in holding a belief and a believer holding a justified belief is *ad hoc*. I shall offer two reasons for thinking that we ought to distinguish between personal and doxastic justification. The first is that drawing the distinction enables us to accommodate the widest range of intuitions possible. If the distinction is not incoherent, and I shall argue that it is not, the desire to accommodate the widest range of intuitions gives us some reason to recognize this distinction. The second is that such a distinction is needed to make sense of the more familiar distinction between the permissible and excusable. If the personal/doxastic justification distinction is needed to make sense of our ordinary practice of normative appraisal, it cannot be out of bounds to appeal to this distinction in the context of the Internalism/Externalism debate.

1. The New Evil Demon Problem

The problem of combining our theoretical views about epistemic justification with the intuitions about undetectable error (i.e., those that support Parity) is commonly referred to as ‘the new evil demon problem’. Nearly everyone thinks it is a problem for Externalism. Many think the problem is insurmountable. Because it seems Externalism clashes so dramatically with our ordinary intuitions, it might be a good idea to start by explaining why someone might think that justification must be an externalist notion.

The rationale for Externalism that I find persuasive begins with the observation that anything that could justify believing p now in light of the information available rather than suspending judgment, it had better justify assuming the risk that comes with committing to a view about p rather than remaining agnostic.⁸ If someone’s belief is adequately justified, there must be a justification for that belief that does what justifications are supposed to. If that justification does what it is supposed to, it justifies closing deliberation prior to the acquisition of additional information. From the subject’s point of view, having come to believe p she will be inclined to think that any further information that comes to light will confirm something she’s taken herself to

⁸ The rationale for Externalism I am offering here is different from the rationale found in earlier defenses of Externalism. Sosa (1985), for example, seems to defend Externalism on the grounds that whatever justifies a belief cannot be the sort of thing that obtains wholly independently from the truth of the belief because of the role that justification must play in the production of knowledge. The argument offered here rests on no assumptions about the role justification plays in the production of knowledge. Motivated in this way, I think Externalism is immune to the objection that externalists have conflated justification and warrant. See Pollock and Cruz (forthcoming) for a version of this objection.

know all along (i.e., that p is true), be perceived as misleading evidence (i.e., evidence that misleadingly suggests something she knows is not true), or show that she has made a mistake in having concluded that p is true on the basis of the information she had. Perhaps such a mistake is excusable, but such mistakes show that what she took to be a justification was no justification at all. A defective justification, you might say, is no more a justification than a decoy duck is a duck. You cannot have a justified belief about p without a justification for believing p . You cannot have a justification for believing p if there are considerations that establish that it was wrongful for you to have judged that p is true on the basis of the considerations available to you when they judged that p is true.

Why think this picture favors Externalism? I might say that you cannot defend Coop's believing that p in light of the fact that Coop had no evidence that supported his belief that p . Because there is no defending him for having believed p , I might conclude that his belief about p was not justified. That Coop cannot be defended in light of how he came to the belief is something we establish (in part) by examining those facts that supervene on his non-factive mental states. No one should deny that a belief might fail to be justified in virtue of how things stand internal to the subject. The reason that this picture is supposed to favor Externalism is that we do not think that beliefs are criticized solely on the basis of considerations available to the believer. If our criticism implies fault, perhaps our criticism is based upon considerations available to Coop. However, not all criticism of someone's belief implies fault. We sometimes we criticize a belief because of considerations we know to be unavailable to believer without implying that the believer ought to be faulted for her mistake.

Some externalists will say that it is part of our ordinary practice of epistemic appraisal to say that Coop should not believe p because he is non-culpably mistaken.⁹ Others will say that it is part of our ordinary practice to say Coop should not believe p because he has unwittingly arrived at that belief by way of a method that is wholly unreliable.¹⁰ If what it is to have a justified belief is to have a belief that can be defended against criticism (whether that is criticism based on considerations available to the believer or available only to those who appraise the believer's beliefs) it seems to follow that the conditions that determine whether a belief is justified include those that supervene on our non-factive mental states but are not limited to such conditions. Justification is an externalist notion because defensibly believing something is an externalist notion.

⁹ I imagine Sutton (2005) and Unger (1975) might say this.

¹⁰ I would expect Goldman (1986) to say something to this effect.

Internalists will likely say that criticisms that appeal to considerations beyond those available to the believer are not criticisms that a justification seeks to address.¹¹ This point, if taken seriously, leads to an odd sort of view. On this view, the justification of a belief is described as something that does not even aspire to show that the subject is right to judge p by showing that considerations not presently accessible to the subject will not show that she was mistaken to judge that p is true on the basis of what was available to her. (If the subject was non-culpably mistaken to judge that p is true on the basis of what was available to her, on many internalist views this shows that there was nothing wrong with the subject's justification.) It does not seem from the subject's point of view that having an adequate justification is neutral on the question as to whether there is some basis for correctly criticizing her belief as mistaken. If it did, theoretical deliberation would not close with the subject taking the available justification to be adequate. Perhaps internalists might say that a justification might succeed in justifying a belief while failing to do what it is intended to do or what the subject thinks it must do in order to justify adopting a belief rather than remaining agnostic (i.e., constitute a defense of having come to believe p on the basis of the available information). The awkwardness of such a view emerges when we realize that this view has it that beliefs might be perfectly justified by defective justifications or justifications the subject is disposed to think of as defective.

As odd as such a view might look, internalists are quick to remind us that while we might say that people should not believe what they do because they are mistaken or because they have unwittingly relied on an unreliable process, when our basis for saying that they should not believe something appeals to considerations beyond those that pertain to the subject's non-factive mental states we will say that the subject was justified in believing p . So, we have to address their new evil demon argument against Externalism:

- (1) Although Coop is systematically deceived, he is justified in believing what he believes about the external world.
- (2) Because Coop is systematically deceived, he does not come to his beliefs on the basis of reliable processes or the sorts of veridical experiences and genuine memories that can serve as the basis for knowledge.
- (C) Hence, none of the conditions beyond those that supervene on our non-factive mental states have any bearing on the justification of our beliefs.

¹¹ This seems to be Bird's (2007) view.

Although the argument might initially seem quite persuasive, the argument's implicit assumptions are far from innocent. Intuitions about counterparts clearly play an important role in the argument, but the argument does not rest on intuition alone. The additional assumptions needed to argue from Parity to Internalism are assumptions we have good reason to reject.

1.1 A Hidden Assumption

In this section, I want to point to two problems with arguing from Parity to Internalism. First, the argument does not exclude certain forms of Externalism unless we introduce additional assumptions. These assumptions we have little reason to accept and good reason to reject. Second, there are intuitions that seem to favor Externalism and seem to be compatible with the intuitions taken to motivate Parity. Just as the externalist faces the task of accommodating or explaining away the intuitions taken to motivate Internalism, the internalist cannot claim that theirs is a view unequivocally supported by intuition.

The argument from Parity to Internalism is invalid. According to Parity, if two individuals are epistemic counterparts, they are justified in believing the same things. According to Internalism, the facts about the justification of an individual's beliefs strongly supervene on facts about that individual's non-factive mental states. Internalism is logically stronger than Parity. For all that Parity tells us, Audrey and Cooper might be justified in believing precisely the same things even if Audrey's justification for her beliefs differs from Coop's. Parity is consistent with a view on which both Audrey and Cooper believe what they do on the basis of evidence that warrants our saying that they are justified in their beliefs while insisting that Audrey's evidence goes beyond Coop's evidence. If Audrey's evidence outstrips Coop's evidence, it seems Audrey has 'external' epistemic reasons for her beliefs that Coop does not and Internalism tells us could not exist. For all that Parity tells us, Audrey's reasons for her beliefs are *better* than Coop's. This comparative judgment is incompatible with Internalism, but it is consistent with Parity. Thus, Parity does not entail Internalism.

To rule out such responses, what the argument for Internalism requires an additional assumption:

Equality: Necessarily, if S and S' are epistemic counterparts, the fact that *q* is relevant to the justification of what S believes iff the fact that *q* is true is relevant to the justification of what S' believes.

Equality comes with costs I'm not sure we should willingly pay. Assume that some evidence is propositional. Assume knowledge of a proposition's truth is sufficient for that proposition's

inclusion in your evidence. Assume false propositions cannot be part of our evidence.¹² Finally, assume that the sceptic is wrong and it is possible to have knowledge of the external world. With these assumptions in place, suppose Audrey knowingly infers r from p . Coop mistakenly believes p and infers that r is true. Audrey's knowledge of p is based on veridical experience. Cooper's belief that p is true is induced by hallucinatory experience. If knowledge of p 's truth is sufficient for p 's inclusion in Audrey's evidence, among the facts relevant to the facts about the justification of her beliefs is that p is part of her evidence. If the falsity of p is sufficient for p 's exclusion from Coop's evidence, there are reasons for Audrey to believe r that are not among the reasons for Coop to believe this. To save Equality, you either have to say that we cannot have knowledge of propositions about the external world, that there can be false propositions included in someone's evidence, or that knowledge is not enough for a proposition's inclusion in someone's evidence. Whatever reason you have for rejecting one or more of these assumptions, it cannot be that these assumptions conflict with Parity or the intuitions that motivate it. These assumptions, although incompatible with Equality, are perfectly consistent with Parity.

I do not offer this in the hopes of refuting Internalism, but merely to show that an assumption needed to argue from Parity to Internalism comes with costs. I know of no one who thinks that false propositions can serve as evidence. It strikes me as exceptionally odd to say of someone something to the effect of, 'Although they know they have hands, that is no reason for them to believe they are not handless'. Yet, if knowledge of a proposition's truth did not suffice for that proposition's inclusion in someone's evidence, I just cannot see how there could be a principled reason to think such a remark must be mistaken.

1.2 Further Intuitions

The first problem with running the argument from Parity to Internalism is that the argument requires an assumption about epistemic relevance that has no clear motivation and comes with real costs. The second problem facing the internalist who wants to claim that theirs is the view supported by intuition is that it seems there are intuitions that support the following thesis:

¹² Williamson (2000) defends each of these assumptions in the course of arguing that our evidence consists of all and only what we know ('E = K' hereafter). A defense of E = K is clearly beyond the scope of this paper. Note that the assumptions listed above did not include what I take to be the most controversial aspect of Williamson's view, which is that knowledge of p 's truth is *necessary* for p 's inclusion in your evidence. This assumption has the unfortunate consequence that being in a Gettier case can have an affect on your evidence. Silins (2005) has argued that E = K conflicts with independently plausible assumptions about the sort of access we have to our evidence. Dodd (2007) has argued that E = K commits us to an infallibilist account of justification that engenders scepticism. I address these objections in Littlejohn (ms. a) and (forthcoming).

Asymmetry: It is possible for there to be a pair of epistemic counterparts, S and S', such that (a) only one of the pair has good enough reason to believe p and have a justified belief about p or (b) S has better reasons to believe p than S' has.

For reasons sketched above, I think it is clear that Asymmetry is incompatible with Internalism. So, if there are some intuitions that support Asymmetry, it is misleading to suggest that Internalism finds unequivocal support in intuition.

Imagine Audrey imagines that her counterpart Cooper has been wholly deceived by a Cartesian demon. She thinks to herself:

- (1) If there is no more reason for me to believe I have hands than there is for Coop to believe he has hands, I should not believe I have them.
- (2) If my reasons for believing I have hands are no better than his, I really have no good reason for believing I have hands.

Intuitively, these strike me as correct. Contrast them with the following thought:

- (3) Even if there is no more reason for me to believe I have hands than there is for Coop to believe he has hands, there is nothing wrong with my continuing to believe I have hands.

I do not think that the oddity of (3) is due to its first-personal formulation. Suppose Audrey is told that Ben and Cooper are both her epistemic counterparts. She is told that Coop is systematically deceived. She is not told whether Ben is systematically deceived or any less reliable in his judgments about the external world than she is. I think it would be right for her to say:

- (4) If there is no more reason for Ben to believe he has hands than there is for Coop to believe he has hands, Ben should not believe this.
- (5) If there is no more reason for Ben to believe he has hands than there is for Coop, I would advise him against believing this. If, however, he was in a position to know that he has hands I would not advise him against believing this.

If your scorecard is anything like mine, you reject (3) but accept (1), (2), (4), and (5). The intuitions that lead someone to accept (1), (2), (4), and (5) support Asymmetry. If my intuitions are correct,

there is intuitive support for both Asymmetry and Externalism. Intuitions being what they are, I cannot assume that they are universally shared. Someone might object that the intuitions I am hoping to elicit are intuitions shared only by infallibilist sceptics or externalists.¹³ Let me say three things in response. First, the intuitions do support Externalism, but they do not support infallibilism. Although these intuitions are consistent with infallibilism, they are also accommodated by fallibilist views such as reliabilism or the knowledge account of justified belief.¹⁴ Second, even if no one with internalist sympathies shares these intuitions, it is important to remember that the aim of this paper is *not* to refute Internalism, much less refute Internalism appealing to considerations internalists accept. The aim is to show how it is possible to accommodate the widest range of intuitions possible within an externalist framework. We cannot deny that people have these intuitions even if those intuitions strike us as mistaken. Confronted by those who have such intuitions, we have to choose between taking them at face value and building our theories around them or explaining them away. I would rather accommodate the intuitions of both internalists and externalists if possible. Third, I do not think we can assume that those who find Internalism intuitive will find (1), (2), (4) and (5) counterintuitive. In trying to elicit the intuitions that supported Parity, we stipulate that the relevant subjects are mistaken in their beliefs. This means that the question as to whether we ought to agree with the relevant subjects is no longer open. This shifts our focus away from a question along the lines of, ‘What are we to think of what they believe?’ to a question more along the lines of, ‘Given that they are mistaken, what are we to think of them?’ In trying to elicit the intuitions I think support Asymmetry, the question as to whether the subject is mistaken is left open as is the question, ‘Given what their reasons are, ought we agree with them in believing what they do?’ By focusing explicitly on the quality of reasons rather than on the subject in the wake of that subject’s mistaken response, I do not think it would be surprising to find those with internalist sympathies or no settled view finding (1), (2), (4), and (5) intuitive.

I wish to make one more point in favor of Asymmetry because I think this sort of argument has not yet surfaced in the literature on epistemic justification. Empirical research

¹³ An anonymous referee raised this worry.

¹⁴ Sutton’s (2005) view on which a belief is justified only if that belief constitutes knowledge amounts to a version of fallibilism as it is typically understood. As it is typically understood, the infallibilist asserts that a belief can be justified only if there are infallible grounds or evidence available for that belief. See Cohen (1988). Because of the factivity of knowledge, Sutton’s view has the implication that there are no false, justified beliefs. However, the view does not have the implication that there are no justified beliefs based on fallible grounds because there is no reason to assume that knowledge requires having infallible grounds for a belief. Thus, Sutton’s view is a version of fallibilism as it is typically understood in spite of the fact that it rules out the possibility of a false belief being justified.

suggests that when it comes to the justification of action, the folk think that the justification of an action depends on more than just what is in the agent's head. I would be surprised if the folk were disposed to think that the justification of action depends on more than just what is in the agent's head but when it came to the justification of those beliefs that lead to action the justification in no way depended on factors external to the subject's perspective. If folk intuition treats actions motivated by non-culpably held mistaken beliefs as excusable rather than justified, why would the folk think of the beliefs that led to these actions as perfectly justified?

An internalist is free to deny the relevance of such intuitions when it comes to selecting between competing theories of *epistemic* justification, but I think that given their theoretical commitments, internalists will have a difficult time making sense of a hybrid view that allows that the justification of action depends (in part) upon features of the objective circumstances while denying that the justification of belief depends on these same features. If the hybrid view does not sit well with the internalist's theoretical commitments, the internalist either has to explain away the folk intuitions about justified action or might help themselves to the sort of distinction I shall argue saves the externalist from the charge that theirs is a view that conflicts with intuition and undermines the new evil demon argument. They can square their view with folk intuition by saying that theirs is a view that concerns *personal* justification and is thus a theory about what it takes for the person to be justified in the face of certain sort of criticism. In so doing, however, they would have to concede that the new evil demon argument is not a cogent argument against Externalism or externalist accounts of *doxastic* justification.

So, let us consider a view that combines an externalist account of justified action with an internalist account of justified belief. On this hybrid view, the justification of an action depends on whether the agent has acted against an undefeated reason against performing that action. That the agent is unaware of the reason might serve as an excuse, it does not show that the reason does not so much as bear on the justifiability of the relevant action. If the agent unwittingly harms another and is not culpable in her failure to appreciate that her actions will harm another, the action is *prima facie* wrongful even though the agent cannot be faulted for having caused the harm. If there is no overriding reason that justifies the imposition of that harm, the action will turn out to be all things considered wrong. However, on this view, if it seems 'from the inside' that there is no overriding reason to refrain from performing the action and the agent cannot be faulted for failing to appreciate that there is such a reason, the belief that leads her to act (e.g., the belief that this is the thing to do) is itself justified.

To see what a view amounts to, let us consider two concrete examples. Ben has fallen behind in his payments to Jacques, a loan shark. Last week, Jacques gave Ben a severe beating

with the warning that if he is not paid within the week, he will kill Ben. Ben starts carrying a gun and spending most of his time in public places in the hopes that Jacques would not attack him if there were people present to witness the attack. One afternoon, Ben is in the diner. The story might continue in one of two ways. The first version of the story goes like this. You see Jacques' twin brother enter the diner. You know it to be Jacques' twin, know that Ben knows nothing of this twin, and know that Jacques' twin is harmless. You know he does not look harmless as he looks just like the very dangerous Jacques. Seeing the man approach, Ben produces his gun and yells 'You won't get me Jacques!' You know that if you do nothing Ben will shoot Jacques' twin on the belief that this is necessary for self-defense and so you lunge at Ben. You knock him down, his shot misses Jacques, and Ben is injured as a result. The second version of the story goes like this. You see Ben enter the diner and Jacques follow soon after. Seeing Jacques approach, Ben produces his gun and yells 'You won't get me Jacques!' You see his first shot miss and Jacques pull out a gun to return fire. As Ben struggles to reload, you hit Jacques with a chair. You knock him down, he loses his gun, and he is injured as a result of your striking him.

It seems that in both situations your reactions were morally justified. In fact, it seems that they were supererogatory. What does this tell us? For one, it suggests that in the first example, Ben lost his right to non-interference. He loses no such right in the second. This suffices to explain the further intuition that you owe Ben nothing for his injuries in the first story but would had you intervened in just this same way in the second. It seems so because the grounds of self-defense were not present in the first example. That the grounds of self-defense did not obtain in the first case suggests that his action was less justified than the action he performed in the second case. Yet, it seems that the relevant mental states were the same in both cases. Thus, it seems that it is not true that the justification of Ben's actions depends wholly on Ben's non-factive mental states.

This much seems to be supported by folk intuition. If this much is right but we were to insist that while the justificatory status of his actions differed in these two cases the justificatory status of his beliefs were the same in both cases, we would have to reject the following principle, linking justified belief to justified action:

Link: If S's belief that Φ -ing is permissible is justified, S's Φ -ing is justified.

The idea behind Link is rather straightforward. Suppose we say that Ben's pulling the trigger just was his shooting Jacques. In turn, his shooting Spike just was his acting on the beliefs and desires that involve, *inter alia*, the belief that the grounds for self-defense justified his shooting. While believing p and allowing yourself to reason from the belief that p are two different things and

Leibniz's Law will not compel us to assign the same deontic status to the believing and the reasoning from the belief, it seems odd to think that there is something about the reasoning from the belief that prevents it from being justified and *nothing* that prevents the belief itself from being justified. Yet, that seems to be what the hybrid view is committed to insofar as it concedes that the acting on the relevant belief, which just was Ben's shooting, was not justified while the relevant belief (i.e., that shooting is the thing to do) is itself perfectly justified. Such a view would require the internalist to say that there is one set of considerations that determine whether it is permissible to let the belief that p figure in deliberation and another that determines whether to believe p . I'm not sure it is coherent to distinguish these considerations in this way.

I think we can show that the costs that come with denying Link are costs internalists should be unwilling to pay. Moreover, I think we can establish Link on the grounds that it is incoherent to draw the distinction between the considerations that determine whether relying on the belief that p in deliberation is justified and the considerations that determine whether believing p is justified. First, it is worth noting that you cannot reject Link without also rejecting the following principle:

Fault: If S's Φ -ing were unjustified, S could be faulted for
having Φ 'd.

If you reject Link, it is possible for someone to have Φ 'd having a justified belief that they ought to Φ even if the Φ -ing itself could not be justified. Surely if someone was not wrong to believe that they ought to Φ or are permitted to Φ and Φ 'd from the right sort of motives, they could not be faulted for having Φ 'd. Whatever we pointed to in order to show that they could be faulted would have to be conditions of which they were culpably unaware or aware of but culpably failed to take proper account of in deliberation. Either way, we could not sustain the judgment that the subject's belief about the deontic status of Φ -ing was itself justified. I don't think internalists would be happy to reject Fault since when it comes to justified belief they often wish to say that the conditions you could not be faulted for are conditions that have no bearing on the justification of your beliefs. While they might say that when it comes to *belief*, conditions for which you cannot be faulted for failing to take account of do not affect justification, but when it comes to *action*, the situation is different. But, it is not clear how they could defend this. It cannot be defended on the grounds that it is part of the concept of justification that the conditions that determine whether something is justified is determined wholly by conditions the subject can be faulted for failing to take account of. They have just denied this. So, it seems that denying Link is costly for internalists insofar as the denial of Link commits you to the denial of Fault.

Not only is the denial of Link costly for internalists, the denial of Link also seems incoherent. On a view that asserts that the justification of action can depend on factors the agent cannot be faulted for failing to take account of but denies that the justification of belief can depend on such factors, it seems that these should be open questions:

- (6) I know that I should believe that I ought to Φ , but what should I do?
- (7) I know that he should believe that he ought to Φ , but should he?

I don't think these are open questions. From the deliberative perspective, there is no intelligible gap between the judgment that you ought to Φ and the judgment that you ought to believe that you ought to Φ . From the perspective of the outside observer, there seems to be no intelligible gap between the beliefs you think the advisee ought to hold about what they ought to do and the actions you think the advisee ought to perform. However, if you combined an internalist account of justified belief with an externalist account of justified action, you would have to allow for the possibility of situations in which an agent ought not perform the actions they ought to believe that they ought to form. You would have no explanation as to why these questions seem closed from both an internal and external standpoint.

If an internalist were to accept Link, they could avoid these difficulties but then their view would conflict with the intuitions that lead people to endorse views on which the justification of an action depends (in part) upon conditions external to the agent's perspective on the situation. Thus, the internalist would be open to the charge that theirs is a view that fails to accommodate intuition much in the way that they suggest the externalist cannot accommodate intuition. Of course, the intuitions that cause trouble for these views are different intuitions, but it is misleading to suggest that Internalism is the view supported by intuition. It is a view that receives support from *some* intuitions, perhaps, but it conflicts with others and is not unique in its ability to accommodate the intuitions that underwrite the new evil demon argument. At least, that is what I hope to show in the next part of the paper.

To sum up, we have seen that internalists that look to motivate their view by appeal to the intuitions backing Parity face a pair of problems. The first is that those intuitions alone cannot motivate Internalism unless combined with a further assumption about epistemic relevance that may prove costly to accept. Reject Equality and the argument from Parity to Internalism fails. The second is that some hold intuitions that appear to be consistent with the intuitions that support Parity but inconsistent with internalist accounts of epistemic justification. First, there were the intuitions about comparative epistemic judgments that suggested that having

a justified belief might require having better reasons than those available to the systematically deceived. Second, there were intuitions about justified action that indirectly seemed to support Asymmetry. As Asymmetry might prove to be consistent with Parity but inconsistent with Internalism, we can now see that the internalist cannot claim that theirs is a view that accommodates the widest range of intuitions possible. Although internalists could deny the relevance of intuitions about justified action when it comes to selecting between competing theories of epistemic justification, to dismiss the relevance of these intuitions the internalists have to reject Link. It seems that such a denial is costly as it forces the internalist to deny that there is any conceptual connection between factors a subject can be faulted for failing to take account of and factors that bear on justification. Moreover, it seemed that the denial of Link is incoherent. In short, the new evil demon argument for Internalism or against Externalism is far more problematic than I think people have appreciated previously.

2. Justification Ascriptions

Taken at face value, a certain kind of intuition supports Parity. Taken at face value, a certain kind of intuition supports Asymmetry. If you thought Parity supported Internalism and Asymmetry entailed Externalism, you might conclude that our intuitions support two incompatible accounts of epistemic justification. The pessimist might conclude that such intuitions cannot be trusted and should not serve as the basis for theory selection. It seems that the intuition that our deceived counterparts are no less justified than our successful epistemic counterparts shows that it is possible for someone to be justified in holding their beliefs provided that things appear a certain way to them regardless of what is actually taking place around them. Other intuitions suggest that the conditions necessary for having a justified belief could only obtain if the reasons someone has for their beliefs are better than the reasons the systematically deceived have for theirs. I think pessimism is unwarranted. We can accommodate both sets of intuitions.

The obstacle to accommodating both sets of intuitions seems to be this. Supposing that Audrey knows p on the basis of veridical perception and Coop believes p on the basis of hallucinatory experiences indistinguishable from Audrey's experiences, we want to say:

- (1) Audrey and Coop are both justified in believing p .

In turn, this is supposed to entail:

- (2) There is sufficient justification for Audrey and Coop's beliefs about p .

We have seen that by denying Equality, a defender of Externalism might accept (1) as well as the inference to (2), but say:

- (3) The justification that Coop has for his beliefs is sufficient.
It is distinct from the justification that Audrey has for her beliefs as the justification Audrey has incorporates all of Coop's evidence and evidence he lacks.¹⁵

The claim that the justification of Audrey's beliefs is overdetermined is, perhaps, true enough. However, this response I think does not quite go far enough. For we have seen that some of our intuitions about counterparts support stronger versions of Externalism that are incompatible with (2). That is to say, we saw in the last part of the previous section that there was some intuitive support for the claim that not only could it be that Audrey's reasons for holding her beliefs are *better* than the reasons Coop had for holding his, but also that only Audrey's reasons were good enough. And while this is consistent with the idea that Audrey's beliefs are supported by more justification than she might need to permissibly hold her belief, it suggests that there might be something right about a version of Externalism on which (2) turns out to be false.

Someone will say that if we reject (2) we must reject (1) as well. And that suggests that we have to reject Parity. In turn, that suggests that internalists are right to say that Externalism clashes with firm intuitions about justification. In this final section I shall argue that the way to sort this mess out is to deny that claims like (2) are logically entailed by claims like (1). If this is right, we can say that externalists are right to reject (2) and their critics are wrong to accuse the externalist of failing to accommodate the intuitions that back (1).

On the standard view of justification ascriptions, (1) entails (2). On the standard view of things, we can say that three kinds of justification ascription are related as follows:

- (SV1) S is justified in believing p iff S's belief that p is justified.
(SV2) S's belief that p is justified iff there is a sufficient justification for S to believe p and this is the basis on which S believes p .

According to (SV1), personal and doxastic justification ascriptions are equivalent. According to (SV2), ascriptions of propositional justification are logically entailed by ascriptions of doxastic justification. If, as I shall argue, we ought to revise this account of the logical relations between various kinds of justification ascription by rejecting (SV1), then perhaps we can make sense of a

¹⁵ This seems to be Williamson's (2000) position. Collins (1997) defends a similar view. Because neither author takes account of the intuitions I have claimed support Asymmetry, I think they see this as the only option for Externalism. It is an option, to be sure, but it is not the only one.

view on which (1) is true and (2) is false. We can try to make sense of a view on which someone can be justified in believing p even if there is not sufficient justification for the belief they hold. Such a view would make sense of those intuitions that supported Parity and Asymmetry.

I shall offer a sketch of an account of justification and explain why someone attracted to such a view of how justifications work will likely think that the standard view oversimplifies matters. Next, I will look at a pair of defenses of the standard view and argue that neither is adequate. I will conclude with an argument for revising the standard view on the grounds that doing so is necessary for making sense of a more familiar distinction between excuses and permissions.

2.1 What Justifications Do

Various things call for justification. You might be called on to justify a decision, or a belief, or someone who made a decision or held a belief. What is it that a justification must do if it is to be successful? On the view I would like to propose, we should think of justifications as defenses:

(JD) To successfully justify something is to successfully defend it.

That which cannot call for a defense against criticism is not something that calls for a justification. If a justification is called for and a defense cannot be given, I cannot imagine that a justification might be given in its place. You might worry that there is more to a justification than just a defense. After all, to say that a belief is justified is to say that there is some positive reason for holding it.¹⁶ While that is true, I think it causes no trouble for (JD). The reason I think this does not cause trouble for the proposal is that if we have identified that someone believes p and we know they have no positive reason for holding that belief, this fact alone will prevent us from offering a proper defense of that particular belief. If you assume that someone's believing p is an instance of epistemic irresponsibility and a defense of that belief is going to involve something that establishes that the believer who holds this particular belief has not been irresponsible, it follows from (JD) that this belief cannot be justified precisely because the believer has no positive reason to believe p . Thus, (JD) pronounces such beliefs unjustified.

(JD) is a thesis about pleas. We're looking at theories about properties and truth-conditions of statements that ascribe such properties. If I'm to take observations about pleas as observations that favor some account of justification, I need some way of linking claims about pleas to claims about properties. One way to link claims about pleas to claims about properties is by means of the 'process-property integration thesis':

¹⁶ An anonymous referee raised this objection.

(PPI) Something is justifiable only if it has the property or properties that would in principle constitute a successful justification of it if cited.¹⁷

If something lacks certain properties, it cannot be defended from criticism or justified and cannot have the status of being justified. We defend things against criticism by citing the features of the thing that show the criticism to be unwarranted. So, on the proposal I am offering here, if we can identify the types of criticism to which something is subject and the conditions under which such criticism can be negated, we can in turn identify the conditions under which the thing is justified.

We can save Externalism from the charge that it is unintuitive by drawing a distinction between the conditions under which a *believer* can be justified and the conditions under which a *belief* can be justified.¹⁸ If such conditions need not coincide, according to (JD), it is possible for a believer to be justified when the belief held is not. To see why the standard view does not allow for such a possibility, assume the intuitions that support Parity are correct and that (1) is true. It follows from this and (SV1) that (2) is true. Whereas externalists will say that conditions beyond those that supervene on the subject's non-factive mental states have to be taken into consideration in determining whether someone's belief is justified, it follows from this and (SV1) that in determining whether the subject is justified in believing these same conditions need to be taken account of. As we know that only the conditions that supervene on a subject's non-factive mental states are needed to determine whether, say, Coop is justified in his beliefs, it would follow that only such conditions can determine whether Coop's beliefs are justified.

To bring out the tension between the standard view and (JD), defenders of Externalism will say that in order to defend someone's belief about *p* from criticism, we have to take account of *both* how things stand with the believer's mental states and in the believer's circumstances. For

¹⁷ Audi (1993: 305). Note that the view is *not* that by asserting that a belief has some property, the belief is justified if the audience accepts that the belief has such properties. Rather, the idea is that if the audience takes your assertion at face value, they will agree that the properties you assert that the belief has constitutes a justification for that belief.

¹⁸ Bach (1985) and Engel (1992) appeal to such a distinction to defend reliabilism. Bach thinks personal justification ascriptions reflect our sense of the goodness of the epistemic 'action' that led to a belief. The trouble with explaining the distinction in this way is that if epistemic actions call for justification, we ought to be able to distinguish between a person being justified in an epistemic action and an epistemic action being justified. If such actions do not call for justification, it is not clear how we can say that a person being justified is to be understood in terms of our evaluation of epistemic action. Engel suggests that the personal/doxastic justification distinction coincides with the distinction between reliable beliefs and excusably held beliefs and adopts a view similar to Goldman (1988). He does not argue that we need to draw the distinction to accommodate intuition or why we ought to identify excusably held belief with a person being justified in holding that belief.

while we sometimes criticize beliefs on the grounds that they are not supported by the evidence of the believer, we often criticize beliefs for misrepresenting how things are or for being arrived at as the result of a process that is not reliable. If a defense of the belief aspires to show that such criticism is unwarranted and a belief is justifiable only if defensible, a belief could only be justified on those occasions where ‘external’ criticism (i.e., criticism that does not appeal to conditions the subject can be faulted for failing to take account of) is unwarranted. However, we do not criticize *believers* on such external grounds. It is no criticism of Coop that he makes a mistake. It is only a criticism of Coop if we say that he is culpable or responsible for the mistake.

Here is the argument from (JD) against (SV1):

- (4) Something can be justified iff it can be defended [(JD)].
- (5) A subject, S, can be defended from criticism for having Φ 'd if she is properly excused for having Φ 'd.
- (6) A subject, S, can be excused for having Φ 'd even if her Φ -ing was all things considered wrong.
- (7) If S's Φ -ing is all things considered wrong, S's Φ -ing cannot be justified.
- (C) It is possible that S is justified in Φ -ing even if S's Φ -ing is not justified.

This conclusion follows from the premises above and it is not clear which premise someone should reject in order to save the standard view. Someone might reject (4), but that requires saying that you can give a defense of something without being able to justify it. What would this defense amount to? Someone might reject (5), but then it would seem that she would say that a subject could be open to criticism for that which she cannot be faulted. In turn, that suggests that the truth of ascriptions of personal justification would potentially depend upon external considerations. If (5) were false, Externalism would be true. If (6) is false, the very notion of excuse must be incoherent. What is to be excused is the wrongness of an action or attitude. The upshot of an excuse is that the wrongness of the action or attitude does not reflect badly on the agent who performed the action or held the attitude. (7) says that an attitude can only be justified if the attitude is not all things considered wrong. To deny this is to allow that conditions can make it such that a belief is wrongfully held but justifiably held. From the deliberative perspective, it seems that it would be very odd to judge that a belief is justified and then wonder whether it is wrong to hold the belief. What was the justification doing for the believer in deliberation if it was not, *inter alia*, establishing that the attitude was not wrongfully held? From an

external perspective, it seems that the very conditions I would use to show that a belief is wrongfully held are the conditions I would use to argue that the belief is not justified.

Perhaps those who defend the standard view will say that the incompatibility of that view and (JD) shows we ought to reject (JD). We shall look at two arguments for the standard view, but I think neither argument gives us good reason for adopting this view.

2.2 Ordinary Language

Someone might say that the appeal to some distinction such as the distinction between personal and doxastic justification is *ad hoc* unless there is evidence of such a distinction found in ordinary language. Suppose I am trying to determine whether to believe p . I consider the available evidence and think to myself:

(8) I am justified in believing p .

Then I ask myself, ‘But what of the belief?’ Sceptical, perhaps, I add:

(9) There is not sufficient justification for believing p .

If (SV1) is false, it is possible for (8) and (9) to be true. Yet, this is an ear-bruising assertion:

(10) I am justified in believing p , but there is no justification
for believing p .

The natural explanation as to why (10) strikes us as contradictory is that it is in fact a contradiction. If it is, then contrary to what I have suggested ‘S is justified in believing p ’ entails ‘S’s belief that p is justified’ and ‘There is sufficient justification for S to believe p ’.

There is no question that (10) seems contradictory. The question is whether (10) is a contradiction. I think that no one could truthfully assert (10). However, I think that what is asserted by (10) could be true. If we formulate it in the third-person, for example, I think there is nothing wrong with asserting:

(11) He is justified in his belief. However, there is no
justification for that belief.

(11) seems no different from:

(12) He believes p and I can defend him for believing that.
The belief, however, is one I cannot defend.

If you think (12) is coherent, note that arguments that would purport to show that (11) is incoherent would show that (12) is incoherent. If you think there is no real difference between a justification and a defense, (12) just says what (11) does.

If you still do not know what to make of the distinction between ascriptions of personal and doxastic justification, consider two exchanges. There has been a recent murder and Cooper has been called in to investigate. Over at the hotel, Audrey and Ben are arguing:

Audrey: Coop believes that Laura's killer is going to be found soon.

Ben: He is a fool. There is no defending a guy like that. He has no good leads, and he should know better than to think he is getting any closer to figuring out what really happened.

Audrey: You have no right to criticize Coop, Ben. If you were in his shoes and did not know how things worked around here, you would think you were getting closer to figuring out who killed Laura. Sure, maybe none of his leads will pan out, but they seem promising. He is surely justified in thinking that he is going to catch Laura's killer.

Meanwhile, Bobby and Donna are talking at the diner:

Bobby: Coop believes that Laura's killer is going to be found soon.

Donna: Is there any justification for thinking that?

Bobby: Not really. He thinks there is, but he has not figured out that this town is full of secrets. He will be chasing blind leads for a while now. Not that he is a fool, mind you. Anyone would think they were hot on the killer's trail given his leads.

I do not think we have to say that Audrey and Bobby disagree. Audrey explicitly believes that Coop is justified in his beliefs and Bobby seems not unwilling to accept that. However, if we can say Audrey and Bobby are not necessarily in disagreement, it is because we think they are making a distinction between an objective assessment of Coop's beliefs which culminates in a judgment about whether there is a justification of his belief and a more subjective assessment that culminates in a judgment about Coop that also uses the language of justification. If we think they need not disagree, we have to revise the standard view.

So, how could it be that (11) and (12) are true even if no one could truthfully assert (10)? When someone acts or believes, there are two kinds of criticism of the action or attitude. Some criticism implies fault. Some criticism implies wrong without implying fault. That is to say that the conditions under which someone can be faulted for having acted or believed are conditions that constitute a kind of wrong that must be absent if the attitude or action is justified. If someone

believes that there is no justification for believing p , they can be faulted for believing p . That fact alone is sufficient to ensure that they are not justified in believing p and there is not sufficient justification for their belief that p . However, if an outside observer sees that there is not sufficient justification for the subject's belief *and* sees that the believer is not in a position to appreciate this, she might defend the believer for having believed what she did while holding off from defending the belief. So, we can explain why (10) could never truthfully be asserted and why its assertion strikes us as contradictory without the assumption that (SV1) is true. Thus, ordinary language does not support (SV1) and the oddity of (10) constitutes no evidence for (SV1). In fact, the dialogue suggests that ordinary language allows for just the sort of distinction I have drawn between personal and doxastic justification.

2.3 The Logic of Justification Ascription

In their work on the logic of justification ascriptions, Kvanvig and Menzel provide a separate argument for the claim that statements that ascribe personal justification of the form 'S is justified in believing p ' entail statements that ascribe doxastic justification of the form 'S's belief that p is justified'.¹⁹ They argue that personal justification can be reduced to doxastic justification as follows. First, they offer us an account of the truth-conditions for ascriptions of personal and doxastic justification respectively as follows:

- (PJ) 'S is justified in believing p ' is true iff S has the property of being the x such that x 's believing of p is justified.
- (DJ) 'S's belief that p is justified' is true iff S's believing of p has the property of being justified.

Those who want us to heed the (alleged) distinction between saying that it is the person who is justified and saying that it is the belief the person holds that is justified will say that what one does in ascribing personal justification is say that the person is such that a certain kind of evaluation is fitting. They claim that (PJ) does justice to this insofar as (PJ) captures the idea that an ascription of personal justification is true in virtue of the person having a property (i.e., the property of being the person such that they have a justified belief). As for (DJ), that clearly captures the idea that an ascription of doxastic justification is true in virtue of the belief's having a certain normative property. Note that the right-hand side of (PJ) could only be true if two conditions are met. First, that someone's believing of p is justified. Second, that someone is the subject, S. If the right-hand side of (PJ) is true, it cannot be that the right-hand side of (DJ) is false as it clearly follows from the

¹⁹ Kvanvig and Menzel (1990). They are critical of Bach's (1985) suggestion that we can distinguish between ascriptions of personal and doxastic justification.

assumption that S is the individual such that S's believing of p is justified that S's believing of p is justified.

If this argument were cogent it would vindicate (SV1). However, this proposal does not capture the spirit of the suggestion that we ought to distinguish personal from doxastic justification. It fails to do justice to the idea that an ascription of personal justification is about the person in saying that what it is to ascribe justification to the person is to see them as being in possession of something that is justified. Moreover, if the argument supported (SV1) it would prove far too much. For the argument would show that for any property, F, such that F can be attributed to either the agent or the agent's attitude, there would be no exceptions to this general schema:

(GS) S is F in Φ -ing iff S's Φ -ing is F.

Suppose while rushing through a crowded market, I knock over a small child. I should have been paying better attention, so we say:

(13) I was responsible for knocking that child down.

It surely does not follow that:

(14) My knocking the child down was responsible.

Or, suppose a morally conscientious person is forced to decide between two competing *prima facie* duties and the kind of conflict we are considering is the kind of conflict whose proper resolution reasonable people can disagree about. It seems that an agent compelled to make a decision might fail to discern which of the two duties takes precedence without being culpable for failing to have acted as she ought to have, which might lead us to say:

(15) Although she made the wrong decision, she is beyond criticism.

(16) Although she made the wrong decision, her decision is beyond criticism.

Whereas (15) seems sensible, (16) is surely incorrect. If the subject's failure to do what she ought is a non-culpable failure, that it was non-culpable suggests that (15) is true. That it was a failure suggests that (16) is true. The argument for (SV1), which relies on (GS), if sound, would show that (16) is a consequence of (15). It seems we can use the language of justification to restate (15) and (16):

(17) Although she made the wrong decision, she was justified in having acted as she did.

(18) Although she made the wrong decision, her decision was justified.

If you think, as I do, that (17) is a restatement of (15) and (18) a restatement of (16), perhaps you will agree that (17) is acceptable, but (18) is not. There would be no question as to whether (17) entailed (18) if (GS) were true or (SV1) were true.

How are we to fix Kvanvig and Menzel's proposal? On the account I favor, to say that 'S is justified in believing p ' is to say that in light of the facts about the situation and the facts about S, S can be defended from criticism. To say that 'S's belief that p is justified' is to say that in light of the facts about the situation and the facts about S, S's belief can be defended from criticism. We criticize agents using hypological concepts so that our criticism of agents has to do with fault and responsibility. We criticize attitudes using deontological concepts so that our criticism of attitudes has to do with whether the belief conforms to epistemic norms or there exist permissions to hold the attitudes in question. If we assume that fault constitutes a kind of wrong while assuming that not all wrongs imply fault, it follows that the conditions that determine whether ascriptions of personal justification are true are a subset of the conditions that determine whether ascriptions of doxastic justification are true. That is to say, because the deontological (i.e., the facts about permissibility) does not supervene on the hypological (i.e., the facts about culpability, responsibility, blameworthiness, praise), the facts about doxastic justification do not supervene on the facts about personal justification.²⁰ As the denial of this supervenience thesis allows us to accept both Parity and Asymmetry, it seems we are moving closer to a view that accommodates internalist and externalist intuitions.

If we say that the notion of personal justification is to be explicated using hypological concepts and the notion of doxastic justification to be explicated using deontological concepts, the truth-conditions for ascriptions of personal and doxastic justification could be understood as follows:

- (PJ2) 'S is justified in believing p ' is true iff S believes p and S should not have expected that there would be an undefeated reason against believing p .

²⁰ The deontological cannot supervene on the hypological if it is possible for two agents to be equally responsible, blameless, excused, or what have you when one and only one manages to believe or do what there is overall reason to do, manage to do or believe what she ought, believe or do what is permissible, or only one believes or acts wrongly. So, if excusable wrongs are possible, the deontological cannot supervene on the hypological.

(DJ2) ‘S’s belief that p is justified’ is true iff S believes p and it is not the case that S should have believed otherwise than believe p .²¹

2.4 Justification, Excuse, and Permission

I have argued that the denial of the standard view is not incoherent and that we ought to revise the standard account of justification by rejecting (SV1). Even if it is not incoherent to deny (SV1), it does not follow that we ought to deny (SV1). One reason to deny it is that by denying (SV1) we can accommodate a wider range of intuitions than we would if we accepted it. Is there any additional reason for rejecting (SV1)? I think so. I shall argue that you ought to reject (SV1) if you want to make sense of the distinction between permissions and excuses. Regrettably, the argument is complicated. I shall argue as follows. We first focus on the cases that are taken to illustrate the distinction between permissible action and excusable wrongdoing. After distinguishing between two ways in which we might remove blame for having failed to do what one ought, we shall see that to capture the necessary distinctions between ways of removing blame and ways of denying wrongdoing we will want to draw on a distinction like the distinction between the subject being justified and that which the subject does or believes being justified. So,

²¹ To see what this amounts to, consider the view defended by Adler (2002), Bird (2007), and Sutton (2005) that knowledge is the norm of belief:

K: You ought not believe p unless you know p .

Combined with this account, we would say that ‘S is justified in believing p ’ is true iff S believes p and should not have expected that her belief would not constitute knowledge. This allows us to say that if her belief about p unexpectedly fails to constitute knowledge, she is not unjustified. However, if we say that to give a justification of a belief is, *inter alia*, to argue that it conforms to the norms of belief, we would say that ‘S’s belief that p is justified’ is true iff S’s belief that p constitutes knowledge. Suppose instead that you agree with Boghossian (2003) and Wedgwood (2007) that the norm of belief is the truth norm:

T: You ought not believe p unless you know p .

We would say that ‘S is justified in believing p ’ is true iff S believes p and should not have assumed that her belief about p would not be mistaken. If we think someone who believes without sufficient evidence ought not have assumed that they would not be mistaken, it follows from T that you cannot be justified in your beliefs without sufficient evidence. To say ‘S’s belief that p is justified’ would require us to say that if p is false and you believe it, you ought to believe otherwise. Of course, it does not follow that you are not justified in your belief. Now, could someone who adopted T explain why *more* than truth is necessary for having a justified belief? I do not see why not. Depending upon your views about fault and permissibility, if you think (as I do) that not all wrongs imply fault but if someone can be faulted for a belief or action the belief or action is thereby wrongful, it follows that an ascription of doxastic justification cannot be true unless an ascription of personal justification is true, too. So, if according to T, the ascription ‘S is justified in believing p ’ can only be true if S has positive evidence to support her belief about p , it follows that the ascription ‘S’s belief that p is justified’ can only be true if S has positive evidence to support her belief about p .

if you believe that the excuse/permission distinction is both coherent and non-empty and appreciate the difference between removing blame by offering an excuse and removing blame by offering an exemption, you should think of an excuse as sufficient for the personal justification without being sufficient for doxastic justification. Intuitively, a justified action is an action such that it is not the case that it ought to have been otherwise. Intuitively, a justified agent is an agent such that it is not the case that she ought to be blamed for that which she is responsible for. Applied to the case of belief, the beliefs that are justified are the ones such that it is not the case that they ought to be replaced by something else (e.g., belief in the negation or suspension of judgment). The believers that are justified are the ones who arrive at their beliefs in such a way that they ought not be criticized for holding the beliefs that they do.

Remember how this all started. We wanted to know whether we ought to say that Audrey and Cooper are equally justified in their beliefs and what significance does this have to the debate between the internalists and externalists. What I would like to suggest is that the reason that we can say (correctly) that Coop is justified in believing what he does is that Coop ought to be excused even though he ought to believe otherwise. The reason that we can say (correctly) that only Audrey's beliefs are justified is that it is not the case that Audrey ought to believe otherwise. Against this proposal, many will either insist that you cannot accommodate the intuition that Coop is justified in believing p by saying that he ought to be excused for believing what he does or insist that his being excused for believing p cannot be what suffices to establish that the ascription of personal justification 'Coop is justified in believing p ' is true.²² I shall try to show that these objections are unwarranted.

My proposal is that when Coop comes to mistakenly believe p on the basis of hallucinatory experience his belief is not justified, but *he* is justified because he ought to be excused for failing to hold a belief for which there is adequate justification, I will begin by looking at what internalists have said in anticipation of my response. Concerning cases of beliefs based on hallucinatory experiences, Audi offers these remarks:

²² I shall argue that if an excuse is warranted, it is true that the agent who is excused from criticism was personally justified in committing the wrong for which they are excused. I shall also argue that the cases of non-culpable error are cases in which excuses are appropriate. There will be some who insist that an ascription of personal justification cannot be true simply because the person ought to be excused. Rather than suggest that there is nothing more to the claim that 'S is justified in believing p ' than 'S's believing p is excusable', I could instead offer an error-theory and say that the reason people have mistakenly said that Parity is true and Coop is justified in his beliefs about the external world which are backed by hallucinatory experiences is that they harbor some misconceptions about when excuses are appropriate and when justifications are.

... given the vivid hallucination, I am in no way at fault for believing what I do, nor do I deserve criticism. 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.

As this passage indicates, Audi thinks that we ought to say that beliefs backed by hallucinatory experiences are justified because we think they are reasonable, think the believer is not to be faulted, and we approve of the belief. Audi believes that there is a difference between justified belief and excusable belief, and to illustrate the difference he asks us to imagine that someone has had a belief 'stamped on the brain', perhaps through a process of indoctrination or brainwashing, and says, "an induced belief which one can do nothing to remove no matter how hard one tries is not thereby justified; it is *excusable*, but not well-grounded".²³ Wedgwood also offers cases of temporary insanity as his example of excusable wrongdoing. If someone shoots and kills another in self-defense, he says, that is a justified killing. If someone shoots and kills another because he has gone temporarily insane, that killing is excusable.²⁴

Note that both authors take as their example of excusable wrongs examples in which an agent has lost control of her attitudes or actions or has had her capacity for responding to reasons seriously compromised so that in their examples of excusable wrongs, the agent's 'response' to the reasons there is nothing that would lead us to affirm the subject's normative competence. When someone's actions or attitudes are justified, however, we think that their actions or attitudes are justified because we think that they are reasonable. I think this is inadequate. To see why, let us focus on the practical examples. There is nothing wrong with saying that if someone kills in self-defense the killing is justified. It is not an abuse of language to say that someone who kills another because they have gone temporarily insane ought to be excused. However, there is a third case to consider. What are we to say about cases of 'imperfect' self-defense? That is, what are we to make of cases in which an agent kills someone believing mistakenly that self-defense justifies such a killing?²⁵

²³ Audi (1993: 28).

²⁴ Wedgwood (2002: 349).

²⁵ Moore (1997) refers to such cases as cases of imperfect self-defense. It is important that we assume in such cases that the subject is motivated by the belief that the killing would be justified by self-defense, knows what sorts of conditions would justify such a killing, is not culpable for

I think it is clear that they are not *inexcusable*. It is clear that if the subject is non-culpably mistaken about the facts of the case (e.g., mistaken in thinking that someone is coming after her with a weapon), she ought to be excused for the killing even if she ought to have acted otherwise. If we think that in cases of imperfect self-defense the killing is excused *and* the killing is justified, that we think of such cases as excusable wrongs forces us to deny:

(JO) If S's Φ -ing is justified, it is not the case that S ought to do other than Φ .

But, for reasons alluded to earlier, (JO) is not the sort of thing we ought to deny. If there were more to Φ -ing when one ought than Φ -ing being justified, you get the odd result that someone might know that Φ -ing would be justified without being in a position to judge that she may Φ or ought to Φ . And if we thought of deliberation as a matter of reflecting on the reasons to come to a judgment about whether Φ -ing is overall justified, denying (JO) forces us to say that deliberation does not so much as address questions about what one ought to do or believe. (Or, alternatively, deliberation is concerned with questions about what one ought to do or believe, but does not directly concern the justification of the actions or attitudes in question.) If the fact that there is sufficient justification for S to Φ were insufficient for putting S in a position to Φ permissibly, it would make sense for an external advisor to knowingly advise S by saying 'There is sufficient justification for you to Φ ' while conceding that she does not know whether S ought to do other than Φ .

Suppose we say that cases of imperfect self-defense are cases of excusable killing and that cases of self-defense are cases of justified killing. What of it? It seems that cases of imperfect self-defense are in some respects like cases of in which someone kills because they have gone temporarily insane and in other respects like cases of justifiable killing in self-defense. Like cases of temporary insanity, the agent who intends to kill in cases of imperfect-self defense has lost the right to non-interference and this reinforces the idea that they ought to act otherwise. Like cases of justifiable killing, however, we can only say that the agent in question ought not be blamed for the killing if we can say that the subject is normatively competent. We would not regard the killing in cases of imperfect self-defense as excusable unless we regarded the agent as reasonable.

It is here that Strawson's distinction between excuses and exemptions is useful.²⁶ He observed that there are three ways to remove blame. If S Φ 's and someone blames her for that,

believing falsely that the circumstances are such as to justify such a killing, and is not motivated by malice or anything but the desire for self-defense.

²⁶ Strawson (2003). Horder (2004) speaks of this distinction in terms of the distinction between denials of responsibility and excuses.

you can either show that blame is inappropriate because S ought to have Φ 'd (i.e., offer a justification for S's Φ -ing), show that blame is inappropriate because S ought not be thought of as a responsible agent (i.e., offer an exemption), or show that blame is inappropriate because while S is a responsible agent it is a mistake to take her having failed to Φ with justification as an indication that she is responsible for the failure (i.e., to offer an excuse in the strict sense). What I shall argue is that the standard view of justification ascriptions stands in the way of making sense of the distinction between excuses and exemptions and that to make sense of this scheme, we will want to do two things. We will want to say that personal justification is insufficient for justified action or belief. We will want to say that nothing more than an excuse is necessary for personal justification.

So, consider our three cases:

- (I) S kills S' because S has gone temporarily insane [Exemption].
- (II) S kills S' because S believes mistakenly, but non-culpably, that S' is trying to kill S and knows that if S' is trying to kill S, shooting S' is the only way to fend off the attack [Excuse].²⁷
- (III) S shoots and kills S' in self-defense [Justified Action].

Suppose a defender of Internalism proposes this scheme:

- (S1) S is justified in pulling the trigger in (II) and (III), but not (I).

If this first scheme is correct, according to (SV1), S's pulling the trigger and killing S' in (II) and (III) is justified. According to (JO), it is not the case that S ought to have done other than pulled the trigger in (II) and (III). But that contradicts the assumption that (II) is an instance of excusable wrongdoing.

To avoid this difficulty, someone might propose this scheme:

- (S2) S is justified in pulling the trigger in (III), but not in (II) or (I).

If we assume (JO), we get the correct result that it is not the case that S ought to have done other than shoot in (III). We avoid the incorrect result that S was permitted to pull the trigger in (II). However, now that we are using the concept of justification to distinguish (II) from (III), we

²⁷ Moore (1997), Robinson (1996), and Zimmerman (2004) would regard imperfect self-defense as an instance of excusable wrongdoing. For an argument that the propriety of such excuses depends on showing the agent's actions in a favorable moral light without thereby offering a crypto-justification, see Gardner (1998).

cannot use it in distinguishing (I) from (II). How could someone accept Internalism, (SV1), and (S2) explain the difference between (I) and (II)? How could they explain the difference between someone who ought to be excused for having killed and someone who ought to be exempted from criticism? It seems that the natural way to mark the difference is as follows. The agent who is to be excused is unlike the agent who is to be exempted insofar as the agent is rational and responsible. If we judge of such an agent that she failed to do what she ought and wish to remove blame, we have to think of her response to the demands she was under as being other than they ought to have been, but reasonable. If the shooter were led by *unreasonable* beliefs to shoot, we would regard her as blameworthy and her action inexcusable. If we say that an agent who is to be excused for having Φ 'd and cannot be exempted for having Φ 'd is such as to have been reasonable in having Φ 'd, (S2) forces us to reject the following claim:

(RJ) If S was reasonable in Φ 'ing, S was justified in Φ -ing.²⁸

Can internalists reject (RJ)? I think not. When we judge that someone was reasonable in having Φ 'd, we tend to focus on how things are from their perspective and ask whether someone with such a perspective on the situation should have known better than to assume that in Φ -ing they would not Φ against an undefeated reason. If we think they should have known better than to assume that, we regard them as unreasonable. If we think they should not have known better than to assume that, we regard them as reasonable. So, if the question as to whether S is reasonable is settled by S's non-factive mental states, but the fact that S is reasonable does not entail that S is justified, we have to assume that the facts about justification fail to supervene on the facts about what is reasonable for S to believe or do. And owing to Internalism about the facts about what is reasonable to do, it follows that Internalism fails as a constraint on what facts bear on what there is adequate justification for believing or doing.

Internalists cannot reject (RJ). Rejecting (JO) would force us to seriously revise our ordinary practice of normative appraisal. Internalists who insist that the standard view of justification ascriptions is correct cannot make sense of the distinctions we want to draw concerning the cases (I)-(III). That is to say, they cannot make sense of the distinctions between exemptions, excuses, and the justification of actions or attitudes. Revising (SV1), it becomes much easier to make sense of these distinctions. To say that someone is to be excused from criticism (and not exempted) is to describe the agent as being reasonable while remaining neutral on the further question as to whether they ought to have acted or believed in ways other than they

²⁸ Cohen (1984) claims that 'reasonable' and 'justified' are synonymous. While I think there is a sense in which the reasonable and the justified come together, there is another sense in which they do not.

actually did. If there is nothing more to show that the person is justified than to showing that the person can be defended from criticism, showing that they are reasonable and affirming their normative competence ought to suffice to show that they are justified. What allows us to maintain the distinction between (a) showing that the person is justified and ought to be excused and (b) showing that what the person did or believed is also justified requires us to reject (SV1). This allows us to say that an action is or an attitude justified only if it is not the case that it ought to have been different. This allows us to maintain the distinction between excusable wrongs and attitudes or actions that justified.

3. Conclusion

Let's take stock. We began by looking at an argument against Externalism. The new evil demon argument takes intuitions most of us share and tries to show that such intuitions are incompatible with Externalism. I have argued that there is no such incompatibility. The additional assumptions needed for a valid argument from Parity to Internalism are problematic. I then suggested that there were additional intuitions that supported Externalism. Either this means that our intuitions are an incoherent mess or that there is some way of making sense of the lot of them. I tried to make sense of them by means of the suggestion that the truth of ascriptions of personal justification of the form 'S is justified in believing p ' depends upon how things are with S's non-factive mental states while allowing that the truth of ascriptions of doxastic justification of the form 'S's belief that p is justified' can be affected by facts external to the subject's mind. This distinction proves useful not only in trying to make sense of our intuitive sense of which justification ascriptions are appropriate but also in trying to make sense of the more familiar distinction between right action and excusable wrongdoing. The major obstacle to this suggestion is that some doubt the coherence of the distinction, but that obstacle has been overcome.

The aim of this paper was not to bring the Internalism/Externalism debate to a close, but to show that an argument that has been taken to bring that debate to a close is far from compelling. The allegation that defenders of Externalism suffer from some sort of tin ear and cannot do justice to our ordinary sense of which justification ascriptions are appropriate is unwarranted. I do think that if externalists were so inclined, they could build an impressive case for their view by appeal to intuition and general theoretical considerations concerning the relationship between justification and other deontological concepts. That will have to wait for another time.²⁹

²⁹ I would like to thank Al Casullo, Nate Charlow, Mark Decker, John Gibbons, Danielle Hampton, Howard Hewitt, Jonathan Ichikawa, Jon Kvanvig, Sarah Sawyer, John Turri, and

REFERENCES

- Audi, R. 1993. *The Structure of Justification*. New York: Cambridge University Press.
- _____. 2002. An Internalist Theory of Normative Grounds. *Philosophical Topics* 1&2: 19-45.
- Bach, K. 1985. A Rationale for Reliabilism. *The Monist* 68: 246-63.
- Bird, A. 2007. Justified Judging. *Philosophy and Phenomenological Research* 74: 81-110.
- Boghossian, P. 2003. The Normativity of Content. *Philosophical Issues* 13: 31-45.
- Brewer, B. 1997. Foundations of Perceptual Knowledge. *American Philosophical Quarterly* 34: 41-55.
- _____. 1999. *Perception and Reason*. New York: Oxford University Press.
- Cohen, S. 1984. Justification and Truth. *Philosophical Studies* 46: 279-95.
- _____. 1988. How to Be a Fallibilist. *Philosophical Perspectives* 2: 91-123.
- Conee, E. and R. Feldman. 2004. *Evidentialism*. New York: Oxford University Press.
- Cruz, J. and J. Pollock. Forthcoming. The Chimerical Appeal of Epistemic Externalism.
- Darley, J. and P. Robinson. 1998. Testing Competing Theories of Justification. *North Carolina Law Review* 76: 1095-1143.
- Engel, M. 1992. Personal and Doxastic Justification in Epistemology. *Philosophical Studies* 67: 133-50.
- Gardner, J. 1998. The Gist of Excuses. *Buffalo Criminal Law Review* 1: 575-98.
- Goldman, A. 1979. "What is Justified Belief?" In *Knowledge and Justification*, ed. G. Pappas. Boston: D. Reidel.
- _____. 1986. *Epistemology and Cognition* (Cambridge, MA: Harvard University Press).
- _____. 1988. Strong and Weak Justification. *Philosophical Perspectives* 2: 51-69.
- Graham, P. Forthcoming. "Theorizing Justification." In *Contemporary Topics in Philosophy 5: Knowledge and Skepticism*, ed. M. O'Rourke, J. Campbell and H. Silverstein. Cambridge, MA: MIT Press.
- Hawthorne, J. and J. Stanley. Forthcoming. Knowledge and Action. *The Journal of Philosophy*.
- Horder, J. 2004. *Excusing Crime*. New York: Oxford University Press.
- Kvanvig, J. and C. Menzel. 1990. The Basic Notion of Justification. *Philosophical Studies* 59: 235-61.
- Littlejohn, C. Forthcoming. From E = K to Scepticism? *The Philosophical Quarterly*.

Brian Weatherson for discussing the issues addressed in this paper. Special thanks to Robert Audi, Matt Burstein, Doug Ehring, Robert Howell, Leo Iacono, Dan Korman, Jonathan Sutton, Brad Thomson, Mark van Roojen, Charles Wallis, and two anonymous referees for incredibly helpful written feedback on various drafts of this paper.

- _____. MSa. Evidence and Armchair Access.
- _____. MSb. Moore's Paradox and the Norms of Belief.
- Moore, M. 1997. *Placing Blame*. New York: Oxford University Press.
- Nelson, M. 2002. What Justification Could Not Be. *International Journal of Philosophical Studies* 10: 265-81.
- Robinson, P. 1996. "Competing Theories of Justification: Deeds vs. Reasons." In *Harm and Culpability*, ed. A. Smith and A. Simester. New York: Oxford University Press.
- Silins, N. 2005. Deception and Evidence. *Philosophical Perspectives* 15: 375-404.
- Sosa, E. 1985. The Coherence of Virtue and the Virtue of Coherence. *Synthese* 64: 3-28.
- Strawson, P. 2003. "Freedom and Resentment". In *Free Will*, ed. G. Watson. New York: Oxford University Press.
- Sutton, J. 2005. Stick to What You Know. *Nous* 39: 359-96.
- Unger, P. 1975. *Ignorance*. New York: Oxford University Press.
- Wedgwood, R. 2002. Internalism Explained. *Philosophy and Phenomenological Research* 65: 349-69.
- _____. Forthcoming. "Normativism Defended." In *Contemporary Debates in the Philosophy of Mind*, ed. B. McLaughlin and J. Cohen. Malden, MA: Blackwell.
- Williamson, T. 2000. *Knowledge and its Limits*. New York: Oxford University Press.
- Zimmerman, M. 2004. Another Plea for Excuses. *American Philosophical Quarterly* 41: 259-66.