

In Defense of Serious Externalism

Thomas D. Senor
University of Arkansas

Section I: Introduction

For the past twenty five years, one of epistemology's greatest bones of contention has been the dispute between internalism and externalism. This despite the fact that every epistemologist acknowledges that knowledge has both internal and external components. Indeed, that it does is trivial. Truth is an external condition if anything is and everyone (or as close to everyone as we get in philosophical discussions) agrees that knowledge entails truth. Conversely, belief is an internal condition that, like truth, is entailed by knowledge. So what, one might naively ask, is all the fuss about? The answer is that the internalism/externalism debate in epistemology isn't primarily about *knowledge* but rather about the specifically epistemic component or components of knowledge. Here too, however, there is nearly universal agreement. Even an epistemologist with the externalist pedigree of Alvin Goldman has generally thought that one's background beliefs are relevant to the question of whether even a perceptual belief counts as knowledge.¹ For example, suppose that there are a pair of identical twins, Judy and Trudy, and that Alvin isn't good at visually distinguishing between the two of them (although he is able to tell that someone is either Judy or Trudy). Alvin is in a store, sees someone who matches his "Judy or Trudy" visual template, and for some reason comes to believe that it is Trudy he sees. Now even if Alvin is right, and even if unbeknownst to him, Judy died more than a year ago and so now anyone that he sees and would identify as Trudy really would be Trudy, Alvin doesn't know that it is Trudy he sees because he possesses a defeater for his belief: namely, his beliefs that he is unable to tell Trudy from Judy, and that it is equally likely that he sees one as that he sees the other. In short, Goldman will allow that one's background knowledge figures into defeating conditions even for a belief that is very reliably formed. Similarly, even the staunchest internalist (and here one can't do better than Keith Lehrer) will insist that if S knows that P, then the internally accessible justification of S's belief must genuinely have certain objective properties (like reliability, trustworthiness, etc) that link S's belief with the state of affairs that makes it true. Yes, on Lehrer's fiercely internalistic view one doesn't know that p (at least in a non-animal knowledge sense of the term) unless one's acceptance system contains a

network of acceptances that inferentially support p.² But knowledge requires not only that the inferential ties be there but that the acceptances from which they are weaved be true, and that the subject be trustworthy.

So if there is all this agreement, why has there been such a literature this past twenty-five years about the disagreement between internalists and externalists? Perhaps the problem concerns not knowledge but justification. That is a very plausible suggestion but, at least on first blush, it turns out not to be as clearly right as one might have supposed. Going back to the view of Goldman, a belief that is reliably formed but the justification of which is overridden not only fails to be an instance of knowledge but also of justified belief. So Goldman's view of justification incorporates an internalist component. Take now the evidentialism of Richard Feldman and Earl Conee: on their account the justification of a belief is a function of the total evidence the agent possesses.³ While this naturally provides a considerable internalistic component to justification, the fact that evidential relations are not subjective, that is are not a function of what the subject *believes* about them but are instead necessary relations between propositions, means that there is an important externalistic constraint as well on evidentialist justification.

So if both externalists and internalists think that there is an external and an internal component to knowledge, to what converts true belief into knowledge, and to epistemic justification, then where do they disagree?

Section II: The Nature of the Dispute

I think the best way to think about the dispute between epistemological internalists and externalists centers on what it is that generates positive epistemic status (since, as we've seen, one can be an internalist or externalist regarding any of a number of different epistemic properties). By "positive epistemic status" I mean any status that is positive from an epistemic perspective; examples of such statuses are "rational," "justified," "responsibly acquired" and "known." More specifically, we can get a better understanding of the fundamental issue that divides internalists and externalists by setting aside defeasibility considerations and the conditions under which an initially justified belief is defeated, and asking what a given theory says about what is required for a belief to have *prima facie* positive epistemic status, i.e., for a belief to have *ultima facie* epistemic status provided that it is undefeated. It is *prima facie* epistemic status that provides the positive lift (so to speak) required for knowledge, justification, etc. *Prima facie* justification (for example) is that which in the absence of justificatory defeat is sufficient for *ultima facie* justification.

To say that the way to mark the internalist/externalist distinction is to pay attention to what a given theory has to say about *prima facie* positive epistemic status is to make a good start at getting clear on the distinction but it is only a beginning. Given time constraints, I won't be able to do more than gesture at a couple points of complication, and then more or less stipulate how we'll think of the matter here. First, what sense of 'internal' and 'external' are we working with in these contexts? For instance, there is an obvious sense of 'internal' according to which psychological states of any kind are internal. So is an adequate internalist constraint one that requires of the epistemic ground that it be a psychological state? Or must the state be conscious? Or maybe "fairly readily accessible" to consciousness as William Alston proposed twenty years ago?⁴ Furthermore, is it sufficient that the ground be internal or must the fact that the ground is justification conferring also be (somehow) internal? Even assuming we resolve these matters how should we define internalism and externalism? Should we insist that they be contradictories? Or could we allow that a theory would satisfy both?

As I said, these are all big issues and a fuller discussion of internalism and externalism would have to address them all. But this is a Presidential Address and not a monograph, so I'll cut to the chase. Having made the *prima/ultima facie* distinction, one can then define externalism as the claim certain external conditions are necessary for justification, for turning true belief into knowledge, etc. On the other hand, if one insists that there is a crucial internalist feature that is necessary for *prima facie* positive epistemic status, then one is an internalist regarding the matter in question. So on this way of construing the matter, one might be both an internalist and an externalist since each position involves a positive requirement rather than a negative one. One might insist that for a belief to be *prima facie* justified, say, the belief must be grounded in a state that is a reliable indicator of the truth of the belief, and thereby be an externalist; yet she might also insist that the ground of the justification must yet be introspectively available to the subject and so in that way be internalist.⁵ In what follows, I will limit myself to the internalist/externalist debate as it regards epistemic justification—and not knowledge. Now there is a great deal of controversy and confusion about what 'justification' is supposed to denote, so let me be clear: by justification, I mean that property that, with true belief and a condition to rule out Gettier examples, is sufficient for knowledge; that is, it is what, with true belief, takes you in the neighborhood of knowledge even if it doesn't leave you at its door.

Here, then, is my suggestion: we should think of externalism as the

view that the epistemic justification of a belief depends at least in part on its externalist pedigree. That is, externalism entails that there is no collection of purely internal (i.e., introspectible) properties that are together sufficient for *prima facie* justification.

Section III: Three Grades of Externalism

Although there are other ways one might cut things up, I want to distinguish three varieties of externalism, beginning with the weakest and working toward the strongest. The first grade I'll dub "minimal externalism." It claims only that there is some aspect of justification that is not in available to introspection. Put more carefully, minimal externalism claims that there is a significant externalist constraint on *prima facie* justification. An example of this kind of externalism would be the view of Alston's as expressed in his paper "In Internalist Externalism" (Alston 1988). There Alston argues that S is *prima facie* justified in believing that p iff S's belief that p is based on an adequate ground, where one part of the adequacy of the ground is that, given it, the objective probability of the belief is high; a second condition of adequacy on grounds is that they be "fairly readily accessible" to introspection. Alston's view counts as minimal externalism because although it does contain a robust externalist condition that is not open to introspection, it also allows and even requires the satisfaction of a robust internalist condition.

The second grade of externalism is what I'll call "serious externalism." On this view, there is *no* internalist constraint on *prima facie* justification. For example, a so-called naïve reliabilism would claim that S is *prima facie* justified in believing that p iff p is produced by a reliable cognitive process. Unlike minimal externalism, Serious Externalism claims not just that there is an external necessary condition of *prima facie* justification but that there is an external sufficient condition as well. Because my goal here is to defend this position, I'll wait and elaborate the view when I consider objections to it.

Serious externalism is serious but not extreme. Extreme Externalism is the position that claims that the sufficient condition for *ultima facie* justification contains no internalist constraints. It's rather hard to find an example of extreme externalism about justification, but a close cousin of this view is Fred Dretske's information-theoretic account of knowledge.⁶ The example of Judy and Trudy discussed earlier can be used to show just how extreme extreme externalism is: in the example, you'll recall, Alvin knows identical twins Judy and Trudy well enough to be able to recognize of any woman that she is either Judy or Trudy, but he can't reliably distinguish between

the two. And while Alvin is aware that he has mixed them up on many occasions and that he has no reason to think his discrimination skills have improved in the intervening time, he nevertheless believes that he's now seeing Trudy (he's not the most careful of believers). Now unbeknownst to Alvin, Judy has met an untimely death. When Alvin next sees Trudy he believes that the woman he sees is Trudy. Since Judy is dead and so no longer a relevant alternative to his seeing Trudy, the signal Alvin receives carries the information that it is Trudy he is seeing and so he *knows* that he sees Trudy.⁷ This despite the fact that if he were rational and he were to carefully reflect on his past lack of success in distinguishing Trudy and Judy, he would think that his visual experience was equivocal and so he would not come to believe that it was Trudy that he sees. Although Dretske thinks that justification has nothing to do with knowledge, we can still see this is as a kind of extreme externalism, albeit about knowledge rather than justification.

Section IV: Why be an Externalist?

Why would someone be an externalist of any stripe? There is, I believe, one motivation that underlies all of these grades of externalism (although that motivation might better support some varieties than others). I said earlier that the epistemic property that I was interested in was that which with truth, belief, and an anti-Gettier condition, was sufficient for knowledge. I follow the tradition and call this "epistemic justification" but I won't pretend that this is any more than a stipulative characterization. Now what sets various notions of epistemic evaluation apart from other types of evaluation is not their domain (primarily, that would be beliefs, although doxastic processes, practices, and perhaps even assertions count too) but their conceptual content. As a great many others have noted, what makes an evaluation "epistemic" is its connection to our twin epistemic goals of believing propositions that are true and not believing propositions that are false. A term that appraises a belief (say) for its success or failure relative to these goals is a term of epistemic evaluation. There are various ways a belief can be evaluated even given these epistemic goals: for instance, it might be the best that the person could do in the circumstances, it might be what the agent nonculpably believed was the best she could do in the circumstances, it might meet some threshold of satisfying the standards, it might maximize the standards, etc. Now this is not the place to take a detailed and well-argued for stand on precisely which evaluation is the right one for my stipulative characterization of epistemic justification. But it is clear that if justification is to be what I said, then it must be quite strongly

truth-conducive. That is, if a belief is to satisfy the condition of being what, in the absence of Gettier conditions, is with true belief sufficient for knowledge, then that belief must not only have a strong subjective link to truth but an objective one as well. Maybe it is formed by a reliable process, maybe it is grounded in a state that is a reliable indication of the truth of the belief, maybe the belief carries the information that the belief is true, maybe it is grounded in veridical and not-misleading evidence, etc. I'm not here arguing for any particular brand of externalism but only suggesting that where it is epistemic justification so circumscribed that is on the line, any belief deserving of that grade of evaluation must bear a significant, objective link to truth. And this kind of objective link to truth is an externalist connection extraordinaire.

Section V: General Motivations for Rejecting Externalism

So much for the motivation for externalism. Why would anyone deny it? That is, what reason might there be for claiming that there are *no* externalist conditions to justification, that all factors relevant to justification must be internal (where this comes to something like 'introspectible')? If one thinks of the epistemic rules governing justification as being action-guiding norms, one will naturally suppose that only aspects of belief formation that are appropriately available to the subject are relevant. Such an internalism might be explicitly responsibilist—that is, construe justification in terms of the responsible conducting of doxastic affairs. Or she might eschew deontic language and think of justification as rule following taken neat. To be unjustified in holding a belief is to fail to follow the appropriate rules but this doesn't imply irresponsibility on the part of the believer. Another possible motivation for the rejection of any externalist constraint on justification is a commitment to what we can call epistemic perspectivalism. The perspectivalist insists that what is relevant for justification is (speaking somewhat loosely here) how the world looks from the agent's perspective. One is justified in a belief only to the extent that she has reason for thinking that it is true. Evidentialism of the sort propounded by Richard Feldman and Earl Conee is perhaps the prime example of the perspectivalist's position: a belief is justified only if it is supported by one's total evidence (where one's body of evidence includes beliefs, perceptual/memory states, etc.).⁸

Having distinguished these two motivations for anti-externalism, I want to put them aside for now. We'll return to them later in the paper.

Section VI: The New Demon World and Clairvoyance Objections to Serious Externalism

Since this is epistemology and positions are often rejected not for deep, theoretical reasons but because they are thought to be prone to counterexamples, I suspect that the main reason that externalism has been rejected is the conviction that there are counterexamples that show that external conditions are neither necessary nor sufficient for justification. And if these examples do what they are purported to do, that will amount to a joint case for the anti-externalist position that externalist factors (at least of a certain sort) are irrelevant to questions of justification. So we will now look at these (well known) counterexamples and see what means the externalist has for responding to them.

Before exploring these two examples, let me say again that the type of externalism I'm fixin' to defend is the *serious* variety; it claims that whether a belief is *prima facie* justified is purely a matter of its objective, non-introspectible properties. The position, then, is that qualities that are near and dear to the two types of internalist views we've considered—namely action-guidance properties and evidential relations—are neither necessary nor sufficient for epistemic justification (at least regarding perceptual beliefs).

There are two primary counterexamples to Serious Externalism: The New Demon World Objection and Laurence Bonjour's Norman-the-Clairvoyant Case. I'll begin by presenting them; after that we'll deal with them in turn.

VI.A: The Objections Stated

The New Demon World: This familiar objection is aimed at all varieties of externalism enumerated earlier.⁹ Consider for a moment a cognitive doppelgänger of yourself who is the inhabitant of a world governed by a Cartesian demon. He or she has just the perceptual and introspective experience you have at the actual world, and inasmuch as is possible, she forms just the beliefs you form in our world (if it helps, you might suppose that your twin has only recently been transported from our world to the demon world). Now consider a given perceptual belief of yours that you take to be as sure as any belief you have that involves the external world—maybe the belief that you have a body or that there is a table in front of you now. Your twin at the demon world has exactly the same history of experiences and beliefs you do, follows the same rules and norms you follow, and she believes that she has a body and that there is a table in front of her. But in her case, the belief is not only false but is unreliably formed since she is the victim of systematic deception. The intuition the proponents of the New Demon World invite us to share is that any belief you are justified

in having is a belief your twin is justified in having. But since her beliefs are emphatically not the result of reliable processes (and are presumably lacking any other suitably externalist feature), then that property can't be relevant to the justification of your beliefs (even if they turn out to be reliably formed). Since what you and your twin share are internal features of the formation and maintenance of your belief, and since there is justificatory parity between the two of you, then no non-internal aspect of your belief can be relevant to its justification. In a word, external factors are not necessary for justification.

That, then, is the New Demon World Problem. The second prominent counterexample is BonJour's famous case of Norman the clairvoyant.¹⁰ Unbeknownst to Norman, he has just come to have a highly reliable clairvoyant cognitive process and it has just produced in him the belief that the President is in New York. Norman has no evidence for this belief, no reason to think that he is clairvoyant, and no reason to believe that there is such a process as clairvoyance. He merely has the process and it has produced in him the belief about the President, and produced it in such a way that it is highly likely to be true. BonJour claims that an externalist of the reliabilist variety will have to judge the belief to be justified, since it can satisfy however strict an external standard one cares to impose. Yet, BonJour and a chorus of others maintain, it is clear that the belief is not justified. But if not, then the externalist condition is not sufficient for justification. So taken in tandem, the New Demon World and the Norman cases seem to many to show that external considerations are neither necessary nor sufficient for justification. Indeed, they are irrelevant.

IV:B Replies to the Objections

What is an externalist to say to these cases? Let's consider the New Demon World first. It is no exaggeration to say that it has been taken with great seriousness by some prominent externalists (Goldman changed his view a couple of times in light of it). Yet it seems to me that the case can be handled rather quickly. Externalism about justification is a family of theses about that which, together with true belief and an anti-Gettier clause, is sufficient for knowledge. To be justified, then, is to be in an objectively strong evaluative position. Which objectively strong evaluative position? Well, one that while being fallibilist (hence the need for the anti-Gettier condition) is nevertheless strongly linked to our attaining truth and avoiding error. Given this characterization of what justification and its conceptual role in the explication of knowledge come to, it is clear that a perceptual belief at a demon world will fail to achieve this status. For the demon is doing a

fine job of making sure that the subject's perceptual beliefs are false. So her beliefs shouldn't receive a positive evaluation with respect to how well she is able to comply with her epistemic goals of truth acquisition and error avoidance. No doubt if we keep in mind standard perceptual beliefs of the sort that we take ourselves to be justified in believing, we can find other epistemic virtues that they embody. They may be responsibly acquired, they may have been formed in accordance with the correct epistemic norms, they may fit her evidence, etc. But they will not be justified in the sense outlined above, and surely it is that sense that has always been the target of the externalist. So not only am I willing to deny the intuition that demon world subjects are justified in their standard perceptual beliefs, but I think it is had only by conflating distinct concepts of epistemic appraisal.

This isn't to say that I think there is no lesson to be learned about justification from these cases. But rather than the falsity of externalism, what the New Demon World shows is that justification doesn't supervene on the cognitive states that I share with my twin at a demon world. No matter how you beef up the internal states of subjects, so long as all the relevant properties could be had by my deceived doppelganger, you don't have a condition that is sufficient for even *prima facie* justification. For my believing a proposition not only fails to have a strong connection to its truth but is instead a virtual guarantee of its falsity.

So much for the demon world. What shall we say about Norman? As I've described the case (and indeed as BonJour has described it), we don't know much about Norman's background beliefs. He doesn't have evidence either that he's clairvoyant or that there is even such a thing as clairvoyance. But this doesn't help much. Is Norman like the average, moderately-educated American who is generally skeptical of claims of clairvoyance, but who hasn't really looked into the matter much? For that matter, does Norman think that his belief is formed by some new, odd faculty? Or does he just find himself believing that the President is in NY without any idea how that belief arose? These missing details would seem to make a rather large difference in how the case comes out. For if there are the makings of a good argument for either the falsity of the belief about the President or the unreliability of its ground, then the Serious Externalist can agree with BonJour that the belief is not *ultima facie* justified because of the presence of a justificatory defeater. So for example, if Norman reflects on his belief about President's whereabouts, is quite sure that he hasn't read or heard anything today (or any other day) that would have produced that belief, and comes to believe that that belief is utterly groundless, then he is clearly *ultima facie* unjustified in holding it. Yet his being *ultima facie*

unjustified doesn't begin to show that Serious Externalism is wrong since it is not a thesis about *ultima facie* justification but only about *prima facie* justification. The Serious Externalist can maintain that the belief is *prima facie* justified but defeated by the considerations evinced above.

In order for the Norman case to be a threat to Serious Externalism it must be clear Norman's lack of *ultima facie* justification is not due to epistemic defeat but instead to a lack of *prima facie* justification. And it's not easy to see how to mess with Norman's background beliefs to get this result. For instance, if Norman has a belief that there are only certain well-established modes of reliable belief formation, and that his belief about the President wasn't formed by any of them, and he has no independent reason for thinking the President is in NY, then he has a reason for doubting his belief and his lack of *ultima facie* justification can be thought, again, to be a function of the presence of defeat rather than the absence of *prima facie* justification. So we need a case in which there is neither an internalistically-accessible justifier nor defeating background evidence in order to have a case against Serious Externalism. In order to do this full justice, it seems to me, we have to suppose that it is an open question for Norman whether there are new and reliable processes popping up from time to time, and that he has no particular opinion about whether he might come to have some new and reliable process. In such a case, he has a belief about the President, he has no reason to be dubious of it (since if he thought he about it, he might reasonably think that this might be produced by a new, reliable process), but he lacks a positive internalistically accessible justifier.

The Serious Externalist will say that if this is how we are to understand the Norman case, then the subject is not only *prima facie* justified but is *ultima facie* justified as well. Because the belief is grounded in a *prima facie* justification-conferring process (i.e., one that puts one in a strong, objective veridical position), the belief will be *ultima facie* justified if it is undefeated. And if I've been successful in spelling out the details of the case and it is plausible that Norman has no defeater for his belief, then we have the kind of example that BonJour was trying (or should have been trying) to develop. I submit that in this revised case Norman is *ultima facie* justified in his belief that the President is in NY. Furthermore, I don't see why the either the responsibilist or the perspectivalist internalist should disagree. Why should we think that Norman is being irresponsible in his belief about the President? His belief isn't the result of listening to those he knows to be epistemically disreputable or drawing a conclusion from premises that, if he only thought a little harder about them, he would see don't actually support the conclusion. He finds that he believes the President is in NY,

on careful, rational reflection he sees has no reason to doubt it and so he continues believing as he does. Why is that irresponsible?

I don't think matters change much if we consider the perspectivalist position. Norman reflects on his cognitive situation, examines his beliefs and perceptual states and finds that he has a positive conviction that the President is in NY and has no reason at all to deny it. Why shouldn't we take this as confirmation that his perspective supports it? True, if we subtract the belief in question then he has no positive reason to accept it. But to insist that his internal perspective must positively support the belief in question is to accept but one, somewhat restrictive form of perspectival internalism that would seem to have problems explaining the positive epistemic status of a good many of our memory and *a priori* beliefs.

I conclude, then, that we have yet to see a good reason to reject Serious Externalism. Not only do the standard counterexamples fail to provide grounds for rejecting it, but it is also consistent with at least some of the more general motivations some epistemologists have had for being internalists.

Notes

¹ The only phase of Goldman's thinking about knowledge in which he (apparently) didn't think background beliefs were of potential defeating significance is Goldman 1967.

² See Lehrer 2000

³ See Feldman & Conee 2004

⁴ See Alston 1988

⁵ This is Alston's position in Alston 1988.

⁶ Dretske 1981

⁷ Dretske has confirmed via correspondence that he accepts this as a case of knowledge.

⁸ Feldman & Conee 2002

⁹ The New Demon World seems to have been thought up independently by Keith Lehrer and Steward Cohen in Tucson and Richard Foley in South Bend. See Lehrer & Cohen 1983, Cohen 1984, and Foley 1985.

¹⁰ BonJour 1985

Works Cited

- Alston, William, "An Internalist Externalism," in *Synthese* 1988, pp. 265-283.
 BonJour, Laurence, "Externalist Theories of Empirical Knowledge," *Midwest Studies in Philosophy* 5, 1980, pp. 53-73.

Thomas D. Senior

- Cohen, Stewart, "Truth and Justification" in *Philosophical Topics* 1984, pp. 279-95.
- Dretske, Fred, *Knowledge and the Flow of Information* (MIT Press, 1981).
- Feldman, Richard & Conee, Earl, *Evidentialism* (Oxford, 2004).
- Foley, Richard, "What's Wrong with Reliabilism?," *The Monist* 1985.
- Goldman, Alvin, "A Causal Theory of Knowing," *Journal of Philosophy* 64, 1967, pp. 357-372.
- Goldman, Alvin, "What is Justified Belief?" in George Pappas, ed., *Justification and Knowledge* (1979), D. Reidel.
- Lehrer, Keith *Theory of Knowledge* (Westview, 2000).
- Lehrer, Keith and Stewart Cohen, "Justification, Truth, and Coherence," *Synthese* 1983, pp. 191-207.