

Externalism Explained

Clayton Littlejohn
Dianoia Institute of Philosophy
Australian Catholic University

Forthcoming in Luis R. G. Oliveira (ed.), *Externalism about Knowledge*. Oxford University Press.

Abstract

This is a defence of externalism about knowledge and also about justification. In this paper, I argue that an important virtue of externalism about justification is that it helps to explain the value of (i.e., importance of) knowledge. I also develop and expand upon the line of argument that I take to be the most promising argument for externalism, one that appeals to morally loaded cases. We need externalism to understand how it could both be true that we're subject to certain widely accepted normative requirements (e.g., to conform to objective standards that apply to all rational agents, to respect the rights of others, etc.) and that we should be guided by our beliefs. Internalist views (along with many orthodox externalist views) imply that these normative requirements don't apply to all rational agents or that rational agents shouldn't be guided by their beliefs even when they believe precisely those things that they ought to believe. Since internalists have given us no good reason to revise our ordinary understanding of the obligations that we might be under and no reason to think that we shouldn't be guided by our beliefs, we have very good reason for rejecting their picture of what we ought to believe.

Introduction

In this chapter, I shall defend externalism about knowledge. It differs from some more familiar externalist proposals because of its deontological character. The view is deontological insofar as it characterizes justification in terms of norms, which I understand as things that specify the conditions under which we should or shouldn't believe something. The view is externalist because these norms have external application conditions (i.e. these norms can sometimes say that we should or shouldn't believe because some external condition obtains or fails to obtain). The fundamental epistemic norm, on my view, tells us that what we should or shouldn't believe depends upon what we are in a position to know. Just as ignorance (according to some) enjoins silence (i.e. we shouldn't assert p if we're not in a position to know that p), it enjoins suspension (i.e. we shouldn't believe p if we're not in a position to know that p). The main focus of this chapter is externalism about knowledge but debates about externalism about knowledge and about justification are connected. They have both been shaped by assumptions about the normative insignificance of knowledge (and external conditions generally) that I believe we should reject.

What is externalism about knowledge? If it were the view that knowledge supervenes upon what's 'in the head' or upon the thinker's non-factive mental states, we wouldn't have much to discuss. Given the factivity of knowledge and some mild anti-sceptical assumptions, it is clear that thinkers can differ in what they know without any interesting internal differences between them. Following Conee, I think we should think of the debate between internalists and externalists about knowledge in different terms:

Externalist accounts of knowledge are those that reach farther into the external world than is needed to provide for Gettier prevention and truth. Some externalist accounts of knowledge replace the justification condition with something external, or count something external as justification. (2004: 79)

This is how I'll understand the nature of the disagreement between internalists and externalists about knowledge. It's not a debate about the supervenience of knowledge on what's in the head or the thinker's non-factive mental states. We'll assume that externalism about knowledge is true iff knowledge doesn't require justification or it requires justification but justification doesn't supervene upon our internal states.

It is possible to have propositional knowledge without justification in the sense that Agnes might know something even if she doesn't justifiably believe it (Kornblith 1999; Littlejohn 2018; Sylvan 2018). For a creature's responses to be justified, they have to be the kind of creature that can have responsibilities, duties, and obligations. I don't think that owls, spiders, or squirrels have duties or responsibilities (e.g. a duty to refrain from entering your home, the responsibility to care for the young). If one comes in under a door or through a window, we don't need to argue about whether this creature's decision was justified. A creature's response to their situation is only justified if their responses can call for a defence and be defended. When a creature cannot be held accountable for the way she responds to her situation, the creature has no responsibilities that they might have failed to meet and so there's no interesting sense in which their responses could be right, consistent with their obligations or duties, or justified.

Responses are justified only when the responding creature can be held accountable but have violated no norms. Norms don't apply to tables or spiders, so even if they do something, their doings don't have or lack justification. At some point, Agnes apparently decided to curl up on the couch to take a nap. If this is her flat and her couch and she's an adult, her decision might have been justified. If she's in the neighbour's flat having come in uninvited through an open window, it might not have been justified. If she's a cat or a dog, however, it doesn't seem to matter which flat she lives in. It doesn't seem that her decision could have been justified or unjustified because she has no responsibilities.

If this is right and the only responses that have the potential to be justified are the responses of creatures that can be held accountable, the internalist view which says that propositional knowledge requires justification commits us to the view that only creatures that can be held responsible for their attitudes have propositional knowledge. In other words, either dogs, cats, spiders, owls, and so on can be held responsible for their responses or they lack propositional knowledge. This, however, seems like a mistake. If a creature can be rationally guided by a fact, this fact could be this creature's reason for Φ -ing. In turn, it seems that when a creature's reason for Φ -ing is a fact, this creature knows the fact obtains (Hyman 2015; Unger 1975). For example, it seems that

you could surmise from a story I tell you that Agnes's reason for climbing through the window was that there was unguarded food left on the table and think this is a reasonable take on what happened even if you can't quite remember whether Agnes is a dog, a child, or a drunk relative. (Having the ability to be guided by facts seems to cut across the divide between those that can be held accountable and those that cannot.) Once we've agreed that Agnes's reason for climbing through the window was that there was some unguarded food left on the table, it seems odd to deny that she knew that there was unguarded food there to be found.

That settles that debate. (The externalists win again!) Readers will hopefully read on. I'll offer a second argument for externalism about knowledge, one that will hopefully shed light on the nature of the connection between knowledge and justification and expose the mistaken assumptions that shape the debate between internalists and externalists about knowledge.

If it's possible for non-human animals to have knowledge without having justified beliefs, this opens up an interesting theoretical possibility that hasn't been sufficiently explored. Think about consequentialist theories of right action. Proper consequentialists think that the good is prior to the right. To fully flesh out the details of their view, they'll need to identify something that has a kind of non-moral value that plays the right-making role. Truth doesn't have a normative dimension or component, but many epistemologists think that it has a kind of normative upshot (e.g. that a belief's being false gives us a reason not to hold it, makes it the case that we shouldn't hold it). It's the kind of thing that could potentially be prior to epistemic rightness and so the kind of thing that might play a right-making role in our epistemological theories. Knowledge, by contrast, was widely assumed to have a normative dimension or component. Knowledge or relations to knowledge couldn't play the right-making role in our epistemological theories if knowledge involved justification because then it would have rightness built in. Knowledge couldn't be prior to epistemic rightness if epistemic rightness were an aspect of it as it is on a familiar view accepted by the internalists. But what if knowledge doesn't have a normative dimension or aspect? What if it's just a relation between an animal and a fact that obtains when that fact can rationally guide the animal's responses, say? What if it were more like seeing an object, a kind of awareness relation that holds between an animal and an external object of awareness? If knowledge didn't have a normative aspect, it would be the kind of thing that would be at least eligible for playing a role in our theory of epistemic right-making. This, to my mind, is an exciting theoretical possibility. We should explore it further. If I'm right, knowledge is to justification what being optimistic is to right action on the utilitarian theory.

1. On the Aim of Belief

According to the veritist, belief aims at nothing but the truth (Joyce 1998; Wedgwood 2002; Whiting 2012). The gnostic disagrees. The gnostic thinks that

belief aims at knowledge (Gibbons 2013; Littlejohn 2013; Sutton 2005; Williamson 2000). How should we decide between these two views?

We might start by replacing this metaphorical talk of aims. We might, for example, replace this talk of aims with talk of norms. The disagreement between the veritist and the gnostic can be seen as a disagreement about whether the fundamental norm of belief tells us that we shouldn't believe falsehoods (NO FALSEHOODS) or that it tells us that we shouldn't believe what we don't know (NO IGNORANCE). We also might replace this talk of aims with talk of functions (Simion 2019). Think about what beliefs are supposed to do. Let's say that a belief is non-defective iff it can do what it's supposed to do and defective otherwise. We can think of this disagreement between the veritist and the gnostic as a disagreement about belief's function and about the difference between defective and non-defective belief.

My preference would be to start with the functions. The norms can then be seen as telling us that we shouldn't form or hold beliefs that are defective, incapable of doing what beliefs are supposed to do. I'll assume that the veritist wouldn't want to concede that the gnostic has a better account of the difference between defective and non-defective belief and that the gnostic wouldn't want to concede that the function of belief is properly described by the veritist's theory.

I've argued elsewhere that if we want to understand belief's functional role, we should think about the ways in which belief provides a basis consisting of reasons for the various doxastic, affective, and practical responses that our beliefs rationalize. While many epistemologists think that a belief's normative properties will be determined by what the belief is based on, we should also recognize that the belief's normative properties also depend upon whether it can provide an adequate basis for downstream responses (Littlejohn 2012). Think about closure principles. If you shouldn't believe p 's obvious consequences, that's a sign that you shouldn't believe p . In having a justification to believe p , this positive standing should radiate outwards towards the things that you know follow from p . We can expand upon this. Suppose p is something that you reasonably hope isn't so (e.g. a certain someone wins re-election, you have a virus). Suppose that you shouldn't be upset that p . Given that p is the kind of thing that reasonable people would find upsetting, this is a sign that you shouldn't believe p . If you should believe p , perhaps you should be upset that p . Maybe it's being proper to be upset that p is a matter both of believing what you ought to believe about p and feeling what you ought to feel about p . Our beliefs (the attitude, not the object) provide a basis (the belief's object, not the attitude) by providing us with propositionally specified entries to various forms of reasoning in which we treat the belief's object as if it's a reason. When you believe what you ought to believe, this removes one of the ways that the downstream responses that belief rationalizes could be unjustified or improper. Belief's reason-providing role should help us figure out the difference between defective and non-defective belief. The defective beliefs don't provide us with reasons. The non-defective beliefs provide us with reasons.

The veritist might think that this way of framing the issue shows the strength of their position. They might think that this view about reasons and function supports their view that true beliefs are non-defective. Think about the familiar idea that the value of true belief can be explained in terms of its practical benefits, the benefits of guiding us to Larissa or, more likely, the kitchen. Even if we think that it's better to be guided by knowledge than mere true belief, say, when our action plans are complex (Gibbons 2001; Williamson 2000), they might insist that true belief is good enough for action. A true belief will, so long as it persists, guide us to where we want to go. If they're good enough for this role, isn't that enough to show that true beliefs are non-defective?

There are a number of problems with this argument. Arguably, these considerations about practical benefits don't support the inference from false to defective belief because empirically adequate beliefs might also be fit for this practical purpose (Kvanvig 2003) and because false representations might nevertheless provide us with the map we need (Nolfi 2021).

The problem that worries me most is that we don't need beliefs or the agent's reasons to move the agent to act. The pragmatic approach isn't helpful for thinking about the difference between defective and non-defective belief because it picks out a functional role (i.e. that of combining with our desires, wants, intentions to guide behaviour) that isn't unique to belief. Credences that aren't beliefs also play this role, so beliefs turn out to be dispensable on this picture. If you're sufficiently averse to rain, even if you think that it won't rain, you might take an umbrella just in case.

Thinking about the different functional roles of belief and credence can help us better understand belief's functional role. In turn, it shows that we shouldn't test competing theories of non-defective belief by thinking about states of mind that are good enough for the purpose of moving us towards some desired goal. What we should consider are functional roles that are distinctive of belief, roles that beliefs can play and credences cannot. Recall Adler's observation about reactive attitudes:

[O]ur reactive attitudes normally do not admit of epistemic qualification. There are all manner of differences between, say, mild and strong resentment ... But there are no attitudes corresponding to a compromise between a negative reactive attitude toward someone (e.g., disgust, revulsion) and a weak epistemic position (e.g., I am slightly more confident than not that ...) to judge whether the claims of that attitude are met. Mild resentment is never resentment caused by what one judges to be a serious offense directed toward oneself tempered by one's degree of uncertainty. (2002: 217)

You cannot be rightly upset with your landlord if you suspect that he might increase the rent again if you're just as confident that he won't. Increasing your confidence to being more confident than not doesn't change things. Being highly

confident on the basis of purely statistical evidence also doesn't change things. To be upset that your landlord has decided to increase your rent again, you have to be convinced that this has happened. To be upset with your landlord for raising the rent, you have to believe that he did that. Believing this is being convinced of it and it is when you're convinced that you're able to be upset for the reason that your rent is being increased again. Belief's job isn't just to move your bodily mass from point to point but to provide you with propositionally specified reasons that can constitute the basis of your decision, constitute your reason for being upset, and convince you that you need to do something about this situation.

Suppose belief's distinctive function is to provide us with propositionally specified reasons that can figure in our reasoning and thereby potentially constitute one of our reasons for feeling, believing, or acting as we do. We can say that a belief is defective if it couldn't contribute a reason to reasoning that could be our reason for ϕ -ing and non-defective otherwise. Which account gives us a better account of when a belief can play this role, veritism or gnosticism? When can an individual's belief (the object, not the state) constitute her reason for ϕ -ing? The veritist and gnostic can agree that A's reason for ϕ -ing can only be that p if A believes p and can agree that it can be that p only if p. The belief and truth conditions are common ground. The differences between these views are these. The veritist thinks that beliefs provide reasons iff they're true, so the truth of a belief is enough and knowledge is not needed. The gnostic thinks that they provide reasons iff they constitute knowledge, so the differences between mere true belief and knowledge mark the difference between those beliefs that provide reasons and those that cannot. When viewed from this perspective, it seems to me that the gnostic view is clearly preferable. When A's reason for ϕ -ing is that p, the fact that p explains A's ϕ -ing and so a consequence of A's ϕ -ing for the reason that p is that A ϕ 'd because p. This, in turn, requires that the connection between A's ϕ -ing and the fact that p is not wholly accidental or coincidental, a fact that's sufficient to refute veritism. There's nothing that true belief brings to the table that ensures that the connection between belief and fact is anything more than accident or coincidence.

Agnes wanted to cheat on her maths exam, so she brought her phone with her and used a maths app that ominously promised to give students the answers they deserved. She was shocked when she learned later that she managed to get but one answer right out of 100. It turns out that the answers were generated at random, so there was no guarantee that they'd get things right and no guarantee that they'd get things wrong. The teacher didn't do her the favour of identifying which answer was correct, so Agnes is now as she was before, completely ignorant as to the answers to the questions on the exam. She believed that the answers were correct as she entered them and in one out of 100 cases her belief was true, but in none of these cases do I think we'd want to cite a mathematical fact as the explanans in explaining why she answered as she did. Contrast this with the case where Agnes finds a working calculator and develops a grasp of mathematical concepts and we might say that (assuming that

she knows that the sum was 36) that the answer was “even” because it was divisible by 2.

If the ability to be guided by the facts that we seem to have in mind when we believe or when we’re convinced is something that we place any value in, it seems that knowledge has a value that’s not contained in true belief that seems to be connected to a distinctive function that belief and belief alone can play. Cases like the one above suggest that when the kinds of accidental connection between belief and fact that prevent a thinker from having knowledge are cases that prevent the relevant fact from constituting one of the agent’s reasons. Even when we expect to find defences of veritism, we sometimes accidentally find better defences of gnosticism. When Lynch, for example, tries to convince us that we value true belief and do so for reasons beyond those that have to do with guiding our actions, he rightly directs our attention to the fact that there’s something that’s missing from life in the experience machine in spite of all the pleasant experiences it forces upon us:

We don’t want to live in the vat, even though doing so would make no difference to what we experience or believe. This suggests that we have a basic preference for truth ... We can put this by saying that I want my beliefs and reality to be a certain way—I want my beliefs to track reality, to “accord with how the world actually is”—which is to say I want them to be true. (2005: 18)

Focus on this last idea. If you want beliefs that track reality, you want true beliefs. This desire for having beliefs that track reality is not satisfied just by having true beliefs, however. Having or tracking the truth is not just a matter of having true beliefs. You cannot have hold of the truth or track reality in Nozick’s experience machine, but we can surely have true beliefs whilst trapped in this machine, whilst cut off from reality, and whilst completely and systematically disconnected from the world outside the mind.

Some readers undoubtedly think that there’s room for something in between veritism and gnosticism. It’s one thing to say that an app that spits out numbers at random cannot put Agnes in touch with the facts or that we cannot track reality whilst trapped in Nozick’s experience machine and quite another to say that we need knowledge to close the gap. It’s at just this point that people with tastes for complicated views like to trot out fake barn cases or cases of so-called environmental luck. They’re supposed to show that we can have the kind of contact with reality that’s lacking from the cases above even if we’re not in a position to know. I don’t find these cases persuasive.

Here’s one diagnosis of what’s going on in such cases. When you see, say, a barn under good viewing conditions and it has that familiar look, the building has a look that’s not distinctive of barns. Non-barns (if they are ringers or convincing fakes) will have the same look and so will trigger the same classificatory dispositions as a real barn. Getting it right under conditions where there are easily encountered ringers will happen pretty much at random much in the way that Agnes’s app delivered the right answer to maths questions pretty

much at random. In neither case do we have a process that classifies things correctly on the basis of the characteristics that trigger the classificatory dispositions, so I don't see any real reason to think that the correctness of the classification of the barn as a barn shows that the subject is in touch with reality in the sense that matters to us. Yes, they see a barn. No, they don't see that the building is a barn. We're interested in contact with facts (i.e. that something is a barn, not a fake), not objects (barns, fakes, owls, etc.), and I suspect that confusions about the objects of perception and knowledge is largely responsible for the muddle that people get into when they think that the beliefs we form in environmental luck cases 'track reality' in spite of failing to be knowledge. Remind yourself that given your dispositions and the distribution of ringers in your environment that you'll classify both barns and fakes at random and it's hard to say that you are attuned to the presence of the property being a barn.

If belief's function is to provide us with reasons that can serve as our reasons for feeling things, believing things, and doing things and reasons are the facts that we normally take them to be, there's a case for thinking that the distinction between knowledge and ignorance corresponds to the distinction between defective and non-defective belief. If the norms that govern belief tell us that what we should believe or should not believe also corresponds to the difference between defective and non-defective belief, we can see why knowledge is the fundamental norm of belief. For these reasons, veritism doesn't give us a plausible account of the function or norm of belief.

2. On Justification

Let's suppose that knowledge is the norm of belief. What does this tell us about justification? I think it tells us that justified beliefs have to be knowledge. This follows if we think that justifications function to identify responses that don't violate the norms that govern these responses:

CONFORMITY: Your belief about p is justified iff
(and because) it doesn't violate any epistemic norms.

This idea is close to one that you'd find in, say, Gardner's (2007) work on the distinction between justification and excuse so I'm surprised when I hear that some people are puzzled by this proposal. When suspected of having violated a law, say, we can try to show that there was no violation (deny the offence) or we can show that there was sufficient reason to violate the norm in question (so that it was wrong but not wrong all things considered). Either defence, if successful, should show that the agent's action was justified. Failing that, we might try to remove responsibility in some other way (e.g. by offering an excuse or an exemption) but what's essential to a justification is that it establishes that the agent shouldn't have done things differently and norms determine when we need to do things differently.

Transposed to the case of belief, the idea is that there are some norms that we shouldn't violate. If you've violated none of them, there's nothing that could threaten justificatory status. If you've violated one of them and there's no

overriding reason that would justify doing so, you lack justification. That's it. That's the theory.

I know that some people think that there must be something more to justification than this. They might think that if, say, we form beliefs irresponsibly or against the evidence we couldn't end up with a justified belief. The thing to say in response is not that such things couldn't matter to justification, but to remember that such things do matter to justification if they're built into the norms. If you thought that the only norms that governed belief had to do with truth or accuracy, this would be a pressing problem for you as you'd have to say that justification required something more or something different to norm conformity, but this isn't my problem because my norms tell us that anything that matters to knowledge matters to justification. If you think that we cannot justifiably believe something by believing against the evidence, my account explains this in terms of a violation of NO IGNORANCE. This kind of objection to CONFORMITY reveals bad background assumptions about the identity of the norms that govern belief, not some insight into the nature of justification (Littlejohn 2013).

Most epistemologists agree that there can be false, justified beliefs, so they reject this:

INFALLIBILISM: if you justifiably believe p , p .

Since INFALLIBILISM is an obvious consequence of NO IGNORANCE and CONFORMITY, they need to reject one or both of these. The epistemologists who reject INFALLIBILISM aren't as clear as I'd wish they'd be on whether they reject all truth-requiring norms or reject the idea that norms matter to justification, but it seems that their theoretical options would be these. First, they might say that whilst it's bad in some way to believe what we don't know, it's not quite right that we shouldn't form these (bad) beliefs (Kelp 2016; Simion 2019). (They can thus remain neutral on conformity and so allow that all justified beliefs conform to epistemic norms, but insist that no norm is violated when we believe falsehoods. All that happens is that we form a bad belief, fail to fulfil an aim or goal, etc.) Second, they might say that while NO IGNORANCE might be true, it tells us what we objectively shouldn't believe. The theory of justification, they'll insist, is concerned with subjective normativity only (Huemer 2001). Third, they might say that justification doesn't require meeting the conformity conditions on doxastic norms. We can justifiably believe what we shouldn't believe (Bird 2007; Ichikawa 2014; Reynolds 2013). What unifies this group is this idea that some or all of the objective conditions that have to obtain for someone to know couldn't be necessary for justification. They think that anyone interested in genuinely normative questions would only be interested in some non-factive notion of justification.

I want to push back against this idea that truth isn't needed for justification. The factive conception of justification we arrive at when we combine NO IGNORANCE with CONFORMITY speaks to important normative issues. My argument starts with this simple idea. We should be guided by our beliefs. This seems like a platitude. Let's unpack it. If we believe

something, we should be guided by our beliefs. The belief's content should figure in reasoning. I don't think that this is true because we ought to treat the contents of our beliefs as a premise in reasoning, but because we ought to be able to treat the contents of our beliefs as a premise in reasoning provided that we believe what we ought to believe. What could the point of a theory of justification be if it's not a theory of when it's appropriate to form or have premise-providing attitudes and when it's not proper to form or retain them?

We can say that there's a kind of wide-scope norm that we're all under some normative pressure to conform to. We ought to see to it that we either rely on our beliefs in reasoning or remove these beliefs from our set of premise-providing attitudes. Let's focus on a special case, the case in which you believe that you ought to ϕ . We can try to capture the content of this platitude about beliefs as follows:

GUIDANCE: You ought to see to it that you don't both believe that you ought to ϕ and respond in an alternative way.

The belief that you ought to head left, for example, is one that either you may hold and follow or you shouldn't hold at all. If you try to head to the right whilst believing that you ought to go left, you violate GUIDANCE. In general, people seem to agree that you're not supposed to do that.

Suppose that CONFORMITY is correct and you justifiably ϕ iff you permissibly ϕ . Notice that GUIDANCE has this surprising implication. If you wonder whether you ought to ϕ , you know that you'll be in one of the following situations: you ought to ϕ , you ought not ϕ , or you are permitted to ϕ and permitted not to ϕ . Now, suppose that you're in a situation in which ϕ -ing is prohibited but believe you should ϕ . Could you also be in a situation in which you justifiably believe that you ought to ϕ ? No, not if GUIDANCE is correct. If such a situation existed, you could justifiably believe that you ought to ϕ even though you'd be required to respond in some alternative way, but then GUIDANCE would fail. If GUIDANCE is correct, so is this:

NEAR NORMATIVE INFALLIBILISM: If you justifiably believe that you ought to ϕ , you may ϕ .

How could this be true?

I've argued that INFALLIBILISM provides the best explanation of NEAR NORMATIVE INFALLIBILISM (Littlejohn 2012; 2014b). The cases now known as 'morally loaded' cases (Williamson 2019) provide a clear rationale for externalism about justification and for thinking that the externalist conception of justification plays a central and important normative role. We can see this by examining views that uphold GUIDANCE but insist upon some non-factive notion of justification. Such views have to treat normative beliefs as a strange exception to the general rule that a justified belief might be false. As I shall argue below, such views are both theoretically strange and fail to vindicate important observations about the distinction between justification and excuse and about the normative irrelevance of normative uncertainty or misleading

normative evidence. We need INFALLIBILISM to understand why we should be guided by our justified beliefs.

Not everyone agrees that guidance is true, but let's focus on those views that try to accommodate GUIDANCE in some form or other. If I might speak loosely and crudely, if you reject INFALLIBILISM but accept GUIDANCE (and so accept NEAR NORMATIVE INFALLIBILISM), you need to explain why some normative beliefs cannot be justified if false given that your view predicts that so many of our beliefs might be justified but mistaken. One idea is that the difference has to do with the kind of evidence we have for normative beliefs. According to Smithies (2015) and Titelbaum (2015), we have maximally strong and undefeated epistemic support for believing certain a priori truths about certain normative matters. In other words, such truths are shiny. They're given the same treatment as, say, logical truths in standard Bayesian accounts of rationality. According to a different way of thinking about these things, the exception is better understood by thinking about the kinds of truths we're talking about. The normative truths are normative, so they have to be the kinds of things that can guide a rational agent (Fox 2019; Gibbons 2009; Kiesewetter 2017; Lord 2018; Oliveira 2018). They couldn't do this if, say, it could be that Agnes ought to ϕ when her evidence provided strong support for thinking that she ought not ϕ or if she were the 'same on the inside' as someone required not to ϕ . Rather than say that the normative truth shines so brightly that no rational thinker could miss it or be mistaken about it, we could say that the rational thinker's epistemic position with respect to the normative proposition helps to determine its truth-value. The normative requirements might be malleable, determined by the epistemic situation of the thinker in question. It might be that Agnes ought to ϕ because her evidence provides sufficiently strong support for the proposition that she ought to ϕ . To make this concrete, we might imagine situations in which Agnes either has very strong evidence for thinking that death can be bad for sentient beings by virtue of being a deprivation or has very strong evidence for thinking that only sentient beings who can think of themselves as creatures that exist over time can be harmed by being killed. Under one scenario, she might be rationally required to believe that hedonism is true and that it is false under the other scenario. Given some assumptions about cows, the shifty view might predict that Agnes would have different obligations in these two scenarios concerning cows (e.g. being required to give greater weight to their deaths in one scenario as opposed to the other). This would be quite different to how the INFALLIBILIST thinks about such cases. Suppose that hedonism is true and cows are sentient beings who don't have preferences about their futures. Under this scenario, it might be impossible for Agnes to justifiably slaughter cows as a way to make money, in which case she couldn't justifiably believe that she ought to slaughter cows as a way to make money regardless of which evidence she happened to have.

Two kinds of case cause trouble for views that combine GUIDANCE with some non-factive view of justification. The first concerns factual ignorance and mistake. The second concerns normative ignorance and mistake.

2.1 The Justification–Excuse Distinction

Let's begin with cases of defence of other and imperfect defence. In the good case, the agent knows that a potential victim is faced by a threat, knows that the potential victim cannot flee or escape the situation, and so knows that force must be used if the potential victim is to be saved from the threat (e.g. a mugger, a loan shark). In the bad case, the agent mistakenly believes that this potential victim is faced by a threat and knows that she cannot escape or flee the situation if this person is indeed a threat. Given the commitments of standard non-factive accounts of justification, it should be possible that she justifiably believes that force would be needed to defend the potential victim. We hold fixed the conditions that the non-factive view deems to be necessary and sufficient for justification when it comes to all the non-normative beliefs and the relevant parties all share the general normative knowledge about necessity, proportionality, and the like. We can zero in on mistaken factual beliefs or factual ignorance to think through the case.

Here's a natural (if not terribly complicated picture) of rights. We each have the right against others that they don't harm or kill us when we pose no threat to anyone else and when our being harmed or killed isn't an unavoidable side-effect of a course of action that the agent might pursue for the greater good. Rights come with correlative duties. If you have the right not to be killed or badly injured, the people you interact with have the duty not to kill or harm you. On this picture, your rights supervene upon factors that are distinct from (and not contained in) the factors that determine whether the people you interact with can (given the resources of any non-factive view of justification) justifiably believe that you pose an imminent threat to some potential victim that they know can only be protected by harming or killing you. On this picture, the duties that an agent has won't supervene upon the factors that determine whether her beliefs are justified. Thus, whilst defenders of such accounts could claim that we each have justification to believe things generally about the rights we have, the duties that they entail, the conditions under which they can be infringed, and so on, the justificatory status of non-normative connecting beliefs are a potential problem for anyone who wants to defend GUIDANCE. In particular cases, the combination of a priori normative knowledge and mistaken but (putatively) justified non-normative connecting beliefs can lead the agent to (putatively) justifiably judge that they ought to do something that, *inter alia*, violates the rights of another person.

It seems that if there can be false but justified beliefs about the nature of the situation (e.g. about whether someone is a genuine threat), GUIDANCE will conflict with the naïve picture of rights from above. After all, GUIDANCE seems to commit us to this conditional:

JUST FORCE: If an agent knows what kinds of situations require her to use lethal force to defend someone from an aggressor and justifiably believes that

these conditions obtain, she can justifiably use this force.

If we combine JUST FORCE with some factive view of justification, we can uphold both GUIDANCE and the idea that mistaken beliefs can excuse but not justify killing or harming in defence of self or other. If, however, we combine just force with standard non-factive views of justification, we have to say that mistaken beliefs can justify the use of lethal force or the imposition of harm, not just excuse it. This draws the distinction between justification and excuse in the wrong place. This is clear when we think about informational asymmetries. In learning that an innocent person was harmed, say, it seems that we don't just learn that something bad came to pass. We learn that reparations need to be made. In seeing an agent use force against an apparent aggressor, outsiders who know whether the apparent aggressor is a threat or not know how they ought to intervene if they were to get involved (e.g. to assist the apparent aggressor if innocent). These facts are difficult or impossible to explain if we say that the apparent aggressor's rights against others depends upon the kinds of non-factive epistemic support they might have for their beliefs.

2.2 Justification and Normative Uncertainty

If there can be false, justified beliefs, there shouldn't be any principled reason why we couldn't have false, justified beliefs about normative matters (e.g. about principles or norms, about values). As we've seen, though, GUIDANCE places limits on the kinds of things that we might have false but justified beliefs about. If, like me, you're sceptical of the suggestion that every rational thinker has undefeated justification provided by excellent evidence to believe basic normative truths and think that rational thinkers might have evidence that supports the relevant basic normative propositions to differing degrees, you might think that the two most interesting approaches to GUIDANCE will either appeal to a shifty view on which changes in the agent's evidence can help change what the agent is permitted to do or INFALLIBILISM. I shall argue that if we accept GUIDANCE, we'll need INFALLIBILISM to understand the relevance of normative uncertainty to permissibility and blame.

Let's consider some examples:

MONKS: Some monks in the grip of views about the permissibility of certain kinds of sex acts more conservative than yours give in to temptation even though they believe they ought to abstain. (Littlejohn 2014b)

Our monks might meet the conditions imposed by non-factive theories of justification for the belief that they ought to abstain and take steps to see to it that others abstain. Under such conditions, non-factive theories of justification will presumably say that they ought to believe that they ought to abstain. This, when combined with guidance, tells us that they ought to abstain. I disagree, of course. I don't think that their beliefs and the epistemic support for them generates any normative pressure towards abstaining. Moreover, it seems quite

plausible that if someone knowingly does what they shouldn't do (failing duress), they should be blameworthy for what they've done. But we should refuse to blame the monks since their actions don't manifest *de re* unresponsiveness—there was nothing that constituted a genuine reason for them to do otherwise that they were both aware of and unresponsive to. Combining GUIDANCE with non-factive theories of justification leaves us in the uncomfortable position that we either need to blame people for doing things that we know don't fail to show sufficient respect for morality's concerns or leads to a kind of scepticism about responsibility in which even known wrongdoing needn't be blameworthy.

Srinivasan offers a similar example that's helpful to consider:

DOMESTIC VIOLENCE: Radha is a woman who lives in rural India. Her husband, Krishnan, regularly beats her. After the beatings, Krishnan often expresses regret for having had to beat her, but explains that it was Radha's fault for being insufficiently obedient or caring. Radha finds these beatings humiliating and guilt-inducing; she believes she has only herself to blame, and that she deserves to be beaten for her bad behaviour. After all, her parents, elders, and friends agree that if she is beaten it must be her fault ... Moreover, Radha has thoroughly reflected on the issue and concluded that, given the natural social roles of men and women, women deserve to be beaten by their husbands when they misbehave. (2020: 8)

Srinivasan says that Radha's belief that she deserves to be beaten is not justified even though it seems that we can draw on familiar epistemological theories to support the opposing verdict. She sees this as a challenge to internalist theories of justification, but I think it poses a more interesting challenge to any non-factive theory of justification. Should we classify Radha's beliefs as justified?

Johnson King disagrees with Srinivasan's verdict:

I don't think that this is such a terrible verdict. For an internalist, saying that someone's belief is justified in no way constitutes an endorsement of their belief-forming circumstances. So, saying that Radha's belief is justified is not an approval of patriarchal ideology — no more than saying that the brain-in-a-vat's belief is justified is an approval of the practice of putting brains in vats. Rather, for the internalist, saying that someone's belief is justified is a way of giving her credit for doing the best she can, epistemically speaking, under whatever circumstances she is in. (2022: 12)

Johnson King is right that we sometimes do the best we can and end up with false normative beliefs. If internalists were happy to say that these beliefs are the result of Radha doing the best she could or that they fit well with her other attitudes, it's hard to take issue with that, but I don't think that these evaluations

track particularly important normative properties. Suppose that we say that Radha's beliefs were justified and we said that justification was a really important normative property because of its connection to rightness, permissibility, and blame. If we said this and we accepted GUIDANCE, we wouldn't just say that Radha's beliefs were justified. We would have to say that the actions and attitudes rationalized by such beliefs were also justified. This, I think, helps us see why the non-factive notions of justification are not normatively central.

Consider the belief that the beatings are deserved and justified and that a woman's duty is to submit to them without complaint. If this belief is justified, GUIDANCE tells us that it should justify the beliefs, actions, and emotions that this belief rationalizes. These would be supported by her (putatively) justified belief about the abuse she's receiving: that she's being treated properly and fairly, that it would be wrong to protest or to use force to protect herself, that she should teach her sons and daughters to assume these roles, that she, too, should use the force the tradition allows to ensure that her children are brought up to continue this pattern of domestic abuse. The emotions that would make sense given these (putatively) justified beliefs might include feelings of guilt for protesting or shame for doing the things that 'merit' the beating but would not include the feelings of anger, indignation, or resentment. When it comes to action, it seems that treating Radha's belief as justified and accepting GUIDANCE would mean that she wouldn't be justified in protesting or using force to defend herself. She would, however, be justified in, say, helping her sister's husband control her sister's life using violence and raising her children to believe, feel, and act in ways that reflect this ideology.

If GUIDANCE holds, it is hard to see how we can fully disapprove of the bad ideology insofar as we agree that it would be right for those who have internalized it to continue to believe, feel, and act in accordance with it. If, however, we try to resist this and say that the actions and emotions supported by Radha's beliefs are not justified, we could deny guidance, but then we deny that the operative notion of justification that attaches to her beliefs does interesting normative work. We'd deny that justification, so understood, rationalizes the downstream responses that belief is supposed to support.

It's also worth thinking about the connection between blame and justification. If Radha's beliefs can be justified in the scenario described above, I suppose the same should hold true for Krishnan. If Krishnan punishes Radha for protesting and believes that this is the right thing to do, should we say that his beliefs are justified? I think it's helpful to consider this passage from Rosen's discussion of moral ignorance and exculpation:

Ignorance is culpable only if it derives from culpable recklessness or negligence in the management of one's opinion. Take any case of action done from ignorance and let it be stipulated that the agent has been utterly scrupulous in policing his own opinion: he has been as careful and as inquisitive and as reflective as a person in his circumstances should be, and yet he has failed to

grasp some crucial fact. I claim that if you bear this stipulation clearly in mind, you will be persuaded in every case that the agent's ignorance is not his fault. (2004: 302)

Rosen seems to think that having justified normative attitudes might require nothing more than meeting certain procedural epistemic duties (e.g. considering salient arguments, follow the reasoning to its conclusion), but let's focus on this idea that Krishnan would be blameless if his belief that he ought to do what he's doing was justified. Rosen might be right that this conditional is quite plausible, but I also think it's clear that Krishnan is culpable for his behaviour. His beliefs, like his actions, reveal a commitment to a set of values that are incompatible with those that morality cares about. He's culpable for what he's done because what he does manifests *de re* unresponsiveness (Arpaly 2002). If so, it seems equally fair to say that he's culpable for his belief that he ought to act this way. The belief also manifests *de re* unresponsiveness. He knows why he abuses Radha and he knows of her suffering, but this doesn't move him to revise his belief that this is what he should do. His belief is thus not in line with the things that could make it right to form beliefs in this field of propositions. The lesson that I'd take from this is that the justification of our actions and attitudes couldn't just be a matter of meeting some procedural epistemic obligations (Alvarez and Littlejohn 2017).

Let me briefly recap. If knowledge is the norm of belief and justification requires conforming to norms, our justified beliefs are only justified because they're knowledge. To resist this view, epistemologists might deny that knowledge is normative for belief (e.g. by saying that it's a goal, an aim, a good thing, or a desirable state) or deny that justification requires conforming to norms, but I think the problems with these views emerge if we think about the rational relations between belief and the responses that beliefs are supposed to rationalize. If our justified beliefs are the beliefs that ought to guide us, our choice of a theory of justification is going to inform our broader views about the line between justified action and excusable wrongdoing or about the normative significance of normative uncertainty. Combining GUIDANCE with standard non-factive accounts of justification forces us to revise our views about rights and revise our views about the normative irrelevance of normative uncertainty to permissibility and blame in undesirable ways. Rejecting GUIDANCE, on the other hand, leads to lots of awkward questions. If our settled view is that we can be required to conform to norms or blamed for violating norms that we justifiably believe don't exist or don't apply to us in the situations we're in, we're externalists of a pretty radical kind. And if we think that this externalism has implications only for the downstream responses rationalized by belief but has no implications for the requirements that apply to belief itself, we're left with this odd and difficult to motivate position that the reasons against feeling or acting in line with our beliefs are often not any kind of reason at all not to form these beliefs. I don't see how this could be a stable or well-motivated position. I think that we need a radical form of externalism about justification to understand

how certain norms apply to us at all and determine what we ought to do, but this might require rejecting any view that's designed to allow for false, justified beliefs.

3. On Externalism About Knowledge

I've argued for two claims. The first is that knowledge is the norm of belief. What we ought to believe is determined by what we can know. The second is that justification requires conforming to norms like the knowledge norm. What we can justifiably believe depends upon what we can believe without violating the norms that govern belief. When these views are combined, we're left with the view that justified beliefs are justified because they constitute knowledge.

If this view is right, what does this tell us about the internalism–externalism debate? Recall Conee's characterization of the debate. He characterizes this as a debate about whether knowing *p* requires meeting the justification condition where the subject's internal states determine whether that condition is met. Given the moderate anti-sceptical assumption that internally indiscernible pairs of thinkers can know different things, if we say that our beliefs are justified because they are knowledge, we're committing ourselves to an externalist view of knowledge. If the normative standard that determines what we ought to believe determines what we have justification to believe and tells us that what we ought to believe depends upon what we can know, knowledge is prior to justification in the same way that the good is prior to the right in consequentialist views.

In framing the debate in the way that he has, Conee seems to have assumed that justification might be part of what 'turns' our beliefs into knowledge and overlooked the possibility that it's knowledge that bestows upon them their justificatory standing. Once we see that what we're responsible for as believers is seeing to it that we don't believe what we don't know, we can see why these debates about the justification condition rested on a mistake.

This mistake had a number of sources. Part of what prevented people from seeing that knowledge plays an important normative role is neglect of cases of animal knowledge where attributions of propositional knowledge seem apt even though normative evaluations of attitudes seem out of place. Part of what prevented people from seeing that knowledge plays an important normative role is the mistaken idea that belief is a state that aims at the truth and (crucially) nothing but the truth. And, of course, part of what prevented people from considering the possibility that knowledge is necessary for justified belief is a failure to take seriously arguments for externalism that weren't part of, say, the reliabilist project of trying to distinguish between beliefs that constitute knowledge and beliefs that fail to do so for Gettier-type reasons. I don't think that we'll find in these arguments convincing arguments for externalism about justification.

My argument for externalism about justification didn't start from the idea that we need something that plays a justification role that helps to eliminate some kind of bad epistemic luck. My argument for externalism started with this

idea that certain norms apply to us and determine how we ought to respond even when we're uncertain about whether they apply to us in the situation we're in (e.g. as is the case with the norms that determine when defence of self and other is permissible) or uncertain about whether the norms pick out normatively important features of the world (e.g. as was the case in MONKS or in DOMESTIC VIOLENCE). If we ought to conform to such norms and we ought to be guided by our beliefs, we ought to see to it that we don't believe falsehoods that would lead us to violate these norms. The only views that vindicate this are views that accommodate INFALLIBILISM.

References

- Adler, Jonathan. 2002. *Belief's Own Ethics*. MIT Press.
- Alvarez, Maria and Clayton Littlejohn. 2017. 'When Ignorance Is No Excuse'. In Philip Robichaud and Jan Willem Wieland (eds), *Responsibility—The Epistemic Condition*, 64–81. Oxford University Press.
- Arpaly, Nomy. 2002. 'Moral Worth'. *Journal of Philosophy* 99(5): 223–45.
- Bird, Alexander. 2007. 'Justified Judging'. *Philosophy and Phenomenological Research* 74(1): 81–110.
- Bolinger, Renée Jorgensen. 2021. 'The Moral Grounds of Reasonably Mistaken Self-Defense'. *Philosophy and Phenomenological Research* 103(1): 140–56. <https://doi.org/10.1111/phpr.12705>.
- Boult, Cameron. 2017. 'Epistemic Normativity and the Justification–Excuse Distinction'. *Synthese* 194(10): 4065–81. <https://doi.org/10.1007/s11229-016-1127-8>.
- Conee, Earl. 2004. 'Externalism, Internalism, and Skepticism'. *Philosophical Issues* 14(1): 78–90.
- Field, Claire. 2019. It's OK to Make Mistakes: Against the Fixed Point Thesis. *Episteme* 16: 175–85.
- Fox, Philip. 2019. 'Revisiting the Argument from Action Guidance'. *Journal of Ethics and Social Philosophy* 15(3). <https://doi.org/10.26556/jesp.v15i3.467>.
- Gardner, John. 2007. *Offences and Defences: Selected Essays in the Philosophy of Criminal Law*. Oxford University Press.
- Gibbons, John. 2001. 'Knowledge in Action'. *Philosophy and Phenomenological Research* 62(3): 579–600.
- Gibbons, John. 2009. 'You Gotta Do What You Gotta Do'. *Noûs* 43(1): 157–77.
- Gibbons, John. 2013. *The Norm of Belief*. Oxford University Press.
- Greco, Daniel. 2021. 'Justifications and Excuses in Epistemology'. *Noûs* 55(3): 517–37.
- Huemer, Michael. 2001. Skepticism and the Veil of Perception. *Studies in Epistemology and Cognitive Theory*, Vol. 69. Rowman & Littlefield.
- Hughes, Nick. 2018. 'Luminosity Failure, Normative Guidance and the Principle "Ought-Implies-Can"'. *Utilitas* 30(4): 439–57.
- Hughes, Nick. 2019. 'Dilemmic Epistemology'. *Synthese* 196: 4059–90.

- Hyman, John. 2015. *Action, Knowledge, and Will*. Oxford University Press.
- Ichikawa, Jonathan Jenkins. 2014. 'Justification Is Potential Knowledge'. *Canadian Journal of Philosophy* 44(2): 184–206.
- Johnson King, Zoë. 2022. 'Radical Internalism'. *Philosophical Issues*. <https://doi.org/10.1111/phis.12235>.
- Joyce, James M. 1998. 'A Nonpragmatic Vindication of Probabilism'. *Philosophy of Science* 65(4): 575–603. <https://doi.org/10.1086/392661>.
- Kelp, Christoph. 2016. 'Justified Belief: Knowledge First Style'. *Philosophy and Phenomenological Research* 93(1): 79–100.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford University Press.
- Kornblith, Hilary. 1999. 'Knowledge in Humans and Other Animals'. *Philosophical Perspectives* 13: 327–46.
- Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Lasonen-Aarnio, Maria. 2019. 'Guidance, Epistemic Filters, and Non-Accidental Ought-Doing'. *Philosophical Issues* 29(1): 172–83. <https://doi.org/10.1111/phis.12166>.
- Littlejohn, Clayton. 2009. 'The Externalist's Demon'. *Canadian Journal of Philosophy* 39(3): 399–434.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press. <https://www.cambridge.org/core/books/justification-and-the-truthconnection/1836026B65551E20D6E4EE623C6066BE>.
- Littlejohn, Clayton. 2013. 'The Russellian Retreat'. *Proceedings of the Aristotelian Society* 113(3): 293–320.
- Littlejohn, Clayton. 2014a. 'Fake Barns and False Dilemmas'. *Episteme* 11(4): 369–89.
- Littlejohn, Clayton. 2014b. 'The Unity of Reason'. In Clayton Littlejohn and John Turri (eds), *Epistemic Norms*, chapter 8. Oxford University Press.
- Littlejohn, Clayton. 2018. 'The Right in the Good: A Defense of Teleological Non-Consequentialism in Epistemology'. In Jeff Dunn and Kristoffer Ahlstrom-Vij (eds), *Epistemic Consequentialism*, 23–48. Oxford University Press.
- Littlejohn, Clayton. 2019. 'Being More Realistic About Reasons: On Rationality and Reasons Perspectivism'. *Philosophy and Phenomenological Research* 99(3): 605–27.
- Littlejohn, Clayton. 2021. 'On What We Should Believe (and When (and Why) We Should Believe What We Know We Should not Believe)'. In K. McCain and S. Stapleford (eds), *Epistemic Duties: New Arguments, New Angles*, chapter 12. Routledge.
- Lord, Errol. 2018. *The Importance of Being Rational*. Oxford University Press.
- Lynch, Michael Patrick. 2005. *True to Life: Why Truth Matters*. MIT Press.
- Madison, B. J. C. 2018. 'On Justifications and Excuses'. *Synthese*, 195: 4551–62. <https://doi.org/10.1007/s11229-017-1418-8>.

- Nolfi, Kate. 2021. 'Epistemic Norms, All Things Considered'. *Synthese* 198: 6717–37. <https://doi.org/10.1007/s11229-019-02485-2>.
- Oliveira, Luis. 2018. 'Deontological Evidentialism and Ought Implies Can'. *Philosophical Studies* 175(10): 2567–82.
- Oliveira, Luis. 2022. Epistemic Consent and Doxastic Justification. In L. Oliveira and P. Silva (eds), *Propositional and Doxastic Justification*, chapter 16. Oxford University Press.
- Reynolds, Steven L. 2013. 'Justification as the Appearance of Knowledge'. *Philosophical Studies* 163(2): 367–83.
- Rosen, Gideon. 2004. 'Skepticism about Moral Responsibility'. *Philosophical Perspectives* 18(1): 295–313.
- Simion, Mona. 2019. 'Knowledge-First Functionalism'. *Philosophical Issues* 29(1): 254–67.
- Smithies, Declan. 2015. 'Ideal Rationality and Logical Omniscience'. *Synthese* 192(9): 2769–93.
- Srinivasan, Amia. 2020. 'Radical Externalism'. *The Philosophical Review* 129(3): 395–431. <https://doi.org/10.1215/00318108-8311261>.
- Sutton, Jonathan. 2005. 'Stick to What You Know'. *Noûs* 39(3): 359–96. <https://www.jstor.org/stable/3506235>.
- Sylvan, Kurt. 2018. 'Knowledge as a Non-Normative Relation'. *Philosophy and Phenomenological Research* 97(1): 190–222.
- Titelbaum, Michael. 2015. 'Rationality's Fixed Point (Or: In Defense of Right Reason)'. In John Hawthorne and Tamar Gendler (eds), *Oxford Studies in Epistemology*, Vol. 5, 253–94. Oxford University Press.
- Unger, Peter. 1975. *Ignorance: A Case for Scepticism*. Oxford University Press.
- Wedgwood, Ralph. 2002. 'The Aim of Belief'. *Philosophical Perspectives* 36(s16): 267–97.
- Whiting, Daniel. 2012. 'Does Belief (Only) Aim at the Truth?'. *Pacific Philosophical Quarterly* 93(2): 279–300.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.
- Williamson, Timothy. 2019. 'Morally Loaded Cases in Philosophy'. *Proceedings and Addresses of the American Philosophical Association* 93: 159–72.
- Zimmerman, Michael. 2008. *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.