

Chapter One. Knowledge as Success from Ability

11-19-08

1. The central thesis.

The central thesis of this book is that knowledge is a kind of success from ability. Let us suppose, with Aristotle, that the intellectual virtues are abilities. Then knowledge is a kind of success from virtue. This is a thesis about what knowledge is. More specifically, and more importantly, it is a thesis about the sort of normative status that knowledge requires. The thesis, then, is that knowledge is an instance of a more general normative phenomena—that of success through ability (or success through excellence, or success through virtue).¹

Adopting this thesis allows progress on a range of epistemology's problems. Some of these are "problems for everyone." That is, they are perennial problems of the field that any adequate theory of knowledge must address. Others are "problems for reliabilism." That is, they are problems that arise for reliabilist theories of knowledge in particular, and that must be addressed if reliabilism is to be a viable approach in epistemology. To that extent, the book can be viewed as an extended defense of reliabilism as a theory of knowledge. In effect, the present approach shows how knowledge can be normative within a reliabilist framework.

Knowledge is a kind of success from ability. This is intended as both an account of knowledge and an account of epistemic normativity. Is it intended as an "analysis" of knowledge? Not in any traditional sense. For example, I do not pretend to give a *conceptual* analysis, where to do so is to analyze some complex concept into more simple conceptual building blocks. Neither do I pretend to give an *ontological* analysis, where to do so is to analyze some property into ontologically prior parts. As Timothy Williamson has pointed out, these sorts of analysis make little sense outside the context of earlier philosophical projects.² Do I intend, at least, to give necessary, sufficient and

informative conditions? No, because to say that knowledge is a kind of success from ability is not to give *sufficient* conditions. On the contrary, it is to give a species without giving a difference. Nevertheless, the account is informative in a straightforward sense: it provides insight into what knowledge is by identifying it as an instance of a more general, familiar kind. Again, the central thesis is that knowledge is a kind of success from ability. Put differently, knowledge is a kind of *achievement*, as opposed to mere lucky success. This locates knowledge within a broader normative arena. Moreover, this is an arena in which we operate with both familiarity and facility. By reflecting on our thinking and practices in this arena, I want to argue, we gain insight and understanding into what knowledge is.

2. Three themes.

The central thesis of the book has now been introduced. In the present section I elaborate on some related themes, organized around this central thesis. In Section Three I provide a brief outline of the discussion to follow.

a. Epistemology as a normative discipline.

When we say that someone knows something we are making a value judgment. We imply, for example, that his or her judgment is preferable to someone else's mere opinion. But then knowledge attributions and the like have a normative or evaluative dimension. Epistemology is a normative discipline.

These claim are hardly controversial—almost everyone in the field would accept them. But in this book I take them seriously and allow them to focus and organize our inquiry. If knowledge has an evaluative dimension—if epistemology is a normative discipline—then a central task of epistemology is to provide an account of the normativity involved.

We may locate that task in the context of traditional epistemological inquiry.

Traditionally, the theory of knowledge has been driven by two central questions: What is knowledge? and What do we know? Corresponding to these questions are two very different projects of epistemology. “The project of explanation” corresponds to the first question and is the focus of Plato’s *Theaetetus*. It asks what knowledge is, and tries to explain the difference between knowing and not knowing. “The project of vindication” corresponds to the second question. It has also been prominent since the early days of epistemology, and is closely related to Pyrrhonian skepticism. It is the project of showing that we have knowledge, in general or in some domain. The first project is concerned to *explain* what knowledge is and how knowledge is possible, whereas the second is concerned to *establish* that knowledge exists.

Contemporary epistemology has clearly privileged the project of explanation over the project of vindication. Even further, there is an emerging consensus that the project of vindication is somehow flawed or misguided, and that the proper task of epistemology is the project of explanation. I heartily endorse this emerging consensus, and will argue in its favor below.³ But for now I want to stress that this book is a contribution to the project of explanation. The project is to understand what knowledge is, and providing an account of epistemic normativity is an important part of that project.

But even here, we can press questions regarding the object of our investigation: Are we investigating knowledge itself, our concept of knowledge, or the term “knowledge” and its cognates? Are we giving an account of epistemically normative properties, our evaluative concepts, or our language? Different philosophers at different times have understood the project of explanation as directed towards each of these. I want to say that a complete theory of knowledge should have something to say about all of them. That is, it falls within the proper domain of epistemology to investigate what knowledge is, how we think about knowledge, and how the language of epistemic evaluation functions, both semantically and pragmatically. Even more broadly, it makes sense to ask what knowledge ascriptions are for and how our practices of evaluation achieve their

purposes. Getting answers to any of these questions increases our understanding, which is the proper goal of philosophical investigation in the first place.

b. Reliabilism.

Simple reliabilism is a powerful view. It explains a wide range of our intuitions regarding what does and does not count as knowledge, it provides elegant solutions to a range of difficult skeptical problems, and it gives a straightforward account of the relation between justification and truth. Yet there is widespread dissatisfaction with reliabilism. Why so?

Much of that dissatisfaction can be summarized this way: Reliabilism is not sufficiently normative. Knowledge is supposed to be a superior state. There is supposed to be something good or praiseworthy about the person who knows, as opposed to the person who has only opinion. But one's beliefs can be reliably formed and yet lack this superior sort of status. Reliabilism seems to leave out the normative dimension of knowledge. More generally, it seems to leave out the evaluative part of epistemic evaluation.

This sort of complaint comes out in the internalism-externalism debate about the nature of epistemic justification. Epistemologists have meant many things by "epistemic justification," but suppose this term refers to something like "epistemic normativity," or "the sort of normative status required for knowledge." Knowledge is (at least) true justified belief in this sense. But why think that reliably formed belief is equivalent to justified belief? *De facto* reliability seems insufficient for the sort of justification required for knowledge. Neither does reliability seem necessary for justification. The victim of Descartes's demon has justified beliefs in some important sense. He seems as justified as we are in believing, for example, that he has a body and that he is sitting by the fire. And yet the victim's beliefs fail miserably from an external point of view—he is altogether out of touch with reality. Externalism in general, and reliabilism in particular, seem to miss

something important about the epistemically normative.

Another way that dissatisfaction with reliabilism comes out is in the “Is that all there is?” syndrome. Over the years, externalism about epistemic justification has gained ground, largely because internalist theories seem to entail unacceptable skeptical consequences. More and more, the choice seems to be between a) externalism and knowledge or b) internalism and skepticism. Faced with this choice, most epistemologists will embrace externalism. But dissatisfaction persists. If ordinary knowledge is externalist, doesn’t that show that we wanted more than ordinary knowledge? Suppose that reliabilism is right, and that knowledge is (something like) reliably formed true belief. What is so great about that? And is that all there is?

Finally, some reliabilists have tried to make a virtue out of necessity, embracing “naturalized” epistemology as an alternative to normative epistemology, or at least as a way to downplay the importance of the epistemically normative. Here the reasoning seems to go something like this: If knowledge is robustly normative then something like an internalist theory of that normativity must be right. But that way leads to skepticism, and to an otherwise inadequate account of our ordinary practices and evaluations. Therefore, knowledge is not robustly normative.

I think that this reasoning is mistaken, and that the mistake is in the first premise. Knowledge *is* robustly normative, but that normativity is not internalist. On the contrary, epistemic normativity is of a perfectly familiar externalist sort, and one that is perfectly “natural” in any relevant sense of the term. Again, knowledge is a kind of success from ability. Put another way, knowledge is a kind of *achievement*, or a kind of success for which the knower deserves credit. And in general, success from ability (i.e., achievement) has special value and deserves a special sort of credit. This is a ubiquitous and perfectly familiar sort of normativity. Thus we credit people for their athletic achievements, for their artistic achievements, and for their moral achievements. We also credit people for their intellectual achievements. Epistemic normativity is an instance of

a more general, familiar kind.

I have been claiming that knowledge is a kind of success from ability, and that this is a plausible and fruitful idea. But here I want to emphasize something else: that the present account of epistemic normativity is tailor-made for reliabilism. For starters, the account is externalist. It makes knowledge and epistemic normativity depend on the knower's abilities, and on relations among abilities, environment and success. These are all paradigmatically externalist factors. But there is a straightforward sense in which the account is also reliabilist: it makes agent reliability an important condition on epistemic normativity. In effect, this book defends an account of knowledge that makes normativity safe for reliabilism. It shows how knowledge can be both reliabilist and robustly normative.

c. Knowledge and understanding.

Finally, a third, related theme recurs in the chapters that follow: that epistemology benefits from a distinction between knowledge and understanding. One benefit of the distinction is that it points to a richer plurality of epistemic goods. Truth and knowledge have epistemic value, but so do understanding, wisdom and other intellectual goods. It should not be unexpected that the values manifested in this plurality are distinctive, and therefore require distinctive treatment in a complete epistemology. A related benefit of the distinction is that it takes significant theoretical burden off the concept of knowledge. If the concept of knowledge serves as a catch-all for every sort of epistemic value, then too much will have to be crammed into it. Put differently, leaving anything of value out of the concept will make it seem that something has gone missing. If knowledge is understood as just one intellectual good among others, however, then the concept of knowledge need not do so much work. In particular, it now becomes possible to accommodate deep-seated intuitions motivating internalism, evidentialism and coherentism while rejecting them as intuitions about knowledge *per se*.⁴ Likewise, and in

related fashion, we get an easy answer to the “Is that all there is?” question. The answer is no—there is understanding, wisdom, and other epistemic goods as well.

Granting that a distinction between knowledge and understanding can be useful, what exactly *are* knowledge and understanding? These are, of course, big questions. I have already said something about knowledge. Details will follow. As far as understanding is concerned, it will be enough to serve the purposes of this book to have a rough and ready understanding of what understanding is. For that we can look to some thoughts by two recent authors, both of whom have called for more focus on the concept of understanding. We can then tie their ideas into a venerable tradition.

First, consider what Jonathan Kvanvig says about understanding.

The central feature of understanding, it seems to me, is in the neighborhood of what internalist coherence theories say about justification. Understanding requires the grasping of explanatory and other coherence-making relationships in a large and comprehensive body of information. One can know many unrelated pieces of information, but understanding is achieved only when informational items are pieced together by the subject in question.⁵

. . . understanding requires. . . an internal grasping or appreciation of how various elements in a body of information are related to each other in terms of explanatory, logical, probabilistic, and other kinds of relations that coherentists have thought constitutive of justification. (192-3)

Wayne Riggs describes understanding in similar terms.

The kind of understanding I have in mind is the appreciation or grasp of order, pattern, and how things ‘hang together.’ Understanding has a multitude of

appropriate objects, among them complicated machines, people, subject disciplines, mathematical proofs, and so on. Understanding something like this requires a deep appreciation, grasp, or awareness of how its parts fit together, what role each one plays in the context of the whole, and of the role it plays in the larger scheme of things.⁶

Tradition has it that understanding is intimately related to explanation, and explanation to knowledge of causes. Thus Aristotle identifies *episteme*, sometimes translated as “scientific understanding,” with knowing the cause of a thing. Commenting on Aristotle’s philosophy of science, R. J. Hankinson writes,

To have scientific knowledge, then, is to have explanatory understanding: not merely to ‘know’ a fact incidentally, to be able to assent to something which is true, but to know *why* it is a fact. The proper function of science is to provide explanations⁷

Here we must take “cause” in a broad sense, to include all of Aristotle’s four causes: material, formal, efficient and final. To understand, on Aristotle’s view, is to have an explanation in terms of these sorts of causes. Put differently, it is to have an answer to different sorts of “Why” questions, in terms of what a thing is made of, what sort of thing it is, what produced a change, or what end was intended.

Putting these various ideas together, a rough and ready notion of understanding emerges: Understanding is knowledge of causes, or something close to it. Understanding involves “grasping,” “appreciating” or knowing causal relations taken in a broad sense; i.e. the sort of relations that ground explanation. This is still rough, but good enough to imply both a) the obvious value of understanding and b) a distinction between understanding and knowledge *per se*.

I will end this section by briefly mentioning a related concern. Some epistemologists think of the intellectual virtues very differently than the way I have been suggesting. Namely, they think of them as traits of character akin to the moral virtues. Examples are intellectual courage, open-mindedness, and intellectual carefulness. Granted that such character traits exist and that they are a kind of intellectual excellence or virtue, what is the place of this kind of virtue in epistemology?

Some authors have argued that character virtues should have pride of place in an account of knowledge. For example, Linda Zagzebski has argued that knowledge can be understood as virtuous true belief, where the virtues are understood as acquired excellences of character akin to the moral virtues.⁸ This seems wrong, however, since there are paradigmatic cases of knowledge that seem not to involve that sort of virtue at all. It is a hard sell, for example, that such virtues are always and essentially involved in cases of perceptual knowledge. That is not to say, however, that the character virtues are not important in other ways. For example, it is plausible that such virtues are often needed to make our cognitive abilities reliable. Put differently, it is plausible that such virtues are sometimes needed to turn mere faculties into excellences. This would give character virtues a place in the theory of knowledge, even if not as a necessary or essential condition on knowledge. Perhaps more interestingly, such virtues might have place in a complete epistemology because they are essentially involved in goods *other* than knowledge, or perhaps because they are valuable in themselves.⁹

3. Outline.

The chapters in this book are divided into three parts. Part One introduces a virtue-theoretic account of epistemic normativity and defends it against alternatives.

A “virtue-theoretic” account of epistemic normativity is one that explains knowledge in terms of person-level excellences. More specifically, it explains the normative dimension of knowledge in terms of person-level excellences. As I said above, I here

understand such excellences as intellectual abilities, or intellectual powers. Hence, knowledge is success from intellectual ability. Put differently: In cases of knowledge, S believes the truth because S believes from intellectual ability. Because abilities are reliable dispositions, the account is a version of reliabilism. Because abilities are person-level dispositions, the account is a version of agent reliabilism. On the present account, then, knowledge and epistemic normativity require reliability, where the seat of reliability is the knower herself.

This account of epistemic normativity is inconsistent with several approaches currently on offer. In Chapter Two, I argue against deontological approaches, understood as theories that understand epistemic normativity in terms of rules or norms. Such theories are inconsistent with a virtue-theoretic account because they *require* that knowledge be rule-governed. A virtue-theoretic account requires no such thing, and so the different kinds of theory place different conditions on epistemic normativity. I argue that a virtue-theoretic account has the advantage here. On the one hand, such an account nicely explains why etiology matters in cases of knowledge. That is, it explains why whether a person knows has something to do with how the person formed her belief. On the other hand, a virtue-theoretic account need not say too much about etiology. Specifically, it need not demand that knowledge-producing processes be rule-governed in the sense required by deontological theories. On the contrary, it can view this as an empirical question for cognitive science rather than a normative question for epistemology.

Deontological theories are at the level of normative epistemology—they try to give a substantive account of epistemic normativity in terms of intellectual rules or norms. Internalism about epistemic normativity is one level up: internalist theories place a restriction on any adequate substantive theory. Specifically, internalist theories demand that epistemic normativity must be understood in terms of factors that are appropriately “internal” to the knower. We can interpret this as a supervenience thesis: the normative

status required by knowledge supervenes entirely on states that are internal to the knower. Different kinds of internalism will differ in their details, but all of them are inconsistent with a virtue-theoretic account. That is because a virtue theory understands epistemic normativity in terms of causal and other modal properties—properties that are paradigmatically externalist. In Chapter Three I argue against internalist theories so understood. A variety of considerations are brought to bear, but the central idea is this: We can imagine cases (psychological twin cases, for example) where there are differences in knowledge-relevant normative status, but no differences in internal states. But then knowledge-relevant normative status does not supervene on internal states.

Having rejected internalist theories of epistemic normativity, we are well positioned to argue against evidentialist theories. Evidentialist theories, as the label implies, try to explain epistemic normativity in terms of evidence and evidential relations. Such theories can be internalist or externalist, depending on what is allowed to count as good evidence. Internalist versions, of course, will place internalist constraints here. In Chapter Four I argue that internalist versions of evidentialism are false because internalism is false. Moreover, externalist versions of evidentialism are unmotivated. More specifically, once we adopt an externalist understanding of what counts as good evidence, we lose the motivation for saying that all knowledge requires evidence. On the contrary, it becomes plausible that there are other ways of satisfying the normative requirements for knowledge—being based on good evidence is only one way among others.

Part One defends a virtue-theoretic account of knowledge and epistemic normativity. Part Two shows how the account can be applied to problems for everyone, or problems that any epistemology must say something about. Two of these problems concern the nature and value of knowledge. The “nature problem” asks what knowledge is. The “value problem” asks why knowledge is valuable. More specifically, it asks why

knowledge is more valuable than mere true belief. Chapters Five and Six take up these questions respectively.

The present account claims that knowledge is a kind of success from ability. Put differently, knowledge is a kind of achievement, as opposed to a merely lucky success. This gives a kind of “framework” for understanding what knowledge is. That is, it gives us a place to start, and to begin to fill in the details. Chapter Five begins to do so by filling in as follows:

S knows that p *if and only if*

1. p is true;
2. S believes that p; and
3. S's believes the truth because S's belief is produced by intellectual ability.¹⁰

The term “because” in clause 3 marks a causal explanation. The idea is that, in cases of knowledge, S's abilities explain why S has a true belief. S believes the truth because, for example, she saw clearly, or reasoned well, or remembered accurately. Put differently: in cases of knowledge, S's believing the truth is explained by S's abilities, as opposed to dumb luck, or blind chance, or something else.

Adopting this idea allows us to explain a broad range of cases, including many cases of ordinary knowledge and many standard Gettier cases. The account does less well, or less obviously well, with other cases however, including the Ginet-Goldman Barn Façade case, some cases of testimonial knowledge, and some cases of innate knowledge. Chapter Five develops the account of knowledge here offered so as to better handle these cases. Even after elaboration, however, we must admit that the account is still not maximally informative. This is because it depends on both semantic and pragmatic

considerations that are not well understood. Specifically, we would like a better understanding of causal explanations and causal explanation language.

The account of knowledge defended in Chapter Five is not maximally informative. It is informative enough, however, to explain why knowledge is more valuable than true belief. In fact, the solution to the value problem falls out of the solution to the nature problem: Knowledge is a kind of success from ability, and we value success from ability more than we value mere lucky success. Chapter Six elaborates on this idea and answers some objections that have been raised against it.

Some of the questions that are raised by Chapters Five and Six concern the relations between knowledge and practical interests. On the face of it, practical interests are relevant to what counts as an appropriate causal explanation. In so far as knowledge attributions involve causal explanations, it would seem that practical interests are relevant to what counts as an appropriate knowledge attribution as well. But how so? Is this merely a matter of appropriate assertion, or can practical interests affect truth-conditions as well? And if practical interests can affect truth-conditions, whose practical interests matter—those of the knower, those of the speaker, or someone else's? Chapter Seven tries to get straight on these questions, and offers some considerations in favor of a contextualist semantics for knowledge attributions. The main idea is this: A central role of knowledge attributions is to identify information for use in practical reasoning. But then knowledge attributions ought to be sensitive to the different interests in play in different practical reasoning environments. These considerations, it is argued, are enough to motivate a contextualist semantics, since different speaker contexts will be concerned with different practical reasoning environments. A corollary of this line of reasoning, however, is that contextualism cannot be put to its usual anti-skeptical work. That is because considerations about knowledge and practical reasoning are also stabilizing considerations: If knowledge is tied to practical reasoning in plausible ways,

then the conditions for knowledge cannot vary widely (or wildly) across speaker contexts. The resources for the standard contextualist reply to skepticism evaporate.

Motivated contextualism does not reply to skepticism. No worries, one might think, because reliabilism does. In other places I have argued as much, showing how reliabilism has easy answers to Cartesian, Humean and Pyrrhonian skepticism.¹¹ Some would say, however, that reliabilism's answers are *too* easy. This is another manifestation of the “normativity problem” for reliabilism. The skeptic wants something—the skeptic *values* something—that reliabilist theories don't deliver. The fact that knowledge is possible—even easy—on reliabilist theories, suggests that reliabilism is not really delivering the goods. Or to put things somewhat differently: If reliabilist theories do deliver knowledge, then it wasn't knowledge that the skeptic was after in the first place. Chapter Eight addresses this sort of concern by focusing on the Pyrrhonian problematic. More specifically, it examines a Pyrrhonian metaphor that likens *unreflective* true belief with grasping gold in the dark. I argue that, on a virtue-theoretic account, we can explain why knowledge is a valuable achievement, even in the absence of the sort of reflection that Pyrrhonism requires. Here I draw analogies to the problem of moral luck, and to an analogous skepticism about the possibility of moral achievement.

Part Two concerns problems that any epistemology must address. Part Three concerns problems for reliabilism in particular. One such problem is the “The Problem of Strange and Fleeting Processes.” Process reliabilism holds that knowledge and justification are to be understood in terms of reliable processes. Roughly, a belief is justified in the sense required for knowledge just in case it is formed by a reliable cognitive process. A variety of cases, however, show that process reliabilism is too weak. Specifically, there are cases where a cognitive process is indeed reliable, but seems too strange or too fleeting to give rise to knowledge. For example, Alvin Plantinga describes a case where S is the victim of a brain lesion, one effect of which is that it

reliably produces the belief that one has a brain lesion. But suppose that S has no evidence for this, or even has evidence that he does not have a brain lesion. It seems wrong to attribute knowledge in the case, even though S's belief is in fact reliable formed. The present view promises to do better here. Intuitively, S's belief that she has a brain lesion is not produced by a cognitive ability, even if it is produced by a reliable process. In effect, the present view places a restriction on which reliable processes can give rise to knowledge: only those that are grounded in cognitive abilities do.

More needs to be said, however, since it is not clear why the dispositions associated with S's brain lesion do not constitute a cognitive ability. In general, we would like a principled reason that distinguishes reliable dispositions that constitute an ability of the knower from reliable dispositions that do not. Chapter Nine tries to do just this by invoking the notion of cognitive integration. It is argued that this helps resolve the cases involving strange and fleeting processes, and helps with some other tough cases for reliabilism as well.

Chapter Ten addresses "The Problem of Defeating Evidence," a problem that has been underappreciated by externalists in epistemology. The problem concerns the way that externalist theories typically approach the concept of defeating evidence, or evidence that counts against one's beliefs. In short, it is not unusual for externalists to adopt an *internalist* account of such evidence. For example, it is not uncommon for externalists to think of defeat relations as *a priori* or quasi-logical, just the way internalists think about evidential relations in general. This is not a happy situation, for now the externalist has given one account of supporting evidence (perhaps in terms of reliability), and a different account of defeating evidence. This is at best schizophrenic and at worse incoherent. Externalists have to do better, and it is argued that the present account helps. Specifically, we can give a principled account of defeating evidence in terms of person-level cognitive abilities.

Finally, Chapter Eleven addresses “The Problem of Easy Knowledge.” As was already mentioned above, reliabilism seems to have easy answers to a range of skeptical problems. In fact, some have argued that reliabilism’s answers are too easy. This sort of problem comes out when we ask whether we know that skeptical scenarios are false. For example, do we know that we are not brains in vats, and if so, how do we know such a thing? According to reliabilism, so long as the belief is reliably produced, we do know such a thing. Moreover, if the world is anything like we think it is, the belief would have to be reliably produced, since there are no close worlds, no close situations, where we get this wrong.

This sort of problem recalls G. E. Moore’s “Proof of an External World.” Moore was no reliabilist, but he seemed to say that we could know that an external world exists, and easily so—we need only note that here is one hand and here is another. But isn’t that too easy? Many philosophers have thought so, and have accordingly rejected Moore’s general stance. In this final chapter I defend Moore by invoking both epistemological and methodological principles from Thomas Reid. In fact, I agree, Moore is merely following Reid on several points, and rightfully so. The upshot is this: Moore has his tongue in his cheek when he claims to “prove” that an external world exists. On the contrary, the skeptic is misguided to demand a proof in the first place, since that is not how one would know such a thing. By adopting a virtue-theoretic account of knowledge, it becomes perfectly clear, and even intuitive, that one knows that external objects exist by seeing them.

The foregoing gives the outline and organization of the book. The chapters have been written so that they are relatively self-contained. Accordingly, it is possible for the reader to follow her preferences and to go directly to one chapter rather than another. The overall argument for the account here defended, however, is given over all of the chapters. That is, the argumentative strategy is to reject alternatives (Part One), and to show how

the account allows progress on a range of epistemology's problems (Parts Two and Three).

.

¹ This thesis has been defended by a number of authors, most notably Ernest Sosa. For the latest development of Sosa's view, see his *A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume 1* (Oxford: Oxford University Press, 2007), and *Reflective Knowledge: Apt Belief and Reflective Knowledge, Volume 2* (Oxford: Oxford University Press, forthcoming). Much earlier than this, Sosa wrote, "knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence." *Knowledge in Perspective* (Cambridge: Cambridge University Press, 1991), p. 277. As early as 1988 Sosa wrote that, in cases of knowledge, one's belief must "non-accidentally reflect the truth of P through the exercise of . . . a virtue." See his "Beyond Skepticism, to the Best of Our Knowledge," *Mind* (1988), p. 184. In the present manuscript I focus on what I take to be the most powerful idea in Sosa's early work, that knowledge is a kind of success from ability, and I attempt to display that idea's explanatory power.

² Timothy Williamson, *Knowledge and its Limits* (Oxford: Oxford University Press, 2000)

³ See especially chapters 8 and 11.

⁴ And to do so without relying on a distinction between animal knowledge and reflective knowledge. Hence this is one place where I depart from Sosa, who makes this latter distinction central to his epistemology.

⁵ Jonathan Kvanvig, *The Value of Knowledge and the Pursuit of Understanding* (Cambridge: Cambridge University Press, 2003), p. 192 .

⁶ Wayne Riggs, "Understanding 'Virtue' and the Virtue of Understanding," in Michael DePaul and Lind Zagzebski, eds., *Intellectual Virtue: Perspectives from Ethics and Epistemology* (New York: Oxford University Press, 2003), p. 217.

⁷ *The Cambridge Companion to Aristotle*, Jonathan Barnes, ed., (Cambridge: Cambridge University Press, 1995), p. 110.

⁸ Linda Zagzebski, *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge: Cambridge University Press, 1996).

⁹ This line of argument is developed more fully in my "Two Kinds of Intellectual Virtue," *Philosophy and Phenomenological Research* LX, 1 (2000), pp. 179-184. For an interesting position along similar lines, see Jason Baehr, "Character, Reliability and Virtue Epistemology," *Philosophical Quarterly* 56 (2006): 161-315.

¹⁰ Clause 3 makes clauses 1 and 2 redundant. I include the redundancies for sake of elegance in formulation. The alternative is something more awkward, such as the following: S knows that p *if and only if* S believes the truth with respect to p because S's belief that p is produced by intellectual ability.

¹¹ For example, see *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry* (New York: Cambridge University Press, 2000); and "Agent Reliabilism," in James Tomberlin, ed., *Philosophical Perspectives 13, Epistemology*. Atascadero, CA: Ridgeview Press, 1999.