Achieving Epistemic Ascent

Sosa's epistemology has long been marked by a desire to avoid unnecessary polarization through compromise that incorporates the insights of opposing camps. In particular he has recently urged us to view both the foundationalist/coherentist and the internalist/externalist controversies in epistemology as *false* dichotomies. Can we find neutral ground between these warring epistemological factions?

The Distinction Between Animal and Reflective Knowledge and the Search for Compromise:

In his book *Knowledge in Perspective* and more recently in his paper "Two False Dichotomies", Sosa stresses a distinction between what he sometimes calls animal knowledge or, following Descartes, *cognitio,* and reflective knowledge, or, again following Descartes, *scientia.* Put very crudely, one achieves animal knowledge simply by getting at the truth in an appropriate (non-accidental) way. The justification constitutive of animal knowledge lends itself to an externalist analysis. Reflective knowledge, the kind of knowledge that philosophers, for example, seek, requires more. It is here that internalist intuitions are most at home. I will argue that there is something profoundly right about Sosa's attempt to make this distinction, and in this paper I will evaluate not only Sosa's suggestion as to how to understand the nature of the ascent from animal to reflective knowledge, but some competing views as well.

Sosa's Conception of Animal Knowledge:

Sosa's account of animal knowledge is essentially a variation on reliabilism, although Sosa prefers to characterize it as virtue epistemology. In "Intellectual Virtue in Perspective" Sosa tries to analyze epistemic concepts like knowledge employing as a conceptual building block the idea of an intellectual virtue. He suggests that we understand an intellectual virtue of a person S as a *relative* concept.[1] S has an intellectual virtue relative to some set of circumstances C, an environment E, an inner nature I, and a field of propositions F, when S has an inner nature I such that if S in C and E and has nature I and if S either believes or disbelieves some proposition P from field F then S is very likely right with respect to P. Because Sosa wants a sighted person to have an intellectual virtue vis a vis believing propositions about his environment that a blind person lacks (even when the blind person avoids error by consistently withholding belief about those facts about the environment that could be known only through sight) he modifies the account to require that in the relevant conditions C and E, the virtue must likely result in true belief. Additional complications are introduced into the analysis to avoid potential counterexamples. Sosa insists, for

[1] **The clearest statement of this view is presented in "Intellectual Virtue in Perspective," in Knowldge in Perspective (Cambridge: Cambridge University Press, 1991).**

example, that even if one has a stable disposition to believe truly in a certain environment, one will forfeit the kind of intellectual virtue required for knowledge if in a relevant broadening of that environment one retains the disposition to form beliefs in the same way, but where the beliefs formed are typically false.[2]

The essence of animal knowledge, then, is true belief that results from intellectual virtue. While we could always make perspicuous the relevant relational claims implicit in an ascription of knowledge, we don't, and Sosa has an interesting discussion of various pragmatic considerations that enter into our assumptions regarding the relevant level of generality presupposed in the implicit choice of descriptions for the relevant circumstances and field of propositions. The basic idea is that we want the intellectual virtue that is essential to knowledge to be *useful* in predicting truth. We want to be able to infer from the fact that S's belief was intellectually virtuous that it is likely to be true, and we want to be able to discover the fact that the belief was the product of an intellectual virtue without having to *independently* establish the truth of the proposition S believes. So, for example, S may have a stable disposition to believe truly that there is a table in front of him when he has visual sensations that result from light reflecting off the surface of a table, but one can hardly use this knowledge of S's disposition to form beliefs together with the fact that S believes that there is a table before him in order to usefully predict that S's belief is true. To discover that S's belief is an instance of the relevant disposition we would need to have *independent* reason for supposing that S is before a table that is causing in S the relevant sensations.

Now I don't propose to evaluate here the details of Sosa's account of intellectual virtue and the way in which it might figure into an analysis of animal level knowledge (cognitio).[3] Rather I want to address the more general issue of whether we need some such externalist account of epistemic concepts, and the question of whether we also need more robust epistemic concepts satisfaction of which allows us to *ascend* from this sort of knowledge and justified belief to the kind of knowledge and justified belief that internalists suggest we seek in our more intellectual moments. I'm also interested in exploring the way in which the lower and higher level epistemic standards connect.

The Appeal of Externalism:

The rise of naturalism and externalism in epistemology is, in part, a reaction to what strikes many as the traditional epistemologist's radical over-intellectualizing of belief formation. The vast majority of Modern philosophers seemed to take for granted

---

2. Sosa illustrate the idea with the cartoon character Mr. Magoo. As some of you may remember, Magoo could form reliably true beliefs about object very close, but retained the disposition to confidently believe various propositions about his physical environment even when the objects in question were far away. Sosa wants to argue that Magoo lacks intellectually virtuous beliefs even about the objects close given that he fails to appropriately withhold beliefs about objects far away.

3. I've discussed a number of concerns, particularly about the requirement concerning broadening, in "Sosa's Epistemology," Philosophical Issues, 5, 1995, 15-27.

that we *infer* the vast majority of what we believe about the world around us. Perceptual knowledge, they claimed, involves inference from truths we know more directly about the character of sense data or appearance. Knowledge of the past involves inference from knowledge of truths about the present content of “memory experience”. Knowledge of other minds involves inference based on knowledge of truths describing the behavior of physical bodies. And the problem of skepticism loomed large on the horizon because these same philosophers held very high standards for what constitutes legitimate inference. One can reasonably believe one proposition P as a result of reasonably believing some other proposition E only if one has reason to believe that there is at least a probabilistic connection between the truth of E and the truth of P.

The traditional epistemologist's reasons for supposing that commonplace beliefs about the external world, the past and other minds must be inferentially justified, if justified at all, are familiar. The traditional epistemologist was a foundationalist who believed that the only way to end regresses of justification was with noninferentially justified belief. Under the influence of Descartes, foundationalist sought to find their "first" truths in infallible belief or infallible justification for belief. But the best justification we can imagine for believing propositions about the external world, the past, and other minds, seems perfectly consistent with those beliefs being false. To avoid a fairly radical skepticism we would need to find justification for the beliefs of commonsense in legitimate *inference* from more secure foundations.

The attempt to defeat scepticism playing by the rules of the traditional foundationalist has a troubled history, but without even worrying about the details of the various philosophical theories attempting to reconstruct a plausible inferential justification for everyday beliefs, we might certainly pause, with Reid, to worry about the fact that we just don't seem to *make* inferences of the sort the traditional view requires. And if we don't even make inferences of the relevant sort, how can we plausibly identify the justification such beliefs enjoy as inferential justification?

It's important not to underestimate the depth of the phenomenological problem. The traditional empiricist foundationalist might be tempted to shrug off the alleged problem by turning from conscious inference from occurrent belief to unconscious inference from dispositional belief. It is undoubtedly true that we don't typically consciously assent to propositions describing the character of fleeting, subjective appearance and infer from those truths propositions about the external world. Nor do we consciously infer truths about the past from propositions describing the occurrence of present memory experience. We often seem to just sense that a good friend is unhappy or pensive without first consciously noting some physical characteristic of their expression or behavior. But none of this implies that we lack the relevant dispositional beliefs about sense experience, memory experience, or physical behavior, respectively. The precise analysis of dispositional belief is a matter of philosophical controversy, but there is surely some sense in which we can truly ascribe to a person indefinitely many beliefs, the vast majority of which are beliefs in propositions that are not currently being entertained. Furthermore, there is no reason to deny that dispositional beliefs can be causally efficacious in producing other beliefs, where the existence of those causal connections can constitute a plausible sort of unconscious inference.

But as I indicated above, the problem is more serious for the traditional epistemologist. As many have pointed out, it is not clear that we have even *dispositional* beliefs about the kinds of mental states that the traditional foundationalist takes to be the subject matter of contingent foundational knowledge.[4] Although I can't argue it here, it does seem to me almost a datum that there *is* such a thing as subjective experience, the occurrence of which does not imply the truth of any proposition about the physical world. But we typically don't *attend* to appearance in the way that would be required to form the *ground* of a dispositional belief. One of the first things an aspiring landscape painter needs to *learn* is the fascinating and subtle differences between the appearances objects present. Ordinary people aren't even very good at recognizing the details of how things look even when those very experiences serve as important *causal* clues in beliefs they form as a result of their experiences. When children or novice painters attempt to represent the way things *look,* they are most often unsuccessful partly because they have a very difficult time leaving the world as they believe it to *be* to attend to the world as it actually *appears.*

"Memory experience" is also notoriously difficult to find phenomenologically. It may exist at some level of consciousness and it may be causally operative in producing beliefs about the past, but there seems almost no plausibility to the claim that our beliefs about the past are caused by *beliefs,* occurrent or dispositional, about the occurrence of memory states. Beliefs about recent past events may be accompanied by various images, but it is far from clear that these images are essential to remembering. Indeed, as Ayer suggested some time ago, there may be no more to remembering some fact than having a true belief caused in the appropriate way by the past event.

It may be more plausible to suppose that we must have noticed something about the friend's behavior in order to reach a conclusion about the friend's mood, but if we take the fact that we can't describe, even to ourselves, what the behavioral clue was as prima facie evidence that we don't have a belief about that behavior, we are again hard-pressed to discover a belief (justified or not) in premises from which we can legitimately infer our conclusion.

Externalist bring to the table a refreshingly undemanding account of both nondifferential and inferential justification. They seem to accommodate a possibility of justified belief that is more in harmony with the phenomenological data. As a species we may have evolved to respond to all sorts of noncognitive stimuli with appropriate beliefs and expectations. And if we have relatively stable dispositions to arrive at the truth in this way, why can't we view the stimulus/response belief-forming mechanisms as the very source of knowledge and justification? We need a concept of animal knowledge, of animal rationality because we *are* animals among other animals. When the young gazelle encounters a hungry lion for the first time, it is indeed fortunate that it does not need to employ inductive reasoning to reach the conclusion that flight would be appropriate. If the world is as we think it is, nature has no doubt taken care of this for the gazelle, and although it *may* involve anthropomorphizing on our part, it is

[4] **See, for example, John Pollock in xxx (forthcoming)**

certainly noteworthy that we describe the gazelle as *knowing* instinctively (without needing to rely on experience) that there is danger present. While human beings are far more complex than gazelles, and may have the capacity to form intentional states that precede such things as flight behavior, it is hardly plausible to suppose that all of *our* beliefs and expectations are at the mercy of our reasoning ability. And this is as fortunate for human survival just as it is for the survival of the gazelle. As Hume so eloquently put the point discussing the way in which humans respond to sensations with beliefs about the external world:

> Nature has not left this to our his choice, and has doubtless esteem'd it an affair of too great importance to be trusted to our uncertain reasonings and speculations.[5]

The fact that we don't take the plausibility of Hume's supposition to reflect on the possibility of making distinctions between reasonable and unreasonable belief, people who have knowledge and people who don't, strongly suggests that we do employ epistemic concepts the satisfaction of which does not require all that much of cognitive agents. We do, of course, find it useful to distinguish people who have capacities to get at the truth in certain predictable ways from people who lack those capacities, and it is the epistemic concepts we employ to mark such distinctions that externalist try to analyze in more formal ways.

I might add briefly to the above remarks a philosophical observation that might be at least tangentially relevant. As a philosopher who has always resisted the invitation to naturalize and externalize philosophically interesting epistemic concepts, I have nevertheless always been struck by how easy it is to play the game of finding counterexamples and proposing solutions to counterexamples from within the externalist's framework. And how would this be possible if we didn't have epistemic concepts that were at least in the neighborhood of proposed externalist analyses? The matter is, for example, quite different from the way in which, say, a dualist would respond to behavioristic analyses of mental states. For the dualist, the behaviorist is so far off the mark that it doesn't even seem remotely possible to "fix" the fundamental problems that pervade the behaviorist's analyses. There seems much more room for philosophical give and take, however, when it comes to thinking about various versions of reliabilism. One *understands* the counterexample Goldman worries about in connection with the person whose reliable memory conflicts with abundant evidence that his memory is unreliable[6] and we can join Goldman in trying to fix the problem. We know what Sosa is worried about with respect to the problem of generality and we can work with him to find a plausible approach to dealing with the problem. Again, all

---

[5] **David Hume, A Treatise of Human Nature, ed. by L. A. Selby-Bigge (London: Oxford University Press, 1888).**

[6] **In his "What is Justified Belief?" in Justification and Knowledge, ed. by George Pappas (Dordrecht: Reidel, 1979, p. 19-20.**

this suggests to me that there is something that the externalist is getting at in the attempt to understand a concept of knowledge and justified belief in terms of a belief's having the relevant genesis.

Dissatisfaction with Externalism--Is That All There is?:

But if we hard-core internalists are occasionally tempted to flirt with externalist analyses of at least some epistemic concepts, we eventually always recoil at the suggestion that there is no more to knowledge and justified belief that what the externalist has to offer. Painting with a broad stroke, I would suggest that the fundamental internalist concern is that having knowledge or justification in the externalist sense doesn't seem to satisfy philosophical *curiosity.* It doesn't seem to provide any *assurance* of the sort the philosopher wants when wondering about the truth of various propositions. Even some confirmed externalist seem to lose their nerve when moving up levels of knowledge and justification. Plantinga, for example, seems content to argue that given his externalist analysis of knowledge and warrant, we *might* have noninferential warrant for believing various propositions about the existence of God. The "might" is surely an epistemic operator. The implication is that we don't *know* anything inconsistent with the proposition that the Christian has the relevant knowledge. But the *contextual* implication of the use of the modal operator seems to make the concession that we also don't know that the Christian *does* have such knowledge. But on most externalist analyses of epistemic concepts, having second-level knowledge that one knows or justified belief that one has a justified belief is not much more difficult that having first-level knowledge or justified belief. If justified belief is reliably produced belief and if beliefs about the past resulting from memory are justified because they are reliably produced, then, pace Alston, the externalist should allow that one can produce a straightforward track-record argument for the conclusion that memory is reliable and get second-level justification for believing that first-level beliefs about the past are justified. If memory and induction are reliable ways of forming belief, one can *remember* remembering various events and *remember* those events having occurred and can then use an inductive argument to generalize that beliefs resulting from memories are usually true. Memory, induction, and perception can be employed together to form the reliable belief that perception is reliable. Memory and induction can be used to form the reliable belief that induction is reliable. If reliability is the essence of justification then achieving knowledge and justified belief at the higher levels seems at arms reach provided that there are the relevant reliable belief-forming mechanisms at the first level. Plantinga's Holy Spirit who aids in producing appropriate beliefs about God and His nature can surely just as easily produce warranted metabeliefs about the activity of the Holy Spirit in producing appropriate beliefs.

Despite all this, and despite having acknowledged that the externalist may have insight into at least some epistemic concepts, we almost can't help retreating to an internalist desire for more intellectually satisfying knowledge and justified belief, at least at higher levels. We can't seem to avoid retreating to conditionals. *If* memory is

reliable then we have justified beliefs about the past. And *if* memory and induction are reliable then we probably have justified beliefs that we have justified beliefs about the past. The externalist has opened the door to the epistemic *possibility* of knowledge and justified belief at the first level, but for some reason we shy away from the claim that it is epistemically impossible that we lack such knowledge. It is not, I think, that the contextualist is right and that when doing philosophy all kinds of skeptical alternatives suddenly become relevant. It is rather, I think, that when doing philosophy and starting to think about the fundamental questions concerning knowledge and justification we get serious and insist that knowledge and justification requires something more than a capacity to get at the truth in certain non-accidental ways.

But what is the ascent we want to achieve in gaining knowledge or justification of the sort that will alleviate philosophical curiosity, provide intellectual assurance and how can we achieve it.

Sosa's Conception of Epistemic Ascent to Reflective Knowledge:

In a number of places, most recently in “Two False Dichotomies,” Sosa suggests an answer to this question of what is necessary to achieve reflective knowledge. He begins by emphasizing the following principle of epistemic ascent:

(KA) If one really knows that p and one considers whether one does, then one must be justified in thinking that one does.

Notice that Sosa does not assert that knowledge implies knowing that one knows, or even that knowledge implies having the capacity to know that one knows. Nor is it clear whether he would accept an analogous principle of ascent for justification:

(JA) If one really believes P justifiably then if one considers whether one does then one must be justified in thinking that one believes P justifiably.

This latter issue is important if we fear regress from Sosa's ascent principles. After all, while one might think it initially plausible to suppose that someone's knowing P requires that person to justifiably believe that he knows that P if he considers the question, do we also want to insist that if the person were to consider the question of whether he justifiably believes that he knows that P he would find himself justifiably believing that he justifiably believes that he knows that P , and so on ad infinitum. In no time at all the higher level beliefs will presumably get too complicated for any normal epistemic agent to keep things straight.

I want to be clear about the nature of the regress I fear from JA. I'm not suggesting that a principle of ascent need involve one in conceptual regress. Just because one thinks that my justifiably believing P would require me to justifiably believe that I am justified in believing P if I consider the question, it doesn't *follow* that that justified metabelief need be *constitutive* of my justifiably believing P. In other words, it needn't be an *analytic* truth that if I justifiably believe P then upon consideration I would justifiably believe that I have such justification. But even if the principle were not

analytic, a principle of ascent concerning justification might require something of which finite epistemic agents are clearly incapable. Now it may be that Sosa would reject JA and for that reason deny that he faces any problematic regresses, but I'm not sure why JA has any less initial plausibility than KA, particularly if we can make a distinction between animal-level justification and reflective justification analogous to Sosa's distinction between animal-level knowledge and reflective knowledge.

Why exactly does Sosa think that KA is plausible? Well , he begins by asking us how we would react to the person who responds to the question "Do you know that P?" by saying, for example, "Maybe, maybe not?" Don't we think that such a person has within his or her doxastic system a kind of disharmony that destroys the possibility of first-level reflective knowledge? That suggestion does seem initially plausible, but I wonder whether one could acknowledge its plausibility without introducing levels of knowledge (animal and reflective). Perhaps, for example, the plausibility of Sosa's claim stems from nothing more problematic than the fact that the subjective certainty condition for knowledge (at least knowledge of the sort that philosophers are interested in, or ordinary people are interested in when they make clear by various inflections that they want to know whether you *really* know that P) looks like it's failing when uncertainty is revealed at the higher level. When the criminal defense lawyer asks the witness if he is absolutely certain that he saw the defendant at the scene of the crime, she will no doubt claim a victory of sorts if the witness responds "Maybe, maybe not, but there's at least a 50/50 chance that I'm certain." That sort of metalevel uncertainty about certainty seems, at the very least, to be in strong tension with the *possibility* of being genuinely certain at the first level. But I'm not sure that the tension is anything other than an epistemic tension. If we suppose that generally people have a kind of unproblematic access to their own occurrent mental states, then if someone is hemming and hawing about whether or not he is certain that P, that's pretty strong evidence that he really isn't certain that P. But that it constitutes strong counter-evidence against first-level certainty needn't seduce us into thinking that a plausible metaphysical account of first-level certainty should reveal a necessary connection between possessing it and having the capacity to be certain that one possesses it. In any event, however we evaluate the intelligibility of uncertainty about genuine certainty, intuitions we have about this are not directly related to Sosa's principle of epistemic ascent. That principle requires that in order to know the person reflecting on the question of whether she knows must have a justified belief that she knows.

What sort of incongruity attaches to a proposed hypothetical situation in which S knows that P but has no evidence that bears on the question of whether or not he knows that P? If incoherence is something that can destroy reflective knowledge, what is the relevant incoherence in the belief system of the person who finds himself with no justification for believing that he knows that P? The answer may depend on just *why* it is that the person lacks justification for believing that he or she knows that P. Certainly, if I am wondering whether or not I know that the stock market will collapse, and I realize that I haven't the slightest reason to believe that I know that the stock market will collapse, that might force me quickly to the conclusion that I don't know that the stock market will collapse. But again, perhaps it depends on why I have no

justification for believing that I possess the relevant first-level knowledge. One perfectly plausible hypothesis as to how second-level facts concerning justified belief about knowledge might destroy the possibility of genuine first-level knowledge concerns reflection on the *truth condition* for knowledge. If I'm thinking about first-level knowledge in an orderly way, I might think separately about the various conditions that seem individually necessary and jointly sufficient for knowledge. And, of course, one of those conditions is the truth condition. I know that P only if P is true. Now if I'm evaluating the possibility of my knowing that P by first examining the truth condition, and conclude that I have nor reason to believe that I know because I have no reason to believe P, then it follows rather straightforwardly on most Gettier-proofed justified true belief accounts of knowledge that I should infer that I don't know that P--I don't know that P because I don't have a justified belief that P, and in its absence, of course, I don't have justified true belief of the sort that could constitute knowledge.

Of course, being unjustified in believing the truth condition for knowledge is satisfied is only one way of being unjustified in believing that I know. As we saw, I might have no justification for believing that I know that P because I don't have any reason for supposing that I am subjectively certain that P. Again, given certain plausible contingent truths about how easy it should be to discover facts about the contents of my own mind, my lacking justification for believing that this condition for knowledge is satisfied would be rather powerful *evidence* to suppose that it wasn't satisfied.

But suppose I do have good reason to believe that P and that I have good reason to think that I am certain that P. On a justified true belief account of knowledge, what else could prevent me from having a justified belief that I know that P? Well, of course, I might not have good evidence for believing that I have good evidence for believing P. In "xx" Sosa claims that "one's belief amounts to reflective knowledge only if one can say that one does know, not just arbitrarily, but with adequate justification." Now if I find myself lacking good evidence for thinking that I have good evidence for believing P it probably *would* be a violation of the rules governing conversational implicature to go around *claiming* that I know that P. In most contexts you are not supposed to make claims if you realize that you don't have good reason to believe what you claim. But from the fact that I shouldn't *say* that I know that P if I don't have good reason to believe that I know that P, it doesn't, of course, follow that I can't know that P without having good reason to believe that I know that P. If I don't have good reason to believe that it is raining outside now (or at least if I realize that I don't have good reason to believe that it is raining outside now), then I shouldn't *say* that it is raining outside now. But from the fact that I shouldn't say that it is raining outside now, it doesn't follow that it isn't raining outside now. Without a justified belief that I've satisfied the conditions for knowledge (whatever they are), I shouldn't claim to know P, but i doesn't follow that if I don't have a justified belief that I know that P, I can't know that P (even after I consider the question of whether or not I do have such evidence).

Summarizing, there may be interesting evidential connections between lacking justification for believing that one knows and lacking knowledge, and there may be interesting connections between lacking justification for believing one knows vis a vis the appropriateness of *claiming* to know, but have we yet been given any reason to

suppose that there is an important kind of knowledge, reflective knowledge, that requires having the capacity to form justified beliefs that we have some other kind of knowledge. Notice that if my understanding of Sosa is correct, then this is the right way to put the relevant question. We are concerned with understanding the conditions under which one can ascend to one sort of knowledge, *reflective* knowledge, by being able to justifiably believe that we have another sort of knowledge, *animal* knowledge. We need a good reason to believe that reflection on the plausibility of a principle of epistemic ascent gives us a good reason to introduce these two concepts of knowledge.

Sosa argues that one is surely epistemically better off if one can answer the question of whether or not one knows with an affirmative "Yes, that is something I do know," than if one must confess "No I don't know that" or if one must admit that "Maybe I do and maybe I don't." So that we that don't confuse issues about what is necessary for having reflective knowledge with what is necessary for appropriately claiming to have such knowledge, let's put the point in terms of justified belief. Let's suppose one provides an analysis of knowledge and that we have some individual S who satisfies the necessary and sufficient conditions for knowledge proposed by the analysis. There are these three possibilities:

1) S might know that P but be justified in believing that he doesn't.

2) S might know that P but have no reason to believe that he knows though also no reason to believe that he doesn't know.

3) S might know that P and also have strong justification for believing that he knows that P.

Of the three epistemic situations S might be in, Sosa argues, 3) is surely *better* than 1) and 2). So if we have an externalist analysis of knowledge that seems to fall short of an epistemic ideal and one is trying to suggest some way in which one can improve one's epistemic position so that one possesses a different kind of knowledge, a better kind of knowledge--reflective knowledge, then isn't it plausible to define reflective knowledge in terms of knowledge *plus*? Reflective knowledge is knowledge plus at least the capacity to reach a justified conclusion that one has knowledge. We can reconcile an externalist analysis of knowledge with our internalist yearnings for something more satisfying by simply distinguishing two kinds of knowledge--animal knowledge understood in terms of arriving at truth reliably (footnote), and reflective knowledge understood in terms of possessing internal justification for believing that one is getting at truth reliably. One can avoid having to choose between externalism and internalism. Furthermore, if Sosa is correct and we should understand the justification that turns animal knowledge into reflective knowledge in terms of coherence, then we may also be able to reconcile our foundationalist inclinations with our implicit recognition of the epistemic importance of coherence, by accepting an externalist, but still foundationalist, account of animal knowledge and combining it with a coherence theory of justification of the sort necessary to turn animal knowledge into reflective knowledge.

An Externalist Response:

I think the externalist will reject Sosa's offer of compromise. The externalist typically offers an analysis of knowledge with an externalist justification condition. The knowledge defined will not, of course, require any logical connection between knowing and having justification for believing that one knows. When Sosa asks the externalist whether it wouldn't be better to know that P and also have the capacity to justifiably recognize that one has such knowledge, the externalist can certainly respond in the affirmative. From an epistemic point of view, it's probably better to know that one knows that P than just to know that P, better to justifiably believe that one knows that P than to just know that P, better to justifiably believe that one has a justified belief that P than to just have a justified belief that P, and so on. When one emphasizes the qualification "from an epistemic point of view," these may just be tautologies. If from an epistemic point of view more knowledge and justified belief is better than less knowledge and justified belief, then the above claims are all obviously true. But one can admit all this and propose precisely the same externalist analysis of the higher-level epistemic states that one provides of the lower-level epistemic states.

Sosa's Epistemic Ladder:

Sosa will reject an attempt to climb an epistemic ladder of ascent from animal knowledge to reflective knowledge by building metalevel knowledge or justified belief on first-level knowledge or justified belief when the knowledge or justified belief at the metalevel is given the same externalist analysis as first-level knowledge or justified belief. And I think he is absolutely right in thinking that this sort of ascent doesn't ever really get us into the better epistemic position we seek. It doesn't really allow us to leave the realm of animal knowledge. In short, Sosa' compromise is to give the externalist an externalist understanding of animal knowledge but require for reflective knowledge something more satisfying. If what I said earlier in this paper is plausible, the problem with leaving one's metaepistemolgy with only the conceptual tools of epistemic concepts understood as the externalist understands them is that we realize that satisfying such concepts doesn't give one the kind of *assurance* of truth one seeks as a philosopher or as any of us who find ourselves reflectively worried about whether or not we really know what we think we know. If I start to wonder whether there really is a physical world with the characteristics I take it to have, by intellectual curiosity isn't affected one way or another by the fact that I happen to be getting at truths about that world in a non-accidental way. Nature or Plantinga's God may have arranged for me to get at the truth when prompted by appropriate stimuli, but that doesn't do me any good at all when it comes to assuring myself that I am indeed getting at the truth. And I believe it is precisely that sort of assurance that reflective knowledge (or reflective justified belief) is supposed to provide.

How does one get the additional assurance that would constitute having reflective knowledge? Sosa's answer is that one gets oneself metabeliefs about the sources of one's beliefs where the metabeliefs cohere in important ways. Put too

crudely, perhaps, Sosa wants to understand reflective knowledge as animal knowledge with coherent belief that one has animal knowledge where the coherence of one's beliefs about the ways in which comes to believe reliably constitutes the kind of justification that will satisfy the internalist demands on knowledge.

But now one must insure that the account of justification that one employs in one's account of reflective knowledge does not *itself* leave one yearning for a more satisfying ascent to yet another sort of justification. The following observations are hardly original. Indeed the most devastating internalist critique of coherence as a source of philosophically satisfying justification was given by BonJour when he himself was a coherence theorist.[7] BonJour argued that coherence without *access* to coherence wouldn't give the internalist the sort of justification the internalist wants. To his enormous credit, BonJour effectively reminded us that there are two sorts of coherence theories--internalist and externalist. One can define a belief's having justification simply in terms of its cohering well with other beliefs in one's doxastic system. Or one can insist that a having justification for a belief requires that one be *aware* of the fact that one's belief coheres with the rest of what one believes. If one understands justification in terms of coherence without requiring access to that coherence, then it seems clear to me that we will now need to make a distinction between "animal" justification, and reflective justification; between justification that is intellectually satisfying and justification that is not. We can surely mimic Sosa's rhetorical questions concerning ascent with respect to knowledge and ascent with respect to justification defined in terms of external coherence. Wouldn't it be somehow better not only to have a belief that coheres with the rest of one's beliefs but to be aware of the fact that one's belief system is indeed coherent? If satisfying reflective epistemic concepts is supposed to put is in a more satisfactory epistemic position, then surely reflective justification requires not only coherence but access to coherence. But access coherence theories face insuperable problems.

How precisely are we to understand access to coherence? Minimally it would involve access to our beliefs and access to logical and probabilistic connections. But "access" is itself a thinly disguised epistemic term. If "access to" means "knowledge of" or "justified belief about" our coherence theory of reflective knowledge or our coherence theory of justification faces vicious conceptual circularity. Notice that Sosa's own strategy for distinguishing between animal and reflective knowledge avoids structural circularity. He can define animal knowledge without invoking the concept of justification (defined in terms of coherence) and he can then define reflective knowledge in terms of justified belief about animal knowledge. But the conceptual circularity will only be postponed if he concedes that coherence without access to coherence doesn't do the job of giving us the sort of justification that would satisfy an internalist. Without access requirements to coherence, however, it's not clear that we have given the internalist anything that would allow the internalist to view the internalist/externalism debate as a false dichotomy.

Coherence theorists who try to incorporate access into an account of justification

---

[7] **In The Structure of Empirical Knowledge (Cambridge, MA: Harvard University Press, 1985.**

may be able to escape conceptual circularity by introducing yet another epistemic concept into their conceptual framework. If coherence without access to coherence cannot constitute philosophically satisfying justification, then why not simply recognize that in addition to coherence one needs to introduce some notion of being aware of, having access to, having direct acquaintance with belief states and relations of coherence? The answer, of course, is that in doing so one will simply cease being a coherence theorist. It does, of course, seem entirely plausible to suppose that we have a kind of unproblematic access to what we occurrently believe and certain logical, perhaps even probabilistic, connections between prepositions we believe. It is revealing that historically coherence theorists just seemed to give themselves knowledge of what they believe (just as contemporary anti-realists just seem to give themselves unproblematic knowledge of the ways in which they represent reality). But what business to they have introducing unproblematic access to mental states. If there is such a thing as direct acquaintance with mental states, and it can constitute a kind of knowledge or justification of propositions made true by those mental states, then one is a traditional *foundationalist* and not a coherentist (though one may, of course, acknowledge certain inferential connections between foundationally justified belief as providing a way of increasing the justification those foundationally justified beliefs enjoy.

But isn't the problem of getting reflective knowledge or justified belief obviously going to arise even for a traditional foundationalist who tries to understand nondifferential justification in terms of direct acquaintance with a fact? I can't really address that question fully here, but I do think the answer is "No." When one is directly acquainted with one's pain while one believes that one is in pain and while one is also directly aware of the correspondence between the thought that one is in pain and the pain, that just is the epistemic state that constitutes genuine reflective knowledge. That just is the epistemic state that satisfies philosophical curiosity, that constitutes philosophical assurance. When one represents the world a certain way and one has the relevant truth-maker for that representation unproblematically before consciousness, there is nothing more one could want by way of epistemic assurance.

Epistemic Descent--Another Approach to Distinguishing Reflective from Animal Knowledge and Rationality:

If there is a moral to be drawn from the above discussion, it is I think that one should *start* one's metaepistemological investigations by trying to discover a kind of knowledge and justified belief that is a good candidate for reflective knowledge and justification. Reflective knowledge and justified belief must be such that when one possesses it one thereby gains the kind of assurance than resolves one's intellectual curiosity. Earlier in this paper, I argued that Sosa is right to mark a distinction between reflective and animal knowledge. I think it is equally important to mark a distinction between intellectually satisfying justification and belief the rationality of which falls short of providing intellectual assurance. Rather than start by trying to understand animal knowledge and justification, however, I think we might more profitably begin with an internalist account of reflective knowledge and justification and work our way down, so to speak, to less intellectually demanding externalist concepts of knowledge and

justified belief by peeling away some of the more demanding conditions on reflective knowledge and justification..

I've argued at some length elsewhere for a very traditional distinction between foundationally justified belief and inferentially justified belief. As my earlier remarks indicate I believe one should understand nondifferential justification in terms of direct acquaintance with facts, representations of those facts and correspondence holding between the representations (thoughts) and the facts. I want to focus here, however, on inferential justification. I argued earlier that the traditional epistemologist maintained demanding standards for inferential justification. Traditional epistemologists (and for that matter, most proponents of contemporary externalist variations of foundationalism) insist that to be inferentially justified in believing P on the basis of E one must be justified in believing E. But they also argued, that an ideal epistemic agent possessing ideal inferential justification for believing P on the basis of E would be aware of either a logical or probabilistic connection between E and P. As we remarked earlier, however, these requirements for inferential justification are so strong that it seems doubtful that ordinary epistemic agents (or, for that matter, most epistemologists) can satisfy them. Philosophers have struggled long and hard to figure out how to legitimately infer truths about other minds from knowledge of physical behavior. There may be some epistemologists who have come up with the correct solution to the problem and who have finally figured out what the relevant probabilistic connection is, but the one thing we can be certain of is that most philosophers have hallucinated a phantom probabilistic connection. We know this because we know that the correct philosophical position concerning evidential connections will always be a minority view. If philosophers can't discover the relevant evidential connection through years of investigation, what chance to ordinary people have satisfying these extraordinarily demanding requirements for inferential justification.

We should emphasize, however, that one of the advantages of distinguishing animal rationality (knowledge) from reflective rationality (knowledge) is that we shouldn't be particularly surprised to discover that we lack reflective knowledge and justified belief. Epistemologists have been preoccupied, one might suggest obsessed, with the goal of defeating skepticism. It just doesn't seem right to suppose that we must conclude from our philosophical frustration that we lack knowledge and justified belief. But the whole point of distinguish animal rationality from the kind of rational belief that satisfies philosophical curiosity is that we can allow the possibility of animal-level knowledge and justified belief while maintaining suitably high standards for the kind of knowledge and justification we seek as philosophers. If Hume was right, we might just be out of luck when it comes to satisfying reason (intellectually demanding reason) with respect to the vast majority of what we believe. But that doesn't mean we will stop trying or that we will pretend to have satisfied reason when the answers to our philosophical question remain elusive.

But what sort of rational belief or knowledge could we possess if we fail to satisfy the more intellectually demanding standards of reflective knowledge and justified belief? What could constitute a kind of inferential justification that falls short of full-blooded reflective inferential justification? Well, it could be the case that we have a complex set of justified dispositional beliefs which together with the occurrence of

various psychological states (sensations and memories, for example) *cause* us to believe various propositions about our environment. Now I have argued elsewhere that one should view *facts* as the relata of causal connection, where a fact is understood in terms of particulars exemplifying relational or non-relational properties at a time. Indeed, I think that the most straightforward generality theories of causation require something like facts to be the relata of causal connection. Hume himself stresses that it is essentially instantiated properties that properly viewed as causes. Once we identify a cause in terms of the exemplification of properties, we have a straightforward way of identifying the regularities that constitute the relevant constant conjunctions--a's being F at t will cause it to be G at t + 1 when it is a law of nature that whenever something is F at one time it is immediately thereafter G.

There is an advantage to the internalist trying to make room for a derivative externalist conception of knowledge and rational belief in allowing that facts are the relata of causal connection. Facts are also the most plausible candidates for the truth-makers of propositions. The very sensory state that (together with background beliefs) causes me to believe that there is a table before me is also the truth-maker for a proposition describing that state. Perhaps we can understand the rationality of the belief that results from a sensory state in terms of the evidential connections that hold between the sensory state, the propositions justifiably but only dispositionally believed, and the proposition that it is the object of the belief that is produced. Reflective inferential justification requires that we be aware of evidential connections between propositions believed. Uneffective inferential justification would require only that the relevant evidential connections obtain where the relata of the connections include not only prepositions believed but propositions that are not believed but are made true by the experiential states that causally contribute to our "conclusions." The resulting account of unreflective justification will be importantly external in that we are allowing that one can have inferential justification without having cognitive access to the justifier. But it also contains elements of internalist, at least if one insists that the causes that are the justifiers must all be internal states. The account of non-reflective inferential justification can be employed in an account of non-reflective inferential knowledge, although the account will be no more unproblematic that other justified true belief accounts of knowledge that need to find additional conditions to avoid Gettier problems.

The above suggestion for externalizing a kind of knowledge and justification that is less demanding that the sort that philosophers seek to satisfy is not new. Haack defends a version of it in explaining the foundationalist elements in her now famous foundherentism.[8] Haack wants experience to play a crucial role in the justification of beliefs but, largely for the kind of phenomenological reasons discussed earlier, she doesn't want knowledge of the external world to rely on beliefs about the character of sensory experience. She also wants to suggest that we can relate the causal role experience plays to an evidential role but defining its evidential role in terms of

---

[8] **Susan Haack, *Evidence and Inquiry* (Oxford: Blackwell, 1993). I suggested that one could accommodate a derivative evidential role for experience in "Inferential Justification and Empiricism," *The Journal of Philosophy,* LXXIII, No. 17, 1976, 557-69.**

evidential connections between propositions describing the experience and the propositions about the physical world believed as a result of those experiences. It seems to me, however, that she has a far too liberal view as to which preposition describing experience one consider in evaluating its evidential role. At one point, for example, she suggests that the relevant preposition describing the propositional/evidential counterpart of the sensory state A has when seeing a rabbit, she offers:

> A is in the sort of percpetual state a normal subject would be in, in normal circumstances, when looking at a rabbit three feet away and in good light. (p. 80)

If one allows the relevant preposition to describe this sort of relational property in assessing the epistemic value of the experience, one might as well go all the way and let the relevant proposition be:

> A is now having an experience that is caused by the presence of a rabbit before me.

A's sensory experience will now provide him with infallible justification for believing that the rabit is present!

Our account of non-reflective inferential justification obviously needs some sort of principle way to choose from among the indefinitely many propositions describing experience the relevant one for the purpose of assessing the epistemic contribution of the experience. Again, this is where the metaphysics of causation might help. An experience exemplifies infinitely many properties, relational and non-relational but only certain properties are such that their exemplification plays a causal role in producing belief. We can take the properties that are causally relevant to be those constitutive of the fact that is the truth maker for the relevant proposition. My belief that there is table in front of me is caused by a visual experience that may have the property of being the kind of experience usually caused by tables under this sort of condition, but reelection on standard epistemological problems strongly suggests that it is only the non-relational intrinsic character of the experience that is relevant to producing the belief. I'd believe precisely the same thing about the table if I lived in a world in which demons typically produce hallucinatory experience. So on the account I'm suggesting the only experiential proposition relevant to assessing the epistemic contribution of the experience would be the proposition made true by the exemplification of the non-relational (intrinsic) properties of the experiential state.

Uneffective Knowledge, Justified Belief, and Skepticism:

I want to make clear that in introducing an intellectually less demanding concept of knowledge and justified belief, I am not asserting that we have knowledge and justified belief of this sort. It *may* be that we have animal knowledge even though we don't have reflective knowledge, but it *may* be that we lack both reflective and animal

knowledge. The modal operator is epistemic and we now have a perfectly natural way of interpreting it. Relative to what we (or at least most of us) reflectively know it is both epistemically possible that we have unreflective knowledge and epistemically possible that we lack it.

Conclusion:

Sosa was fundamentally correct in suggesting that we don't have to choose between inernalism and externalism. Furthermore, he was fundamentally correct in suggesting that the ground for compromise was to be found in a distinction between kinds of knowledge, and, I would add, kinds of justification. I have strong reservations about the attempt to understand intellectually satisfying knowledge by layering justification understood in terms of coherence upon animal knowledge. Instead I would try to find a kind of animal knowledge by stripping away some of the more intellectually demanding conditions on inferential knowledge while leaving in place the fundamental role of evidential connections.