

# Access Externalism

JOHN GIBBONS

This paper argues for externalism about justification on the basis of thought experiments. I present cases in which two individuals are intrinsically and introspectively indistinguishable and in which intuitively, one is justified in believing that *p* while the other is not. I also examine an argument for internalism based on the ideas that we have privileged access to whether or not our own beliefs are justified and that only internalism is compatible with this privilege. I isolate what I take to be the most plausible form of privileged access to justification and show that it is compatible with externalism.

The notion of knowledge is an externalist notion. Since truth is an independent necessary condition on knowledge, there can be two individuals who are the same on the inside, duplicates with respect to their introspectively accessible properties, one with a true belief and one with a false belief. This could make the difference between one who knows and one who does not. If you think of knowledge as justified, true belief plus *X*, it is clear that the fourth condition, whatever it is, is an externalist notion. You could have two individuals with the same beliefs and experiences (and whatever else is both epistemically relevant and introspectively accessible), one who infers from a false premiss and one who infers from a true premiss (Gettier 1963), or one who is in fake barn country and one who is in real barn country (Goldman 1976). These intuitively external facts could make the difference between one who knows and one who does not. If you define 'warrant' as whatever it is that makes the difference between knowledge and true belief (Plantinga 1993, p. 3), then warrant, since it captures the fourth condition, is an externalist notion.

If there is a characteristically epistemic internalism, it has to be internalism about good old-fashioned justification. Whatever relations hold among the various epistemic notions, justification is the only thing left to be an internalist about. It could be that warrant is the more basic or more important notion. It could be that justification is not necessary for knowledge (Alston 1988). It could be that knowledge is in principle unanalysable so any talk about a fourth condition rests on a mistake (Williamson 1995). These are all important issues. They're just not

mine. My question is whether justification as we ordinarily understand it is a purely internalist notion. I will argue that it's not.

Internalism is a supervenience thesis. I will be primarily concerned with what is sometimes called 'Access Internalism', the view that justification supervenes on introspectively accessible properties of the believer, though I think the arguments work against the view that justification supervenes on the internal on any natural understanding of 'internal'.<sup>1</sup> The idea behind internalism is that only internal properties are *directly relevant* to justification. The internalist can readily agree that pretty much any fact could be indirectly relevant to justification. If the external fact that  $p$  causes an experience that  $p$ , and this causes you to believe that  $p$ , the external fact influences not only the existence of the belief but also its justificational status.

According to the internalist, this influence is necessarily indirect. If you had had the experiences in the absence of the fact, this would have been just as good from the point of view of justification, at least according to the internalist. On that view, an external fact can only make a difference to justification by, for example, causing, or resulting in, or somehow making a difference in the internal facts.<sup>2</sup> If you can have the internal fact without the external fact, and there are differences in justification only when there are differences in the internal facts, it is tempting to suppose there is some sense in which the internal facts are doing all the work. The presence or absence of an external fact, all by itself, is never directly relevant to justification. External facts only make a difference to justification by making a difference to the introspectively accessible facts.

Since internalism about justification is a supervenience thesis, the strategy for arguing against it is fairly straightforward. You just tell some stories. If you can have two people who are the same on the inside, in the relevant sense of that expression, but different on the outside where, intuitively, one of them is justified but the other is not, then internalism about justification is false. If internalism is false, externalism is true. No matter how important the inside is, some of the stuff on

<sup>1</sup>I get this terminology from Alston (1986). Sometimes the idea is that justification supervenes on that which is knowable on the basis of reflection, that is, introspection or a priori reasoning. If we leave aside cases of the contingent a priori, this comes to the same thing. Pryor (2001) calls this view 'Simple Internalism'.

<sup>2</sup>I apologize for the disjunction, but causation is not the only relation that generates indirect relevance. Supervenience is another. Suppose that justification supervenes on the mental and that the mental supervenes on the physical. Even if those physical states were intrinsic to you, they would not be internal to you in the epistemic sense. You cannot, in general, know through introspection that you are in physical state  $P$ . Brain states and the like are relevant to justification for the internalist but only indirectly so. They make a difference by making a difference to the inner life.

the outside matters as well. The presence or absence of an external fact, all by itself, is directly relevant to justification.

The first part of this paper consists largely of the counterexamples to internalism. The second part of the paper is primarily concerned with one particular argument for internalism. I focus on this one because I think it is the most plausible. The basic idea behind the argument is that only internalism explains or is consistent with our privileged access to the justification of our own beliefs. I will argue that we do have a certain kind of privilege with respect to justification. But since the kind of externalism I favour explains and is consistent with this kind of privilege, the argument for internalism does not go through. Since the kind of externalism I favour accepts an accessibility requirement on the facts that make us justified and delivers privileged access to the facts about justification themselves, I call it 'Access Externalism'.

There is one further thing of a preliminary nature before we get to the stories. I take it as fairly obvious that there are things we ought to believe and things we ought to know. If accepting the legitimacy of the epistemic 'ought' is sufficient for accepting a deontological conception of justification, then I think we should all be deontologists, and I do not mind the label. But I will show that this minimally deontological conception of justification leads directly to externalism, not internalism.<sup>3</sup> If, on the other hand, something more substantive is required for being a deontologist, then I may or may not be one. If not, I will not miss the label either.

While the labels do not matter, here is what does. There are things you ought to know and things you ought to believe. If you ought to know that *p*, then you ought to believe that *p*.<sup>4</sup> The notion of what you ought to believe is closely connected to one notion of what you are justified in believing. If I say, 'You are justified in believing that *p*', this is ambiguous. On the stronger reading, it entails that you do believe that *p* and that this belief is justified. On the weaker reading, neither of these things follows. On this reading, it means, roughly, that you have available all the justification you need for believing that *p*, whether you

<sup>3</sup> Both Plantinga and Alston argue that the deontological conception of justification leads to internalism. See Alston (1986) and Plantinga (1993) Ch. 1. Some internalists argue this as well. See, for example, Ginet (1975), pp. 28–36, and Chisholm (1977), pp. 14–5. Goldman (1999) argues against the derivation of internalism from deontologism.

<sup>4</sup> I don't think your obligations are closed under entailment. You ought to know that Hitler was evil. If you know that Hitler was evil, then Hitler was evil. But it doesn't follow that it ought to be the case that Hitler was evil. The particular entailment, if you ought to know it then you ought to believe it, holds because both knowledge and belief are epistemic notions subject to the same epistemic 'ought'.

believe it or not and whether you believe it for the right reasons or not. If you ought to believe that  $p$ , you are justified in believing that  $p$  in the weak sense.

While I will rely on your intuitions about when people ought to know and believe things, nothing I say will rely in any way on anyone's intuitions about when blaming someone is appropriate. The appropriateness of the action does not completely determine the appropriateness of blaming. If I stop blaming you because I forgive you, something that happens every day, this is a situation in which I think you did something wrong, but I do not think blame is appropriate. Forget about blame, and think about what ought to be.

### 1. Externalism

So, finally, here's a story. The other morning, I went downstairs to make a mushroom, jalapeno, and cream cheese omelette. I had checked the night before to make sure we had all of the ingredients. Since Sunny rarely eats breakfast, it was reasonable for me to believe that the ingredients were still there. I went to the refrigerator and pulled out the eggs and mushrooms. While chopping, I firmly believed that I would soon have a mushroom, jalapeno, and cream cheese omelette. Unfortunately, in plain sight on the door of the refrigerator, there was a note. 'We're out of cream cheese.' I didn't notice the note, but I should have. After all, this is where we leave notes of this sort in our house. I thought I was having cream cheese for breakfast, but I should have known better. If I should have known that  $\text{not-}p$ , then I'm not justified in believing that  $p$ .

In a nearby possible world, things are the same inside me, but slightly different inside and outside the refrigerator. In that possible world, there is cream cheese in the fridge and no note on the door. My belief that I would have cream cheese for breakfast was based on the same reasonable grounds. Unlike the first case, there was no evidence (or potential evidence) available to me to override these reasonable grounds. So in this case, I am justified in my belief about breakfast. If the mere possibility of unnoticed evidence is inconsistent with justification, then we are justified in a lot fewer things than we think. In fact, I am inclined to say that in this case I know what I will have for breakfast. If knowledge, or at least this kind of knowledge requires justification, we get the same result: two people the same on the inside, different on the outside, where one is justified and the other is not.

The details of the case don't matter that much. The structure looks like this. In the first story, there is some evidence available to you.

Maybe we should call it ‘potential evidence’. In any case, there’s a set of facts such that if you were aware of them, you would no longer believe that *p*. But it’s not enough that there are such facts. They have to be, in some intuitive sense, available to you. When the facts are both relevant and available, we can say that you are epistemically responsible for them.<sup>5</sup> When the note is on the door, I am responsible for seeing it. But if Sunny had written the note and absentmindedly stuck it in her pocket and brought it to the studio, it would not be available to me, and I would not be responsible for knowing about it. Epistemic responsibility as I understand it is not a responsibility for doing something. It is a responsibility for knowing something. Not all obligations, even moral obligations, are obligations to act. You ought to care about the well-being of others, even if caring is not an action. To say that you are epistemically responsible for the fact that *p* is to say, roughly, that whether you are aware of the fact or not, you ought to know that *p*.

So in the first story, you have available, unnoticed evidence.<sup>6</sup> It’s easy to get from there to the claim that you should have known better. Then we apply the principle that you are not justified in believing things you should know are false. You can tell the second story in two different ways. In the original version of our story, the evidence did not exist. But in our modification of the story, where the note is in Sunny’s pocket, the evidence exists but is not available to me. Since what I am epistemically responsible for depends on both the existence of the evidence and its relation to me, and since neither of these things is completely determined by my introspectively accessible properties, neither is justification.

Oddly enough, not everyone is moved by my original version of the story. So I will show you how to construct a range of stories, and you can use the one that suits your intuitions. We begin with standard-sized notepaper stuck to the door of a refrigerator. Now gradually increase the size of the note. Make the letters larger, brighter, and more colourful until you get to a billboard-sized note with the letters written out in bright lights. Somewhere along the line, probably well before you get to the billboard, you will get to the point where you think that whether I

<sup>5</sup> There has been a lot of work done on the notion of epistemic responsibility. Usually the emphasis is on our epistemic responsibilities to act or to believe, while our responsibility to know certain things often receives less attention. See, for example, Kornblith, (1983); Greco, (1990); and Miller (1995).

<sup>6</sup> The connection between knowledge and unpossessed evidence is familiar enough. See Harman (1980). Harman’s stories are cases of justified true beliefs that do not constitute knowledge because of the unpossessed evidence. My stories show that when the evidence is something you ought to possess, this can affect justification as well.

noticed or not, I should have. If I failed to notice something I should have noticed, then you are inclined toward some negative epistemic evaluation. I should have known better.

Now this one story does not tell you anything about my epistemic character, or habits, or dispositions. Maybe this kind of thing happens all the time with me. But maybe I am usually quite epistemically virtuous, despite this one brief lapse. So your reaction to the story, your negative epistemic evaluation, cannot be a judgement about my character, habits, or dispositions. It is, as it appears to be, an evaluation of a particular belief. I believed I would have cream cheese, but I should have known better.

We can connect this evaluation, the claim that I should have known better, with the notion of justification by way of two principles. The first principle says that you shouldn't believe something you ought to know is false. Of course, we can evaluate beliefs from any number of different perspectives. We can evaluate them morally, practically, and aesthetically. But if we read the 'shouldn't' and 'ought' in the principle as expressing a purely epistemic evaluation rather than a moral or all-things-considered evaluation, then the principle is fairly obvious. If, from the epistemic point of view, you ought to know that not- $p$ , then, from that point of view, you shouldn't believe that  $p$ .<sup>7</sup>

The first principle connects two normative notions, what you ought to know with what you should or shouldn't believe. The second principle connects these notions with that of justification. The second principle says that if you shouldn't believe that  $p$ , then you are not justified in believing it if you do. I don't know if doing your epistemic duty or fulfilling your epistemic obligations is sufficient for justification. I am inclined to guess that it is not.<sup>8</sup> But I am pretty sure that failure to fulfil your obligations or to do your duty is inconsistent with justification. As long as this talk about duties and obligations is the right way to understand 'should' from the epistemic point of view, this is what the principle says. Putting the two principles together, we get the claim that if you ought to know that not- $p$ , then you are not justified in believing that  $p$ .

<sup>7</sup> Here's one way of looking at it. If you ought to know that not- $p$ , then you ought to believe that not- $p$ . Since you shouldn't believe a contradiction, if you should believe that not- $p$ , you shouldn't believe that  $p$ .

<sup>8</sup> This depends, of course, on what your duty is. If you ought to know that  $p$ , and you do what you should, then you are justified in believing that  $p$ , at least if justification is necessary for knowledge. But sometimes, even if  $p$  is false, we say that you should believe that  $p$ , and fulfilling this obligation might not be sufficient for justification in the stronger sense. You might believe what you should, but believe it for bad reasons.

We have to be careful in telling our stories. It is tempting to suppose that if two people are indistinguishable with respect to their introspectively accessible properties, then each is as careful as the other. I do not think that is true,<sup>9</sup> but for the moment, suppose it is. If each is exceptionally careless, then neither is justified. If both are exceptionally careful, it is hard to see how the first person would miss the available evidence. But we spend most of our lives in between these two extremes. In our stories, we need a degree of attentiveness where if the world cooperates, then you are justified, and if not, you are not. I think this degree of attentiveness characterizes the overwhelming majority of our waking lives. Even if not always, most of the time, you can imagine the possibility of something you did not notice but should have. If the mere possibility of this state of affairs is inconsistent with justification, then we are not justified in believing very much. If it's the existence of the state of affairs that is inconsistent with justification, then justification is determined, at least in part, by external things.

All of this is just a recipe for counterexamples. It tells you what to look for in a story. The work is done by the stories themselves. Here is another fairly straightforward story. Sometimes we forget things we ought to remember. Excited about the prospect of having my favourite kind of omelette the next morning, I check the refrigerator for ingredients. I discover, to my dismay, that we have no cream cheese. So I decide to settle on swiss. The next morning, my intention to have an omelette is still with me, but I have forgotten the cheese situation. While chopping the mushrooms, I believe, as it were, out of habit, that I will have a mushroom, jalapeno, and cream cheese omelette. But I should know better.

If internalism about justification is the view that my present justification for believing that *p* is determined by what is presently introspectively accessible to me, then the view is in trouble. Imagine a possible world in which I check the refrigerator the night before and there is cream cheese. In the morning, there is nothing I have forgotten that I should have remembered. I remember the fact that we have cream cheese without remembering the event of checking. So I am introspectively indistinguishable from myself in the other story. In this story, not only does my intention to have an omelette stay with me, but also my intention to have a particular kind of omelette stays with me. Whether I remember this or not, my belief that I will have that kind of omelette is grounded in my having checked the ingredients the night before.

<sup>9</sup> Subconscious desires and other introspectively inaccessible mental states can influence your degree of carelessness.

The stories so far have been intended to show that the presence or absence of some external facts can make a difference to the justification of belief. This leads to externalism about justification. If there are external facts that you ought to know about whether you are aware of them or not, then you can have two people who are the same on the inside with different epistemic obligations. If you can have two people who are the same on the inside but have different epistemic obligations, it could be that one of these people satisfies her obligations while the other does not, even though they believe the same things. And if the belief that  $p$  is the result of a failure to fulfil your epistemic obligations, if it's an irresponsible belief, then you shouldn't believe that  $p$ , and you are not justified in believing if you do.

The fact that some external facts are directly relevant to justification in this way does not suggest that they all are. The facts have to be available to you in some sense, and we need to say something about this notion of availability if it is not understood in terms of introspective access. And availability is not sufficient, since there are all sorts of irrelevant facts that you are in a position to know, though you are not under any obligation to know. In this paper, I must set aside the important question of what makes a fact, mental or otherwise, relevant to the justification of a belief. Here, I focus on the question of what kind of access we must have to the relevant facts. In this, I take it that I am following the internalist. Internalism is a supervenience thesis. It does not say that everything in the supervenience base is relevant. It says that everything relevant is in the supervenience base. If you lack some belief about how you don't feel, the introspectively accessible fact that you don't feel that way may be irrelevant to the justification of your beliefs.

So some (but not all) external facts are directly relevant to justification. So justification does not supervene on the internal. So internalism is false. Only (but not all and only) accessible or available facts are relevant. Whether the facts are mental or physical, internal or external, or necessary or contingent is less important than whether the facts are accessible. This raises two questions. First, why is there any restriction to accessible facts? And second, what kind of accessibility is involved if not introspective accessibility? We turn to these questions in the second part of the paper.

## 2. Internalism

### 2.1 Access

My argument against internalism is fairly straightforward. What you are paying attention to does not determine what you should be paying attention to. What you think you should be paying attention to does not determine what you should be paying attention to either. Since what you ought to know or notice is determined both by what you are like on the inside and on what is going on around you, you can have two people who are the same on the inside, but different on the outside, and who have different obligations. So you can have intrinsic or introspectively indistinguishable duplicates, one who fulfils and one who fails to fulfil her epistemic obligations.

I think we should take these epistemic obligations seriously. If you should know that  $p$ , you are not justified in believing not- $p$ . You shouldn't believe something you ought to know is false. And if you shouldn't believe something, you are not justified in believing if you do. Since the claim that you should know that  $p$  is factive, since that normative claim entails that  $p$ , a difference in the facts can result in a difference in justification without leaving a trace on the inner life. So internalism is false. So externalism is true. Who would have thought any different?

I think that the intuitive idea behind internalism is that we have some kind of privileged access to the facts about justification. While we could, through self-importance or lack of attention, make a mistake about whether one of our own current beliefs is justified, it is harder to imagine a case in which we make a non-culpable mistake about the justification of our own beliefs.<sup>10</sup> But if justification did depend on external facts, it is hard to see how we could be privileged about it, since we do not have privileged access to the external world. This fairly straightforward argument for internalism is the focus of this part of the paper.

The internalists are right in stressing our privileged access to justification. The question is how exactly to understand this privilege. Though I am officially an externalist about justification, I must confess that I am unable to shake the intuition that my twin in a vat has a large number of justified beliefs.<sup>11</sup> I have tried to shake this intuition, but

<sup>10</sup> Plantinga (1993) Ch. 1 discusses this motivation for internalism. He seems to accept the inference while rejecting the deontological conception of justification.

<sup>11</sup> I mean, of course, a recently envatted twin, one whose thoughts are about the same things that my thoughts are about.

nothing seems to work. When I look for grounds to support this intuition, the best I can come up with are some claims about accessibility. My twin has done the best that he could, and there is nothing accessible to him or available to him to override his firm conviction that he has a hand. Since there is no way he could have known, we cannot say that he should have known any better.

Now one pair of individuals who are introspectively indistinguishable and alike with respect to justification does not entail that all such pairs are alike. My stories show that not all such pairs are alike. But the stories do not sever the connection between justification and access. They simply rely on an ordinary, commonsense notion of accessibility according to which some, but not all of the facts about the external world are accessible to you. The facts about my refrigerator are accessible to me but not to you, though I can make them accessible to you just by telling you about them.

What is the ordinary, commonsense notion of accessibility? I think it is the notion of what you are in a position to know. This differs from what you can know or are capable of knowing in two different ways. You might be able to find out, and so come to know, all sorts of obscure facts about arachnids by consulting a reliable source. But you are not already in a position to know them. You are, presumably, in a position to know what colour the walls are in whatever room you happen to be in, at least if you can see them from where you are. You might have to think about it for a moment or look up from where you are, but if you do not have to significantly change your epistemic situation in order to find out, then you are in a position to know. Of course, the vagueness of 'significantly' will infect the notion of what you are in a position to know. But this is exactly what you would expect from a commonsense notion.

In addition to things you can find out that you are not in a position to know, there may be facts you are in a position to know that you are unable to face.<sup>12</sup> If a baby falls in the river, and you are the only person in a position to save her, you might be so paralysed by shock that you are unable to save her. You ought to save her, and, as it says in the story, you are in a position to save her, but there is at least some inclination to say that you cannot save her. If you believe that  $p$  as a result of wishful thinking, and the evidence that not- $p$  is staring you in the face, we say that you ought to know that not- $p$ . If the desire that  $p$  is strong enough,

<sup>12</sup> Here I disagree with Williamson (2000), p. 95 when he says that if you are in a position to know that  $p$  then no obstacle must block your path to knowing that  $p$ . Though there is some disagreement, I am relying on the same basic idea that Williamson is talking about.

there may be some inclination to say that you cannot know that not- $p$ . ‘Ought’ might imply some kind of ‘can’, but an inability due to irrationality does not preclude an obligation. You are in a position to know what is going on in the room around you, and you are in a position to know what is going on in your own mind.

If you cannot bring yourself to face the unpleasant fact that  $p$ , when the fact that  $p$  is out there in the world but staring you in the face, we might say that you are in a position to know that  $p$ , and so this is the kind of fact for which you may be epistemically responsible. If this is along the right lines, internalism is along the wrong lines. If the kind of accessibility relevant to justification is anything like the ordinary, commonsense notion of accessibility, internalism is false. There is no non-sceptical reading of ‘what you are in a position to know’ that restricts its extension to internal facts. Accessibility without internalism gives you Access Externalism.

### 2.2 Privileged access

Though this weak kind of accessibility requirement might explain intuitions about brains in vats, it cannot be the rationale behind internalism. Perhaps the need for privileged access comes from thinking about justification in terms of ‘ought’ and ‘should’ and the need to bridge the gap between objective and subjective obligations. Suppose, if you can, that it is an objective fact that you ought to  $A$ . Since this is an objective fact, you could be wrong about it. So suppose that while you ought to  $A$ , you believe that you ought not to  $A$ . Surely there is something wrong with  $A$ -ing when you think you shouldn’t. But there is also something wrong with failing to  $A$  when you should  $A$ . Now if your belief that you shouldn’t  $A$  is completely unjustified, if, for example, you should have known better, then this kind of ignorance will not get you off the hook. But what if you have a justified false belief that you shouldn’t  $A$ ? What do we say then?

At this point, we might distinguish objective from subjective obligations.<sup>13</sup> You have an objective obligation to  $A$  when you ought to  $A$  (when you really ought to  $A$ ). And you have a subjective obligation to  $A$  when you are justified in believing that you (objectively) ought to  $A$ . If we are interested in rationality or epistemic justification, there is a strong temptation to emphasize your subjective obligations. Suppose you are justified in believing that you ought to believe that  $p$ . Isn’t this reason enough to believe that  $p$ ? Is there any sense of ‘should’ that is relevant to justification in which you shouldn’t believe that  $p$ ? You might

<sup>13</sup> For a number of ways of making this distinction, see Alston (1985).

get the facts about  $p$  wrong, but even so, isn't this a justifiable error? If there is no such sense of 'should', then at least as far as justification goes, part of the distinction between objective and subjective obligations collapses. If you are justified in believing that it is a requirement, then it really is a requirement. All subjective obligations are objective obligations.<sup>14</sup>

If subjective obligations are obligations then subjective permissions are permissions. If you are justified in believing it is permissible, then it is. If we think of a justified belief as an epistemically permissible belief, then justification satisfies what we might call the JJ Thesis.

(JJ) If you are justified in believing that you are justified in believing that  $p$  then you are justified in believing that  $p$ . ( $JJp \rightarrow Jp$ )

This is a fairly strong form of privileged access to the facts about justification. For most contingent propositions, justification does not entail truth. But for contingent propositions about the justification of your own current beliefs, according to one version of (JJ), justification does entail truth.

The JJ Thesis captures the idea that you cannot make a non-culpable error about the justification of your own beliefs. Of course, mistakes are possible, but if you are not justified in believing that  $p$ , and you think you are, then you should have known better, that is, your second-order belief is not justified. In short, if not- $Jp$  then not- $JJp$ , which is the contrapositive of (JJ).

The truth of (JJ) is not obvious. It needs to be argued for. The dustbin of philosophy is filled with purely formal epistemic principles almost all of which are false. Given different interpretations of the 'J', there are different readings of (JJ). The stronger reading of the 'J' means that your belief that  $p$  is justified, while the weaker reading means that you have available all the justification you need whether you make use of it or not. I'll concentrate on two readings of (JJ), the stronger reading where 'J' is given the strong reading throughout and the weak reading where it has the weak reading throughout.

I am pretty sure that the weak reading of (JJ) is true. If you are permitted to believe that you are permitted to believe that  $p$ , you are permitted to believe that  $p$ . The initial permission, like all the others, is not

<sup>14</sup>The subjective/objective obligation distinction precisely parallels the subjective/objective cat distinction. There's an objective cat in front of you iff there's a cat in front of you. There's a subjective cat in front of you iff you are justified in believing there's a cat in front of you. The idea that all subjective obligations are objective obligations is the best explanation of the fact that while no one (who isn't confused) takes subjective cats seriously, people do take subjective obligations seriously. Subjective cats are not cats. But perhaps, subjective obligations are obligations.

an all-things-considered permission. It is a purely epistemic permission. Any moral or practical considerations that are not epistemically relevant are ignored. But it is an all-epistemic-things-considered permission. Whether you are permitted to believe that  $p$  is one of the epistemic things that needs to be considered when trying to figure out whether you are permitted to believe that you are permitted to believe that  $p$ . If the consequent is false, if it is not the case that you are permitted to believe that  $p$ , this will outweigh any considerations in favour of the claim that you are permitted to believe that you are permitted to believe that  $p$ . Though the reading of (JJ) is weak, the notion of privileged access is still substantive.<sup>15</sup> It does not follow from the fact that you are permitted to believe that you have a hand that you have a hand. 'You are permitted to believe', like 'you ought to believe' is not factive. But, according to (JJ), beliefs about your own justification, or your own epistemic obligations are special.

I don't know whether the stronger reading of (JJ) is true. What is surprising, to me at least, is just how plausible it is. But let's take a look at the dialectical situation. I am considering an argument for internalism. From the point of view of that argument, the more privileged access we have, the better things look for the internalist. As an externalist, it is tempting to simply assume for the sake of argument that the stronger reading is true. As we will see, externalism is compatible with both versions of (JJ), and this is what matters as far as the argument for internalism goes. But since I think this notion of privileged access is sufficiently interesting for its own sake, I will try to show how to respond to one kind of counterexample to the stronger reading of (JJ).

Suppose that a reliable and trustworthy source convinces you that you are justified in believing that  $p$ , so you have prima facie, second-order justification. But your belief that  $p$  is the result of wishful thinking. You believe that  $p$  for bad reasons, and while you may have reasons to believe that  $p$ , they are not the reasons for which you believe that  $p$ , and you would continue to believe that  $p$  no matter what other reasons you had or didn't have. Since the first-order belief is caused and sustained in a bad way, it is not justified.

Here is one way to deal with the apparent counterexample. If you believe that  $p$  on the basis of bad grounds, then you should know that they are bad grounds. Wishful thinking happens every day, and in the overwhelming majority of cases, the people who are doing it do not

<sup>15</sup> Substantive, but not extreme. We're not infallible about justification, and the condition of being justified is not, as far as I can tell, a luminous condition. For luminosity, see Williamson (2000), Ch. 4.

know that they are doing it. But none of this shakes my conviction that from the epistemic point of view, you shouldn't believe on the basis of wishful thinking.

If you *A* when we would have thought that you shouldn't *A*, it is not enough to say that you did not know you were doing it. If you should have known you were doing it, culpable ignorance will not get you off the hook. If you want to say that people shouldn't engage in wishful thinking, then you need the claim that if people do engage in wishful thinking, they should know that they are. And in general, if you think that people shouldn't believe for bad reasons, then you need the claim that if they do, they should know better.

If we take our epistemic obligations seriously, we get out of the counterexample. You have *prima facie*, second-order justification for believing that you are justified in believing that *p* because of the trustworthy source. But if you should have known that your belief that *p* was not justified, if you should have known that it was based on bad grounds, then all things considered, you are not justified in believing that you are justified in believing that *p*. The absence of first-order justification, or the fact that you should have noticed its absence keeps the second-order *prima facie* justification from constituting all-things-considered justification.<sup>16</sup>

The basic idea behind (JJ) is that the second-order process of evaluating your own beliefs is not independent of first-order inquiry. From your point of view, the question of whether you are justified in believing that *p* is not independent of the question of whether *p* is true. If the questions were independent, any pair of answers would be acceptable in some situation or another. But you should never believe a proposition of the form

(*p* but I am not justified in believing *p*).

<sup>16</sup> All right, one more. Suppose you have a justified but false epistemological theory that says that if you are in *C* then you are justified in believing that *p*. Since the theory is false, there could be someone in *C* who isn't justified in believing *p*. Suppose that person believes the epistemological theory and knows she's in *C*. This person could justifiably conclude that she's justified in believing that *p*. This is *Jp*. But we assumed that her first-order belief wasn't justified. This is not-*Jp*.

The reply goes like this. If the theory is false, there could be someone in *C* who's not justified in believing that *p*. But it surely doesn't follow that no one in *C* is justified in believing *p*. People with independent reasons for believing *p* might be in *C* but not be justified in virtue of being in *C*. I'm inclined to think that someone who's justified in believing the theory (if you are in *C* then you are justified in believing *p*) and who knows she's in *C* has pretty good reason to believe that *p*. You have to consult your intuitions here. It's no good to simply rely on (JJ) at this point. But I do think this reaction to the story is fairly natural. If (JJ) was incompatible with the possibility of a justified but false epistemological theory, that would be another story. Thanks to Al Casullo for the objection and Mark Van Roojen for the reply.

Nor should you believe anything of the form

(*p* is false, but I am justified in believing it).<sup>17</sup>

The process of second-order evaluation seems fairly minimal, even when it occurs because it is nothing, or at least very little, over and above the process of first-order inquiry. How do you figure out whether other people are justified in believing things? By figuring out what is accessible to them. How do you figure out what you are justified in believing? By figuring out what is accessible to you. But this just is figuring out what is true in your neighborhood.

If the second-order process of evaluation contains or should contain the process of first-order inquiry as a part, then any facts relevant to the justification of the first-order belief are relevant to the justification of the second-order belief. In fact, it is tempting to suppose that first-order justification itself is directly relevant to the justification of the relevant second-order belief if you form one. In the ordinary case, your second-order belief is justified because your first-order belief is justified. If that is how it ought to go, there's no mystery why second-order justification is sufficient for first-order justification.<sup>18</sup>

### 2.3 From access to privilege

So I am looking for a plausible argument for internalism about justification because I am sure the internalists are right about something. The argument, as I see it, has two steps, and so far, we have only looked at the first step. This is the idea that we have privileged access to the facts about justification. We could not make a non-culpable mistake about whether we are justified. The mistake comes in the second step. The mistake is to think that the only way we could have privileged access to justification is if internalism is true. We might have privileged access to our own mental states, but we do not have privileged access to the external world. For any external proposition, we could make a non-culpable mistake about that. So, if justification did depend on the external facts, we would not have privileged access to it. But we do. So it doesn't.

<sup>17</sup> Propositions of this form are Moore paradoxical. See Moore (1962), p. 277. Adler (2002) makes a great deal of this version of Moore's paradox, though his idea is not that you shouldn't believe propositions of these forms. Adler thinks it's impossible to believe propositions of these forms.

<sup>18</sup> I'm not suggesting that you ought to have a second-order belief about justification for every first-order belief you have. But I am suggesting that if you do form the relevant second-order belief, then you are epistemically responsible for the facts about first-order justification. If that's right, then the only way the second-order belief could be false is if you fail to know something you ought to know, in which case, the second-order belief would be unjustified.

I think that this is a plausible argument for internalism, though not, of course, a sound one. There is a problem understanding the claim that if justification did depend on the external facts, we would not have privileged access to it. Suppose that justification depends on the contents of your beliefs and that the contents depend on the neurological facts or the atoms in the void. We do not have privileged access to these latter facts, at least not under these descriptions, but this does not look like a good reason to deny the dependence of justification on the mental, or the dependence of the mental on the physical, or the transitivity of dependence.

Let's just suppose that the internalist can fix this.<sup>19</sup> There is still room for disagreement. Let's look a little more closely at the claim that for any external proposition, we could make a non-culpable mistake about that. This has a straightforwardly true reading. For any proposition  $p$  about the external world, there is a possible epistemic situation  $e$  in which I make a non-culpable mistake about  $p$  in  $e$ . If you think that brains in vats are justified, you can think of a situation in which you make a non-culpable mistake about whether you have a hand. But what is this to me? If I'm not in  $e$ , if I'm in  $e'$ , then maybe I couldn't make a non-culpable mistake about  $p$  in  $e'$ . If you and I are in different epistemic situations, then your culpability or lack thereof is irrelevant to mine.

So how do we individuate epistemic situations? Why, epistemically of course. Whatever an epistemic situation is, the following should hold true. If  $a$  has justification for believing that  $p$  in  $e$  and  $b$  does not have justification for believing that  $p$  in  $e'$ , then  $e$  is not the same epistemic situation as  $e'$ . You just don't get for free the claim that if  $a$  and  $b$  are the same on the inside then  $a$  and  $b$  are in the same epistemic situation. This is to beg the question. But I do not have to beg the question against the internalist. I have already argued against internalism on the basis of my intuitively plausible stories. If I have not yet convinced you that internalism is false, I do not expect this to do it.

So suppose there is a note on the door. I didn't notice, but I should have. I am not justified in believing that I will have cream cheese for breakfast. Now suppose that there is no note on the door, and I am justified in my belief. It follows from our condition on non-identity of epistemic situations that these are not the same epistemic situation, despite the fact that I am the same on the inside in both cases. Could I

<sup>19</sup> It's not clear that the notion of direct relevance will help here. If someone argues that the physical cannot be directly relevant to the mental because we have privileged access to the latter but not the former, it's tempting to suppose that something has gone awry.

make a non-culpable mistake about the note in the first situation? No. If I believe there is one, I do not make a mistake. If I fail to believe something I should have known, then I am culpable. The only way to imagine a non-culpable mistake is by imagining me in a different situation.

If epistemic situations are understood epistemically rather than qualitatively, then external facts play as much role in determining what situation you are in as they do in determining whether or not you are justified. Perhaps not just any external fact can be directly relevant to your epistemic situation in this way. Perhaps only accessible facts can play this role. But here, the ordinary, commonsense notion of accessibility will do. The clearest case of an external fact that makes a difference to your epistemic situation is one you ought to know. If you ought to know that  $p$  in  $e$ , then you cannot make a non-culpable mistake about  $p$  in  $e$ , though you might make a non-culpable mistake about  $p$  in some other situation. That's why the external fact that  $p$  makes an epistemically relevant difference.

But if you ought to know that  $p$  in  $e$ , it is tempting to suppose that you *can* know that  $p$  in  $e$ . But this just means that you are in a position to know that  $p$  in  $e$ . In the first version of my story, I was in a position to know there was a note on the door. I just happened not to notice. When the note is in Sunny's pocket, I am no longer in a position to know, and, as a result, no longer obligated to know. Since we are in a position to know a great deal about the external world, this notion of accessibility does not give you internalism. But it does give you what the internalist wants. It gives you privileged access. The idea behind (JJ) is that if you are not justified in believing  $p$ , but you think you are, then you should have known that you are not. If the facts that make for justification were truly inaccessible, rather than merely external, we would not be able to say that you should have known better.

There is nothing mysterious about combining externalism about justification with the sort of privileged access that I have identified. One way to read (JJ) is to think of it as the claim that first-order justification is a necessary condition for second-order justification. If any external facts are relevant to the determination of the first-order justification, those very same facts are relevant to the determination of your second-order justification. If you want a model for this, ask any externalist about content. Of course we have privileged access to the widely individuated contents of our own thoughts. Whatever determines the content of the first-order thought also determines the content of the second-order thought.<sup>20</sup>

<sup>20</sup> See, for example, Burge (1988), or Gibbons (1996).

But do we have privileged access to the external world? Is that the premiss of the internalist's argument that I am rejecting? No, at least, we do not have the same kind of privileged access to the external world that we have to the facts about justification. To test (JJ) for truth, you have to evaluate the antecedent and consequent with respect to the same epistemic situation. This is why you do not have privileged access to whether someone else is justified or whether you were justified last week. This is also why the same set of external facts can be relevant to both first-order and second-order justification. But you have to evaluate (JJ) with respect to all possible epistemic situations. If the connection between first-order and second-order justification holds across the board, then (JJ) is true.

Let  $p$  be a proposition about the external world, perhaps the proposition that you have a hand or the proposition that there is no note on the door. I claim that there are some epistemic situations in which you couldn't make a non-culpable mistake about that. Any situation in which you ought to know that  $p$  will do. But the connection between justification and truth does not hold across the board. If there is any possible epistemic situation in which someone does make a non-culpable mistake, then justification does not entail truth.

So you can have privileged access to the facts of justification even if external facts are relevant to the determination of that justification. But with privilege comes responsibility. You are epistemically responsible for more than just what goes on in your own mind. You are epistemically responsible for certain features of the chunk of the world around you, but it will be different chunks and different features in different epistemic situations. Though you are not responsible for everything that is accessible to you, everything you are responsible for is accessible. This suggests the following supervenience thesis.

(AE) Justification supervenes on what you are in a position to know.

If you can know something on the basis of introspection, then you are in a position to know it. So internalism, at least Access Internalism, entails (AE). So internalists should not worry that (AE) is false. How far is (AE) from more familiar versions of externalism? There are differences, but it is not as far as you might think.

Most of us think we know that perception is reliable. If we are in a position to know that perception is reliable, then it is consistent with (AE) that this fact about the reliability of perception, not just our knowledge of this fact, is directly relevant to the justification of our beliefs. Suppose, on the other hand, that someone has a reliable faculty

of clairvoyance but is in no position to know that the faculty is reliable (BonJour 1980). Are the clairvoyance-based beliefs justified? Suppose, to make life difficult for the externalist, that they are not. This causes no trouble for (AE) since, according to the story, the fact about reliability is outside the supervenience base. Since the believer is in no position to know about the reliability, according to (AE), the reliability cannot be directly relevant to justification.

I do not pretend to have proven (AE). I present it here for your consideration as an alternative to internalism that is intended to capture some of its spirit. If you are moved by my stories, but you are still a fan of accessibility, (AE) may be the smallest reasonable step away from internalism. In order to have privileged access to the facts about justification, we need access, ordinary, everyday access to the facts that make us justified. If (AE) is true, then as far as justification is concerned, the basic epistemic line is around what you are in a position to know. Given the significance of this line, accepting the relevance of some external facts will not commit us to accepting the relevance of them all. Any fact we are not in a position to know is epistemically irrelevant.

There was a note on the door of my refrigerator. I didn't notice, but I should have. This is the thin end of the wedge. It shows that the traditional line around the mind is not as epistemically significant as the internalist suggests. Access Externalism is the most natural replacement for internalism. It includes an accessibility requirement on the facts that make you justified, and it is compatible with both strong and weak claims about our privileged access to the facts about justification themselves. Access Externalism is also compatible with the direct epistemic relevance of the reliability of perception, the proper functioning of our faculties, and other traditional externalist favourites, as long as we are in a position to know that these facts obtain. Though one of the edges of the wedge is thin, the natural replacement for internalism suggests that the internalists weren't just off by a bit.<sup>21</sup>

1042 Oldfather Hall  
University of Nebraska  
Lincoln, NE 68588  
USA  
jgibbons2@unl.edu

JOHN GIBBONS

<sup>21</sup> I would like to thank Alex Byrne, Al Casullo, Joe Mendola, Jim Pryor, Mark Van Roojen, Susanna Siegel, Ralph Wedgwood, Steve Yablo, and an anonymous referee for *Mind* for helpful comments on earlier versions of this paper.

## References

- Adler, Jonathan 2002: *Belief's Own Ethics*. Cambridge, MA.: MIT Press.
- Alston, William 1985: 'Concepts of Epistemic Justification' in Alston 1989, pp. 81–114. Originally published in *The Monist*, 68.
- 1986: 'Internalism and Externalism' in Alston 1989, pp. 185–226. Originally published in *Philosophical Topics*, 14.
- 1988: 'Justification and Knowledge' in Alston 1989, pp. 172–82. Originally published in *Proceedings of the World Congress of Philosophy*, 5.
- 1989: *Epistemic Justification*. Ithaca: Cornell University Press.
- BonJour, Laurence 1980: 'Externalist Theories of Empirical Knowledge' in *Midwest Studies*, 5, pp. 53–73.
- Burge, Tyler 1988: 'Individualism and Self-Knowledge'. *Journal of Philosophy*, 85.
- Chisholm, Roderick 1977: *Theory of Knowledge*. 2nd edn. Englewood Cliffs: Prentice Hall.
- French, Peter et al. 1980: *Midwest Studies in Philosophy* 5. Minneapolis: University of Minnesota Press
- Gettier, Edmund 1963: 'Is Justified True Belief Knowledge?' *Analysis*, 23, pp. 121–3.
- Gibbons, John 1996: 'Externalism and Knowledge of Content'. *Philosophical Review*, 105.
- Ginet, Carl 1975: *Knowledge, Perception and Memory*. Dordrecht: D. Reidel.
- Goldman, Alvin 1976: 'Discrimination and Perceptual Knowledge'. *Journal of Philosophy*, 73, pp. 771–91.
- 1999: 'Internalism Exposed'. *Journal of Philosophy*, 96, pp. 271–93.
- Greco, John 1990: 'Internalism and Epistemically Responsible Belief'. *Synthese*, 85, pp. 245–77.
- Harman, Gilbert 1980: 'Reasoning and Evidence One Does Not Possess' in French et al. (1980), pp 163–82.
- Kornblith, Hilary 1983: 'Justified Belief and Epistemically Responsible Action'. *Philosophical Review*, 92.
- Miller, Richard 1995: 'The Norms of Reason'. *Philosophical Review*, 104.
- Moore, G.E. 1962: *Commonplace Book: 1919–1953*. London: Allen and Unwin.
- Plantinga, Alvin 1993: *Warrant: The Current Debate*. Oxford: Oxford University Press.
- Pryor, James 2001: 'Highlights of Recent Epistemology'. *British Journal for the Philosophy of Science*, 52, pp. 95–124.

Williamson, Timothy 1995: 'Is Knowing a State of Mind?' *Mind* 104,  
pp. 533–565.  
———2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.