

## Epistemic Internalism, Philosophical Assurance and the Skeptical Predicament

### *Introduction:*

It is a particular pleasure to contribute this paper to a volume honoring Al Plantinga. I have always viewed his work as a model of how to do philosophy and I have learned a great deal from him over the years. It is as a result of philosophical conversation with both Plantinga and his former student Michael Bergmann that I have come to the (always painful) conclusion that I need to revise some of the things that I have said in print. This paper is an attempt to do just that.

In *Metaepistemology and Skepticism* I implied that the fact that externalists, to be consistent, should allow “track record” arguments in support of their belief that they have first-level justification<sup>1</sup> is a kind of *reductio* of their position. I said the following:

You cannot use perception to justify the reliability of perception! You cannot use memory to justify the reliability of memory! You cannot use induction to justify the reliability of induction! Such attempts to respond to the skeptic’s concerns involve blatant, indeed pathetic, circularity. (177).

---

<sup>1</sup> Throughout this paper I’ll be talking about justification. I’m convinced that what I call justification is the same thing that Plantinga calls warrant. I certainly don’t attach to the concept of justification normativity of a sort that Plantinga successfully argues is irrelevant to knowledge. In any event, those happier with talk about warrant may translate my remarks into that terminology if they choose.

The above still seems right to me and, I hope, will strike you as plausible.<sup>2</sup> If one embraces some version of externalism such as reliabilism, if one embraces the view that the reliability of a belief-producing process is *sufficient* to generate justified output beliefs (provided that the input beliefs, if any, are justified), then one will need to embrace what I take to be the absurd view that one can use reliable methods of forming belief to justify belief that they are reliable (and thus, that their output beliefs are justified). As I noted in the book, even many externalists seem to get cold feet when it comes to bootstrapping their way up to justified metabeliefs about reliable processes using those very processes. Alston (1993), for example, specifically rejects the legitimacy of track record arguments. Plantinga is harder to read on this issue, but it is interesting to note that in his (2000) defense of Christianity he is certainly hesitant about claiming to have inspired warranted belief that he has inspired warranted belief.

If I had stopped with the remarks quoted above, I wouldn't have anything to retract. I went on, however, to generalize (always a dangerous move in philosophy):

The fundamental objection to externalism can easily be summarized. If we understand epistemic concepts as the externalist suggests we do, then there would be no objection in principle to using perception to justify reliance on perception, memory to justify reliance on memory, and induction to justify reliance on induction. *But there is no philosophically interesting concept of justification or knowledge that would allow us to use a kind of reasoning to justify the legitimacy of using that reasoning.* (180, emphasis added)

As Michael Bergmann (2000) has pointed out there seems to be (at the very least) some tension between what I said above and what I also said about direct acquaintance, a relation that is crucial to my own internalist account of what constitutes noninferential justification of a sort that satisfies philosophical curiosity. On the view I defended, one has noninferential justification for believing P when one has the thought that P while one is directly acquainted with the fact that P and the correspondence between the thought and the fact.<sup>3</sup> When answering the hypothetical question of how I would justify my belief that there *is* such a relation of acquaintance, I answered as follows:

If I am asked what reason I have for thinking that there is such a relation as acquaintance, I will, of course, give the unhelpful answer that I am acquainted with such a relation. The answer is question-begging if it is designed to convince someone that there is such a relation, *but if the view is true it would be unreasonable to expect its proponent to give any other answer.* (77, emphasis added).

---

<sup>2</sup> For an excellent discussion of some of these issues, one which helped me understand my own views better, see Cohen (2002).

<sup>3</sup> I also suggested that it might be possible to have a noninferentially justified belief based on acquaintance with a fact very similar (but not identical with) the fact that P.

Well, what is good for the goose is good for the gander. Why shouldn't a reliabilist, who has, of course, offered a quite different account of noninferential (and inferential) justification, respond in precisely the same way? If the reliabilist is asked what reason he has to believe that there are belief-independent, unconditionally reliable processes, why shouldn't that reliabilist respond by claiming to have a reliably formed belief that they exist, a belief whose reliability involves relying on a successful track record of using those very processes. The answer is question-begging if it is designed to convince someone that there are such reliable belief-forming processes, but if the account of justification is correct, what other answer would one expect to get from a proponent of the view?

In what follows, I want to explore what I take to be a more felicitous way of presenting the alleged *reductio* against externalism. Our discussion will seek to get clear about precisely what internalists want and why it is that they believe that externalists do not succeed in analyzing a philosophically satisfying concept of knowledge or justification.

### *What do Internalist's Want?*

In a number of important books and papers, Barry Stroud has tried to make clear what a philosophically satisfying account of knowledge must produce.<sup>4</sup> He usually puts the goal of the philosopher in terms of understanding. Our task as philosophers is to develop a philosophically satisfactory understanding of knowledge in general, or knowledge within a certain specified domain. While he concedes that one can, of course, *scientifically* study human cognition in the same way that one can study any other natural phenomenon, he argues that scientific investigation could never yield results that satisfy the philosopher. That is because there are certain constraints on how one can legitimately study or investigate knowledge *in general*. But what are those constraints according to Stroud? As far as I can tell, understanding human knowledge involves coming to know what we know and how we know it. But to be philosophically satisfying (particularly from a first person perspective) our investigation must meet the following conditions:

1) In trying to understand whether and how we know various propositions in a given field of knowledge, we cannot presuppose that we know or even reasonably believe any propositions alleged knowledge of which we are investigating. As a result, we are forbidden from employing as *premises* any proposition knowledge of which we are trying to understand. So if we are trying to figure out how, if at all, we know propositions about the past, for example, we couldn't use as a premise in reaching our conclusion any truth about the past.

2) In trying to understand how we know various propositions in a given field of knowledge, we cannot presuppose the legitimacy of any of the *methods* we employ in coming to believe propositions of the sort in question, and therefore cannot use any of those methods in studying the knowledge in question. So, for example, if we are trying to figure out how we know

---

<sup>4</sup> These include Stroud (1984), and a number of papers contained in Stroud (2000).

propositions about the external world through perception, we cannot use perception to facilitate our understanding.

There is something very seductive about the above constraints. Historically, philosophers who have taken the problem of skepticism seriously seem to have just taken for granted 2). It was viewed as a the worst sort of question-begging to attempt an inductive justification of induction, or a perceptual justification of the veridicality of perception. Although the problem of memory was not discussed nearly as often, it would presumably be equally illegitimate to employ memory in the attempt to certify the legitimacy of relying on memory.

It is not hard to see why Stroud finds skepticism so difficult to avoid given the above constraints. If the epistemologist's ultimate goal is to understand *all* knowledge, knowledge *in general*, and to do so within the constraints posed by 1) and 2), it doesn't take a pessimist to see clouds on the horizon. To understand knowledge *in general* we would need to satisfy ourselves that *all* of our methods of arriving at conclusions are legitimate and we would need to do so without using *any* of those methods! Even if we were to arrive at a purely a priori knowledge of the legitimacy of epistemic principles, we would have left philosophically mysterious a priori knowledge—there would still be one source of knowledge that we haven't been able to study philosophically. In a striking comment that really led me to think about some of these matters afresh, Plantinga (2001, 390) suggested that if the internalist insists on something like 2) in the quest for epistemic security, then even if there were a God, that God would be unable to have knowledge of the sort the internalist wants. That seems right to me.

#### *More Modest Internalist Goals:*

I have never been comfortable with the emphasis many internalists place on the importance of access to knowledge and justification. As I said, Stroud often seems to locate the epistemologist's target as second-level knowledge (or understanding). Other internalists seem to think that having first-level knowledge or justification is inseparable from having second-level knowledge or meta-justification. But as I have argued elsewhere global access internalism seems to raise immediately the specter of vicious regress.

We should make at least the following distinctions among various sorts of access requirements for knowledge and justification. (In what follows I'll focus on justification, but what I say will apply *mutatis mutandis* to knowledge.)

#### 1) Global Actual Access Internalism:

The global actual access internalist claims that in order for S to be justified in believing P, S must have access to the fact that he has that justification. The most natural interpretation of "access" here is knowledge or justified belief. Internalists typically also want the meta-justification in question to be noninferential or introspective. So the claim is that in order for some set of conditions J to constitute S's having justification for believing P, S must have an introspectively justified belief that J exists. The modal status of this claim is crucial if it is to be even intelligible. It seems hopeless to argue that this principle is an analytic truth. It doesn't even make sense to suppose that J, *by itself*, constitutes S's having justification for believing P only if

one *adds* to J S's access to the fact that J obtains. That's tantamount to claiming that J constitutes S's justification for believing P only if doesn't really constitute S's justification for believing P! It is really only J + access to J (call that A) that constitutes S's justification for believing P. But of course (J + A) doesn't really constitute S's justification for believing P either. Global actual access internalism implies that one must have access to (J + A) call that access A\*. But (J + A + A\*) won't constitute S's justification either for one must have access to that condition... , and so on *ad infinitum*.

To escape this problem the global access internalists must claim that their principle is some sort of synthetic necessary truth. Given what genuine justification is, there is a necessary connection between possessing it and realizing (knowing or justifiably believing) that one possesses it. Just as P's being true implies that it is true that P is true, even though its being true that P is true is not, presumably, *constitutive* of P's being true, so the access internalist we are discussing might hold that S's having justification implies that S is aware of that justification without that second-level awareness being constitutive of the first-level awareness. While this sort of strong access requirement might not lead to *conceptual* regress, it might still seem to lead to *vicious* regress. Every justified belief requires the having of an infinite number of ever more complex meta-beliefs.

## 2) Global Potential Access Internalism:

It might seem that the specter of regress is less ominous if we shift from a requirement that having justification requires actual access to that justification, to the claim that having justification requires only potential access to it. Again, most internalists want the potential access to be introspective (or noninferential). There are also a number of different interpretations of the critically important concept of potentiality upon which the view relies. Just as with actual access internalism, it is still important that the potential access internalist not regard second-level potential access requirements as *constitutive* of first-level justification. Rather, to avoid conceptual regress the claim must be that having justification is by its nature *tied* to something else, the possibility of accessing (introspectively or noninferentially) that justification. Here one can see the natural connection between what one might call internal state internalism, and potential access internalism. The internal state internalists takes a person's having justification for a belief to be an internal feature of the person.<sup>5</sup> Internal states might in turn be viewed as nonrelational properties of mind. Traditional foundationalists have often taken it to be a mark of

---

<sup>5</sup> We must be careful to recognize here the distinction between having justification for a belief and a belief's being justified (a distinction I will return to later in the paper). There can be justification for me to believe some proposition P, but unless I base my belief that P on that justification, most philosophers will deny that the resulting belief is justified. The basing relation is often construed as causal, and most philosophers will reject the idea that the relation of causing is in any sense purely internal to the person whose beliefs are caused. So to be consistent, internalists shouldn't claim that a belief's being justified is a purely internal matter, though they could consistently claim that there being justification to believe a proposition is a purely internal matter.

nonrelational properties of mind that they be introspectively accessible. Thus we can see how an internal state internalist might also end up being a potential access internalist.

While the threat of regress on potential access internalism might seem less severe, it is nevertheless present. To be sure, having justification for believing P doesn't require having an infinite number of ever more complex higher-level beliefs. But it still seems to require the possibility of forming infinitely many ever more complex higher-level beliefs. God might be up to that task, but it is not clear that you and I are.

### *Philosophical Assurance:*

I think it would be a mistake to dismiss access internalism solely on the grounds that it is fertile ground for skepticism. The philosophical enterprise is by its nature odd. Philosophers ask questions about that which is simply taken for granted by non-philosophers. Those of us who are parents remember fondly the also sometimes frustrating days when our young children would ask a seemingly endless number of "Why" questions. "Why is the sky blue?" "You give some sort of answer. It's X and X things appear blue? Your kid then wants to know why X things look blue? The looming regress of "Why?" questions inevitably ends with an impatient parent responding "That's just the way it is," a response that, no doubt, did little to satisfy the child's curiosity. The epistemologist, Stroud argues, wants to know why we can legitimately conclude that a certain way of forming belief is legitimate, and the epistemologist's philosophical curiosity isn't going to be satisfied by being told at *any* stage of the game that it just is. It's always possible then that the epistemologist is led by philosophical curiosity on a quest that can only end in failure. But if access requirements so obviously lead to skepticism, one might want very strong reason to believe that they are independently plausible.

### *Access Requirements and Defeaters:*

In a recent paper Michael Bergmann (2002) tries to explain why higher-level requirements for having justification can *seem* so plausible despite the fact that the externalist's rejection of them is in the end correct. His suggestion, in short, is that *when one considers* the question of whether or not one's belief is reliably produced (or produced in a trustworthy way), *then* the fact that one disbelieves in that reliability, or even withholds belief with respect to reliability, *does* render unjustified the lower-level belief in question. He argues, however, that disbelief and even withholding of belief requires entertaining the proposition in question. With respect to any given proposition, one can believe, disbelieve, withhold belief, or do none of the above. According to Bergmann, if one simply fails to consider the question of whether one's lower-level beliefs are reliably produced, their reliable source may render them justified even in the absence of any doxastic attitudes concerning their source.

In response to Bergmann, I questioned whether the mere subjective attitude (even if wildly irrational) a person takes to the legitimacy of a belief-forming process should affect the epistemic status of the belief formed by that process. But one might plausibly argue for a somewhat different view about potential defeaters. Let's make the traditional distinction between there being justification for a belief and a belief's being justified. There can be justification for S to believe P even if S doesn't believe P, or believes P, but not as a result of

possessing justification (not by *basing* the belief on the justification possessed). By contrast, to have a justified belief that P, there must not only be justification for the belief, but the belief must be *based* on that justification. Now whether or not an aging person believes that his beliefs based on memory are now unreliable, and whether or not the person in question even *considers* the question of whether his beliefs based on memory are now unreliable, shouldn't we conclude that if there is good reason for that person to believe that the beliefs are unreliably produced, that defeats whatever justification he might otherwise have had for believing propositions about the past based on memory?

Of course, if we succeed this way in securing the relevance of metalevel justification concerning the *unreliability* of a belief-forming process to lower-level justification of the belief produced through that process, you know what is coming next. As Bonjour argued in *The Structure of Empirical Knowledge*, it would seem plausible to suppose that justification for *withholding* belief concerning the justificatory status of a lower-level belief (or the legitimacy of the process that produced it) is equally relevant. With or without my *consideration* of the question of whether my belief that p is justified, if it would be rational to withhold belief with respect to my having justification for believing p, i.e. if the rational thing to conclude is that it is no more likely than not that my belief that p is justified, surely that fact defeats whatever justification I might otherwise have had for believing P? But haven't we now got all the way to metalevel requirements for justified belief? These metalevel requirements involve justification at the higher levels--they are not particularly concerned with what the person actually believes at the higher levels.

Bergmann argues that there are four possibilities with respect to the higher-level doxastic attitudes one takes towards the appropriateness of a lower-level belief. One might 1) believe that the belief is appropriately formed, 2) disbelieve that the belief is appropriately formed, 3) withhold belief with respect to whether or not the belief is appropriately formed, or 4) have no attitude whatsoever with respect to whether or not the belief is appropriately formed (because one hasn't even considered the matter). But it seems to me that there are only *three* possibilities with respect to the epistemic justification there is for one to adopt higher-level attitudes towards the justification of lower beliefs. There might be 1) justification for believing the proposition that one's belief that p is justified, 2) justification for disbelieving the proposition that one's belief that p is justified, or 3) justification for withholding belief with respect to whether or not one is justified in believing p. It is very tempting for the internalist to argue that either 2) or 3) defeats one's justification for believing p. Therefore unless there is justification to suppose that one's belief that p is justified, one is unjustified in believing p.

The above argument is very attractive and does provide at least *prima facie* plausibility for a justificatory access requirement even if that requirement leads to skepticism. But there is still the worry that we are being led too quickly to a requirement for justification that simply ensures at the outset a victory for skepticism. There is still the concern that as we move up levels we get to propositions so complex that it will be impossible for creatures like us to even entertain them. If we can't entertain a proposition, then it is not clear in what sense there can be justification *for us* to believe the proposition. Borrowing again from Bergmann's idea, however, we might be able to explain why it is so tempting for traditional foundationalists to suppose at least that whenever they have a first-level *noninferentially* justified belief, there is also noninferential justification for them to believe that they have that first-level justification.

Whenever they succeed in *raising* the question of whether they have lower-level justification for believing a certain proposition, they *will* find noninferential justification for believing that they have noninferential justification for believing that proposition. That is because their ability to raise the question presupposes that they *are* capable of entertaining at least that metalevel proposition describing justification for a belief about justification. The second-level acquaintance with the first-level acquaintance that is partially constitutive of first-level justification is typically available when one has the conceptual sophistication to entertain the proposition made true by the fact that one has second-level acquaintance with first-level acquaintance. When one is directly acquainted with one's pain in a way that yields noninferential justification for believing that one is in pain, and one has the ability to formulate the question of whether one is justified in believing that one is in pain (formulation of which involves entertaining the relevant proposition), one will find oneself directly aware of the fact that one is directly aware of the pain.

Of course, one shouldn't infer from the fact that whenever one looks for something one finds that it has a certain characteristic, that it necessarily has that characteristic. One shouldn't infer from the fact that whenever one looks for justification for believing that one has a noninferentially justified belief (when one does) one finds the relevant higher-level justification, that there is a necessary connection between having the lower-level justification and having justification for believing that one has it. Consider, by analogy, a pathetic argument for an extreme sort of anti-realism--the view that there is no reality that does not necessarily involve a representation of it. "Go ahead," the anti-realist argues, "give me an example of an unrepresented fact." You won't be able to do so, of course. As soon as you choose some fact as an example, you will have thereby represented it. It should be obvious to everyone, however, that from the tautology that all represented facts are represented, it does not follow that it is necessary that all facts are represented. Similarly, even if it were necessarily the case that whenever we consider the question of whether we have noninferential justification, we find that we are justified in believing that we have it, it doesn't follow that there is any necessary connection between having a noninferentially justified belief and having justification for believing that we have that justification.

There is another obvious fact that complicates the issue. In trying to *discover* from the first person perspective what the connection is between having a certain kind of justification for believing P and having justification for believing that we have that justification, we need to start with uncontroversial examples of justified beliefs. But what makes a belief an uncontroversial example of a justified belief is presumably the fact that we have strong justification for believing that it is justified. So again it will be a trivial truth that whenever we have an *uncontroversial* example of a justified belief we will find that we have justification for believing that we have a justified belief.<sup>6</sup> That analytic truth, however, will not secure a necessary connection between having justification and having access to that justification.

Seeking philosophical assurance by moving up levels seems destined to result in

---

<sup>6</sup> This is a problem from the first-person perspective. One can still try to generate examples of someone's having justification without having the relevant meta-level justification by looking at second- and third-person ascriptions of justified belief.



disappointment. We will either get to a point at which we can no longer formulate the relevant question because it has become so complex, or we will simply get frustrated or bored and abandon the project. Again, one might conclude that epistemologists are simply doomed by the nature of their philosophical quest to a life of philosophical disappointment. But is there any other way for the internalists to succeed in their search for philosophical assurance?

*The Source of Philosophical Assurance:*

Why do I think that one can't use memory to justify the legitimacy of using memory and perception to justify the legitimacy of using perception, but I do think that one can use acquaintance with acquaintance to justify the existence of acquaintance? The answer, in one sense, simple. On the view I accept, facts about what we are acquainted with are by themselves sufficient for having philosophically relevant justification; facts about what we are caused to believe (reliably or not) by memory and perception are not. Autobiographical reports are all well and good, the externalist will reply, but what's that got to do with philosophical argument. The externalists (who are usually also foundationalists) are perfectly happy with their own accounts of (external) conditions that they claim are sufficient for justification.

What I want to suggest is that one should test the plausibility of a claim about what is *genuinely* sufficient for having justification by exploring the implications of that claim when moving up levels. Specifically, as I suggested in the quotes with which I began this paper, it seems to me that reliabilists, for example, *ought* to have no qualms about using a way of forming a belief to justify one's belief that that way of forming beliefs is legitimate. Either the reliability of the belief-forming process is enough, by itself, to yield justified output beliefs or it is not. If it is, then it is no matter what level of belief one is interested in justifying. So if memory and induction are reliable, then through memory and induction I can justify my belief that memory is reliable. I remember seeming to remember doing certain things and I also remember doing them. If induction is a reliable way of forming beliefs about generalizations, I can conclude on that basis that my beliefs about the past based on memory are reliably produced and thus justified. As I said before, it is striking that even many proponents of reliabilism can't quite bring themselves to argue that this is a legitimate way to justify belief that memory is reliable. To be sure, they might argue that *if* memory is reliable then we can form justified beliefs about the reliability of memory this way, but they feel uncomfortable simply asserting that they have justified belief about the reliability of memory formed in this way. Why? Because at some level they realize that in asserting the critical antecedent of the conditional claim they go beyond what they are in a position to assert *qua* philosophers trying to satisfy philosophical curiosity.

The matter is, I think, quite different with what I call acquaintance. I stub my toe and I believe that I am in excruciating pain. What justification do I have for thinking that I'm in pain? How do I know that I'm in pain? My answer is that I am directly aware of the pain itself—the very truth maker for my belief. The pain is “there” transparently before my mind. The thought that is about the pain and the pain that is its object are both constituents of the conscious mental state that I call acquaintance. When all this is so, we are in state that is all that it could be by way of satisfying philosophical curiosity. What more could one *want* as an assurance of truth than the truth-maker there before one's mind? When one is directly acquainted with pain as one entertains the proposition that one is in pain, there seems to me to be no need, no point, in

moving up a level and asking about the justification one has for believing that one is in this state. It is not that one can't ask the question. The question is well-formed and there is, of course, a readily available answer. Just as acquaintance with pain was a completely satisfying way of assuring oneself that one is in pain, so acquaintance with this acquaintance with pain is a completely satisfying way of assuring oneself that one is acquainted with pain.

But again, I would emphasize that it doesn't strike us as even relevant to explore the second-level question as a way of getting a *better* sort of assurance that one is in pain. Why would it? If I'm right, what is relevant to getting the assurance one wants as a philosopher is getting the *pain* itself before one's consciousness.<sup>7</sup> In the second-level act of acquaintance the pain is present before consciousness again as a constituent of a more complex state of affairs, but having it before consciousness in *that* way is no better, so to speak, than having it there as an object of first-level awareness.

The matter is quite different, I think, with belief-forming processes that may or may not be reliable (or that may or may not be functioning properly, or "tracking" facts). Am I noninferentially justified in thinking that I am in pain when I stub my toe? The reliabilist, for example, says that I am *provided that* my belief is caused by a process that is unconditionally reliable. The philosopher can't *resist*, at this point, asking the obvious next question. But *is* my belief caused in the right way? The question is irresistible not because one *in general* needs second-level justification in order to have first-level justification. The question is irresistible because having a belief caused in a certain way when we don't know whether or not it is caused in that way is clearly not something that would give us assurance of truth. Strangely enough, some externalists seem to become aware of that fact themselves when they try to apply their analyses at the next level. They realize that track-record arguments aren't really getting us anywhere when it comes to giving us the assurance we seek. The appropriate moral to draw, however, is that if you can't live with a track-record argument given your claims about what is genuinely sufficient for having justification, then you should abandon those claims. Furthermore, if you feel the *need* to move up a level to satisfy philosophical curiosity, that too is an indication that you should reconsider your view about what is sufficient for (philosophically relevant) justification at the first-level.

Richard Fumerton  
University of Iowa

---

<sup>7</sup> I'm focusing here on philosophically satisfying noninferential justification. The story is much more complicated for inferential justification. For inferential justification, what I think the internalist wants is direct acquaintance with probabilistic connections between evidence and conclusion. Wanting this and getting this are two quite different matters.

## References

- Alston, William. 1993. *The Reliability of Sense Perception*. Ithaca: Cornell University Press.
- Bergmann, Michael. 2000. "Externalism and Skepticism." *The Philosophical Review* 109, 113-28.
- . 2002. "Defeaters and Higher-Level Requirements." Unpublished.
- Cohen, Stewart. 2002. "Basic Knowledge and the Problem of Easy Knowledge." *Philosophy and Phenomenological Research*, forthcoming.
- Fumerton, Richard. 1996. *Metaepistemology and Skepticism*. Boston: Rowman and Littlefield
- Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.
- . 2001. "Internalism, Externalism, Defeaters and Arguments for Christian Belief," *Philosophia Christi*, Series 2, Volume 3, 379-400.
- Stroud, Barry, 1984. *The Significance of Philosophical Skepticism*. Oxford: Clarendon Press.
- . 2000. *Understanding Human Knowledge*. Oxford: Oxford University Press.