Epistemology Idealized

Robert Pasnau
University of Colorado
pasnau@colorado.edu

Epistemology today centrally concerns the conceptual analysis of knowledge. Historically, however, this is a concept that philosophers have seldom been interested in analysing, particularly when it is construed as broadly as the English language would have it. Instead, the overriding focus of epistemologists over the centuries has been, first, to describe the epistemic ideal that human beings might hope to achieve, and then go on to chart the various ways in which we ordinarily fall off from that ideal. I discuss in detail two historical manifestations of idealized epistemology — Aristotle and Descartes — and then consider how this perspective might make a difference to the discipline today. In the end, an idealized epistemology points toward a normative, prescriptive rather than descriptive enterprise.

1. Epistemology as lexicology

I begin with an historical puzzle, and end with the suggestion that something is missing from recent work in epistemology. The puzzle concerns the large gap between epistemology today and epistemology as it was done for most of its history. Today, epistemology primarily concerns conceptual analysis, first and foremost the analysis of knowledge itself. Historically, however, this is a concept that philosophers have seldom been interested in analysing, particularly when it is construed as broadly as the English language would have it. Instead, the overriding focus of epistemologists over the centuries has been, first, to describe the epistemic ideal that human beings might hope to achieve, and then go on to chart the various ways in which we ordinarily fall off from that ideal. My suggestion will be that conceiving of epistemology in this way might make an important difference to how we think about the subject today.

I say that epistemology today centrally concerns conceptual analysis, but in truth its principal concern has often seemed linguistic: Under what conditions is it true to say, in English, that someone knows something? Of course, one is free to choose between the formal and the material mode of exposition, but the field of

Mind, Vol. 122 . 488 . October 2013
doi:10.1093/mind/fzt093 Advance Access publication 12 February 2014

epistemology in particular has often seemed to labour, in modern times, under the formal yoke, wishing to understand the structure of our central epistemic concepts, but settling for the conventions of our ordinary epistemic vocabulary. Something about epistemology as it is currently conceived seems to lure its practitioners into the trap of lexicology.

Epistemology is prone to lexicology because it lacks other sorts of stable grounds on which to build. Philosophers of mind do not belabour the ordinary English meaning of 'consciousness', but rather begin by disambiguating, and thereafter make no fuss over the word itself. It is quite unclear, in contrast, whether 'know' is similarly ambiguous. Theories of causation likewise do not care about the word 'cause' — they take as their starting point various paradigm cases, and try to make those come out right. It is not so clear, however, what counts as a paradigm of knowledge, inasmuch as the most familiar cases, such as $2+2=4$ and *knowing this is a hand*, are notoriously problematic and heterodox in character. Unable to gain traction in these sorts of familiar ways, epistemology tends toward lexicology. Few want such an outcome, but the modern history of the subject repeatedly displays this pattern, and it is unclear what alternative there is.[1]

Quine, four decades ago, famously proposed one such alternative — that epistemology be naturalized 'as a chapter of psychology and hence of natural science' (Quine 1969a, p. 82). This suggestion has of course been enormously influential and fruitful, but seems to take philosophers away from the distinctively normative questions that arise concerning knowledge — questions over whether a given class of beliefs are worthy of the appellation 'knowledge' — and into the purely descriptive domain of how we in fact form beliefs. I mean to describe another approach, one that retains the normative core

[1] Of course, alternative approaches abound. Prominent recent examples include Foley 1987, which turns from knowledge to epistemic rationality; Craig 1990, which focuses on social function; Stich 1990, which would replace analytic epistemology with a pragmatic approach to cognitive success; Williamson 2000, for which knowledge is a basic mental state; and Alston 2005, which turns from knowledge to justification, and then argues for a pluralistic conception of the latter. Others are perfectly happy with epistemology as lexicology. Consider e.g. Ludlow 2005, which celebrates 'the new linguistic turn in epistemology' (p. 12), and remarks that 'first, and most obviously, any investigation into the nature of knowledge which did not conform to some significant degree with the semantics of the term "knows" would simply be missing the point … [E]pistemological theories might be rejected if they are in serious conflict with the lexical semantics of "knows"' (p. 13). This seems just as doubtful as saying that ontology — the study of what things there are in the world — must conform with the semantics of the term 'things'.

of the subject but without the lexicological tendencies.[2] Rather than conceive of epistemology as the quest for the minimal conditions that a belief must satisfy to cross a certain threshold, I suggest we turn our attention toward ideal conditions — not the absolute ideal that a god might achieve, but the ideal epistemic position for a human being, given the powers we have available to us and the kind of world we live in.

To suggest something of what an idealized epistemology would look like, I will begin with two cases — Aristotle and Descartes — where this project has actually been carried out. With that historical material in mind, I will then turn to our modern era, and what difference this approach might make today. Ultimately, I will argue that even those who seek precise criteria demarcating the domain of knowledge would benefit from some account of what the epistemic ideal is. For without a sense of what human inquiry can and should ultimately aim at, it is hard to see how we can hope to come to agreement about when our beliefs should count as warranted, justified, well founded, and so forth. Without an idealized epistemology, in other words, all we can expect to produce is a more-or-less-detailed report on our actual linguistic practices — epistemology as lexicology.

2. Aristotle's ideal theory

No matter how many times your undergraduates may tell you, do not believe them when they write that philosophers have been attempting to define knowledge for centuries. Worse than being a cliché, this is not even true. To be sure, there was Plato, in the *Meno* and the *Theaetetus*, and the casual follower of philosophy's history might be forgiven for supposing that his example carried forward more or less continuously until the present day. In fact, however, Plato has always been more honoured than imitated, and that is particularly the case

[2] The view that epistemology is normative at its core is perhaps controversial. At present this view seems widely accepted, and Riggs 2006 even speaks of 'the value turn in epistemology'. It has been questioned whether the concept of knowledge is normative at all — see e.g. Fumerton 2001, but for a defence see Owens 2000 and, more recently, Weatherson 2008. Regardless, however, of whether either knowledge or justification should be analysed in normative terms, it may be that epistemology broadly construed should be centrally concerned with normative questions about what we ought to believe. This is as much as I here assert.

Quine himself, it should be added, denied that his naturalized approach was intended to remove normative questions from epistemology (see e.g. Quine 1990, pp. 19–21). For a good overview of the naturalized program in epistemology, see Kornblith 1999.

with respect to his interest in definitions. The Platonic dialogues, especially but not only the early dialogues, are interested in defining all sorts of things: knowledge, piety, friendship, courage, justice, statesmanship, and so on. Most of these definitional projects no longer interest us. Although philosophers still sometimes think about friendship and courage, and regularly think about justice, it is rare to find attempts at definition.

Knowledge is the exception. It is only recently, however, that the quest to define 'knowledge' has been perceived as an interesting philosophical question. From Aristotle through the Middle Ages and well beyond, philosophers took an interest in carefully circumscribing one or another particular kind of cognitive grasp of reality — perception, imagination, assent, deduction, etc. — but showed little interest in trying to define the broad category of knowledge. That English contains this very general word of positive cognitive appraisal did not strike philosophers, even those who spoke English, as calling for any special definitional inquiry.[3]

The case of Aristotle is particularly illuminating, both because it comes so soon after Plato, and even more because it would dictate the way philosophers conceived of epistemology for two millennia. Like Plato, Aristotle devotes an entire treatise, the *Posterior Analytics*, to the theory of knowledge, or *epistēmē*. The results he arrives at, however, are utterly different from the sort of picture suggested in the *Theaetetus*. Whereas it might plausibly be thought that Plato, at least in this dialogue, is pursuing our modern lexicological project — considering ordinary usage of the term 'knowledge', and pursuing necessary and sufficient conditions for its satisfaction[4] — it is clear from the very start of the *Posterior Analytics* that this is not

[3] Locke may strike the reader as a counterexample, inasmuch as he does, very prominently at the beginning of Book IV of the *Essay*, define knowledge as 'the perception of the connection and agreement, or disagreement and repugnancy of any of our ideas' (IV.1.2). The example is, however, anomalous. Despite its evident vulnerability, it was not accorded much attention by Locke's contemporaries. It also should be read in the context of Locke's steady effort, throughout the *Essay*, to stake out an English philosophical vocabulary to replace the 'uncouth, affected, or unintelligible' Latinate vocabulary of the Aristotelians (*Essay* epistle p. 10).

[4] For doubts about this, see Burnyeat 1981, pp. 133–6 and Kaplan 1985, pp. 351–3. See also Benson 2000, who argues that the early dialogues should be read as concerned not with knowledge, but with understanding. In contrast, however, see Fine 1990, pp. 114–15, who warns against the over-hasty assumption that Plato's topic is not knowledge in our sense. It should also be noted that in other dialogues, and in particular in the *Republic*, one might see Plato's interest as much more oriented toward the epistemic ideal.

Aristotle's aim. No conversation with an ordinary Athenian, no matter how one-sided, could plausibly have elicited the result that knowledge concerns a proposition that is necessary and universal, known on the basis of an affirmative demonstration in the first syllogistic figure, the premises of which are necessary and explanatory of the conclusion. This is not what even the most erudite Athenian could have meant by '*epistēmē*', before Aristotle came along, and if this is what *epistēmē* is then we would have to conclude that it is something that hardly anyone has ever had, in any domain.

But if the *Posterior Analytics* is not analysing the meaning of 'knowledge', then what is it doing? In what sense is this an epistemology at all? One line of answer to these questions has been to find some other English word that better fits Aristotle's project, the most prominent such suggestion being that this is a theory of *understanding*.[5] Clearly, this is a promising idea about how to translate '*epistēmē*' in the context of the *Posterior Analytics*. The agent who comes to understand a proposition in the way Aristotle describes goes well beyond simply *knowing* that proposition. One can come to know quite well that vines shed their leaves, for example (*Post. An.* II.16), simply by observing it happen. But someone who grasps the general truth in the way Aristotle describes — on the basis of necessary principles grounded in the vine's essence — might plausibly be said to have a better *understanding*. Even so, useful a suggestion as this may be for purposes of translation, it does not go very far toward explaining what Aristotle is after. Is he simply engaged in his own lexicological project, trying to understand a Greek word for which 'understanding' is the closest English counterpart? Presumably there is something special about '*epistēmē*', as the *Posterior Analytics* conceives of it, that makes it worthy of being singled out for special treatment. To be sure, understanding is eminently worthy of study. But why study this, rather than knowledge? And why develop the details in the way Aristotle does? Are the arcane details of his demonstrative method really intended as necessary and sufficient conditions for understanding a thing?

The traditional reading of the *Posterior Analytics* takes it to be a theory of scientific knowledge. This is how '*epistēmē*' is often translated, and the treatise itself is almost always described in these terms, as offering a theory of scientific knowledge or scientific

[5] For this suggestion, see Burnyeat 1981. This is also how Barnes translates *epistēmē* and its cognates.

understanding.[6] It is odd that this should be so, however, because it is apparent on even casual inspection that the treatise's scope is much broader than science as we now conceive of it. Although scientific examples figure prominently, they are not its exclusive focus. The method is evidently meant to apply to mathematics too; there are, indeed, as many mathematical examples as scientific ones.[7] There is also no reason to think that the method is scientific rather than philosophical. Indeed, it is not even clear how we would mark the divide between science and philosophy in this pre-modern era. To be sure, various ancient authors use the plural noun '*epistēmai*' to refer more or less to what we now think of as the sciences.[8] But it is highly misleading to describe the *Posterior Analytics* as a treatise on science, given how much more broadly the theory is meant to apply. A theory that does not discriminate between science and mathematics on one hand, and science and philosophy on the other, is surely not a theory of science in our sense at all. Aristotle scholars will perhaps defend themselves on this point by insisting that, of course, they are using the term 'science' in the broad sense of the Greek '*epistēmē*'.[9] But once that is said, it becomes clear that it explains nothing at all to characterize the treatise as scientific in its concerns.

To describe the *Posterior Analytics* as a theory of science is perhaps most charitably regarded as shorthand for the more complex idea that it aims at an account of *systematic theoretical knowledge*—the sort of thing that one does in mathematics and philosophy just as much as in the sciences. One may speak of the project synecdochically as *scientific*,

[6] The older Oxford Translation (by Mure) renders '*epistēmē*' as 'scientific knowledge.' The idea that the *Posterior Analytics* is a theory of scientific knowledge can be found in virtually any discussion of the topic. Irwin 1988, to take one prominent example, offers this general characterization: 'the *Posterior Analytics* describes the structure of a science and of the content of scientific propositions' (p. 118). And Taylor 1990 remarks that 'the *Posterior Analytics* … gives a detailed account of the conditions necessary and sufficient for the achievement of *epistēmē* in the context of an exact science, but this appears to the modern eye as at best one kind of knowledge, *scientific* knowledge, among others … ' (p. 116). Both Burnyeat and Barnes, though shying away from the term 'knowledge', take for granted that the treatise's subject is science, and Burnyeat even ultimately allows 'that in the end it will not do too much damage to go back to the traditional rendering of ἐπιστήμη as "scientific knowledge"' (p. 132).

[7] See the tabulations in Barnes 1969, pp. 129.

[8] See e.g. Plato, *Republic* VII, 522c. But arithmetic and geometry are also paradigms of *epistēmai* (see e.g. *Post. An.* 75a39, 76b8–9).

[9] See e.g. Barnes 1969, p. 123: '"science" is here of course to be understood in the broad sense of the Greek "ἐπιστήμη"'.

but that is just because, as is so often the case when dealing with Greek philosophy, we do not have the right term in English for conveying what Aristotle is after. There is, however, more to be said here than this. Commentators have almost unanimously latched on to the notion that the *Posterior Analytics* offers a theory of science because they have not seen any other sort of enterprise in the vicinity that the treatise could be concerned with. It is not just that we lack a word to talk about systematic theoretical knowledge, but that we lack any place in our conceptual scheme for the study of such a thing. However, philosophers do of course study the nature of science. Hence it has become an *idée fixe* in the recent literature that this is what the *Posterior Analytics* does.

Regardless of how the topic of the *Posterior Analytics* is to be characterized, there is a further puzzle concerning its methodological prescriptions: that the method described seems both impractical and in fact unpractised by Aristotle. If possessing *epistēmē* requires grasping first principles and essences, then it seems unlikely that we have achieved this condition in more than a few domains. (Mathematics would be the most likely place to find such a methodology in place. But it is unclear whether mathematical proofs satisfy the requirement that one know a proposition through principles that explain the reason why it is so.[10]) Perhaps unsurprisingly, Aristotle's own writings, including the *Posterior Analytics* themselves, contain no examples that satisfy all the necessary requirements. To be sure, his many examples serve individually to illustrate one or another dimension of the theory, but each seems incomplete in one regard or another. The prescribed method, then, seems to be one that he himself is incapable of fully putting into practice.[11]

[10] Aristotle clearly thought that it did satisfy this condition; see e.g. *Post. An.* II.11, 94a27–35. There was controversy over this claim, however, both in antiquity and among later commentators. For discussion see Barnes 1993, pp. 92–3, 107–8. For the later history, see Gilbert 1960, pp. 86–92, and de Pace 1993.

[11] Thayer 1979 refers to 'the extraordinary fact that no one completely satisfactory scientific syllogism fulfilling the requirements Aristotle prescribes for scientific demonstration can be found in *Posterior Analytics*' (p. 100). Barnes 1969 goes even farther, remarking that 'in the whole of the Aristotelian *corpus* there is not, as far as I am aware, a single example of a demonstration' (p. 124). Others have, in the face of this worry, sought to find more extensive signs of demonstration in the *corpus*. Focusing in particular on the biological works, Gotthelf 1987 makes a persuasive case that their structure 'is at least amenable' to the framework of the *Posterior Analytics* (p. 178). (Other essays in that same volume reach a similar conclusion.) But the point remains that, if we take the *Posterior Analytics* seriously as a guide to what *epistēmē* requires, it is hard to see anywhere in Aristotle where *epistēmē* has been achieved.

All of these puzzles dissolve when one reads the *Posterior Analytics* as describing an epistemic ideal. Aristotle characterizes his subject matter as *haplos epistēmē*, unconditional or unqualified knowledge, in contrast to various lesser forms of knowledge, which he is willing to count as *epistēmē*, but which are in one way or another deficient.[12] As I understand the project, these lesser kinds of knowledge are simply the ordinary sorts of knowledge that human beings regularly do possess, and unqualified knowledge is the ideal state that we should aspire to, even if its attainment is extremely difficult. This explains why Aristotle elsewhere seems not to practise what he preaches — not because it is not his goal, but because it is an idealized goal. For this reason, too, the *Posterior Analytics* itself gives us little more than fragments of what *haplos epistēmē* is supposed to be. This is not the perverse failing that it might seem, because Aristotle is describing the ideal aim of inquiry, rather than something that he himself is in a position to do.

Aristotle nowhere says explicitly that *haplos epistēmē* is an ideal that he has not yet realized. But he comes close in the *Metaphysics*, where he remarks:

> The study of truth is difficult in one way, in another easy. A sign of this is that no one is able to attain it completely, nor entirely misses it. But each individual says something concerning the nature of things, so that while he may individually contribute little or nothing, from the collaboration of all there comes a great amount. It is like the proverbial door: who can fail to hit it? In this respect it is easy; but being able to grasp the whole and not only a part makes the difficulty clear. (II.1, 993a30–b7)

Aristotle does not use the term '*epistēmē*' here, but it seems clear enough that the study (*theoria*) he has in mind is just what he had described in rigorous detail in the *Posterior Analytics*. For what is most distinctive about Aristotle's conception of *epistēmē* is his insistence that it involve a grasp not just of a single isolated proposition, but of the whole causal and inferential network of propositions that lie behind it. Aristotle's ideal theory therefore requires a grasp of the

Barnes 1969 (and 1993, pp. xii, xviii–xx) has suggested that Aristotle is offering not a method for acquiring knowledge, but a method for presenting it in systematic form. But, as others have observed (e.g. Burnyeat 1981), this hardly seems to solve the puzzle, inasmuch as Aristotle's method of presentation seems just as far from his prescriptions as does his method of discovery. It seems hard to believe, moreover, that the complex strictures of the *Posterior Analytics* are intended primarily as a strategy for presenting one's research.

[12] For example, 71b10, 72b30, 73b17, 74a33.

whole door, not just a part, and what he tells us here is that it is easy to make a contribution to *epistēmē*, but very hard to achieve the complete ideal. Indeed, 'no one' is able to do that.[13]

As soon as one considers the possibility that Aristotle is offering not a theory of 'knowledge' — even for a special refined domain of inquiry — but an account of the ideal limit of human inquiry, it becomes easy to see why Aristotle would insist on the various characteristics of *epistēmē* that he requires. It is certainly plausible to suppose, for instance, that a truly ideal understanding of any aspect of reality would not just grasp the fact, but grasp the reason why the fact obtains. Those requirements that are not self-evidently ideal are eventually defended in later chapters of the treatise, in precisely the way we should now expect. In I.24, for instance, he offers a lengthy series of arguments for why *epistēmē* should be of the universal rather than the particular. How do we decide? By determining which method is superior. Thus, 'if you know something universally, you know it better as it holds than if you know it particularly. Hence universal demonstrations are *better* than particular demonstrations' (85b13–15). The next chapter argues in similar fashion for why *epistēmē* should be based on affirmative rather than negative premises: not because the latter fail to yield knowledge or understanding or science, but because they are simply 'worse'. And so I.26 continues by showing that positive arguments 'are better' than arguments cast in the form of a *reductio*. Obviously, arguments that are deficient in these respects can nevertheless increase our understanding significantly. In other places Aristotle happily recognizes a wide variety of cognitive states that fall short of the ideal described here, such as grasping particulars and retaining them in memory, reaching conclusions in a non-explanatory way (*hoti* rather than *dia ti*), and achieving practical wisdom in action. We might reasonably describe all of these as kinds of knowledge. The point of the *Posterior Analytics* is simply that they are not ideal.

The question remains of *why* it is so interesting to work out the ideal goal of systematic theoretical inquiry. This still needs arguing for — and I will take the matter up in section 4 — because the significance of such a project has been quite eclipsed in modern

[13] This is not to say that no one ever will, or ever has achieved the ideal. Indeed, Aristotle takes the striking view that philosophy has in the past 'probably often been developed as far as possible and then perished' (*Meta.* XII.8, 1074b10–11; see also *Politics* VII.10, 1329b25–35). But this reflects not so much any particular optimism regarding past cultures, but only his insight into the implications of maintaining, as he does, that the world's past history is *infinite*.

epistemology. For now, though, I think it is clear enough *that* this is Aristotle's project. As soon as one formulates the notion of an ideal epistemology, it becomes completely natural and even obvious to suppose that this is what Aristotle is up to.[14] Certainly, this is how the *Posterior Analytics* was traditionally understood. Albert the Great, in the prologue to his commentary on the treatise, remarks of *scientia* (= *epistēmē*) that 'this is the end and the most perfect and the sole unconditionally desirable thing among the logical sciences'.[15] And when Thomas Aquinas, in his commentary, comes to analyse Aristotle's definition of '*epistēmē*' (71b10), he begins with the remark that 'to have *scientia* of something is to cognize it perfectly'. From this principle he derives the various features of the definition: that what is known in this way must be necessary, and that it must be grasped through a grasp of its cause.[16] Although it was commonplace among later Aristotelians to distinguish between Aristotle's *haplos epistēmē* (*scientia simpliciter*) and various weaker forms of knowledge applicable to everyday life, it seems not to have occurred to any of them, no more than it did to Aristotle himself, that a precise analysis of 'knowledge' in the ordinary sense would be philosophically interesting.

[14] Accordingly, seeds of this idea can readily be found in the existing literature. Lesher (2001), for instance, argues against Burnyeat that the achievement of *epistēmē* goes beyond understanding, and is in fact 'a complete grasp of a subject' (p. 49), amounting to 'expert knowledge' or 'disciplinary mastery' (p. 54). This is certainly in the neighbourhood of my proposal. These remarks, in turn, are perhaps not so far from what Burnyeat himself thinks, inasmuch as he too concludes that 'Aristotle's thought is concentrated on the τέλος, the achieved state of understanding which is the end and completion of the epistemological process' (1981, p. 133). This is, quite precisely, my own view. Taylor (1990), similarly, holds that '*nous* + *epistēmē* is the ideal type of knowledge … ' (pp. 121–2). But because none of these scholars conceives of idealized epistemology as a central epistemological project in its own right, they are unable to give these thoughts their proper prominence and significance. Thus Burnyeat continues to think of the treatise as 'a contribution to the philosophy of science' (p. 97, and see n. 6 above), as does Taylor (n. 6 above), and Lesher puts all of the weight on the too-narrow idea of *epistēmē* as expertise within a discipline.

[15] Albert the Great, *Analytica posteriora* I.1.1 (ed. Jammy, vol. 1).

[16] Thomas Aquinas, *In Post. an.* I.4 (1882, vol. I.2). John Duns Scotus similarly invokes the perfection of *scientia* when he seeks to explain the strict requirements of the *Posterior Analytics* (*Additiones magnae* prol. 1.1 [1639, vol. XI.2]). See also, from the sixteenth century, Francisco Sanches, who defines *scientia* as 'the perfect understanding of a thing' (*Quod nihil scitur*: Francisco Sanches 1988, p. 200).

When Aquinas characterizes Adam's knowledge in the Garden of Eden, he does so precisely in terms of the *Posterior Analytics* framework (e.g. *Summa theol.* 1a 94.3). Such ideal knowledge is just one of the perfections taken from mankind after the original sin (see Reynolds 2006).

On the medieval distinction between perfect *scientia* and lesser forms of knowledge, see Pasnau 2010.

Ultimately, via these scholastic authors, one does arrive at science. Our modern use of that term arose in the seventeenth century, from the Latin '*scientia*', understood at the time as simply the Latinate analogue to Aristotle's '*epistēmē*'. But this does nothing to confirm the idea that Aristotle is giving us a philosophy of science. On the contrary, the explanatory order runs in the opposite direction. Aristotle's project is to construct a theory of the epistemic ideal, which would eventually, two millennia later, furnish a label for the project of modern science — even as the new 'scientists' were calling into question various aspects of the Aristotelian approach. (With 'science' appropriated for these special purposes, some other English term was required for more ordinary cases, and of course usage settled on 'knowledge'.)[17] That modern science grew out of such aspirations for the ideal tells us something about the relationship between Aristotle's theory and science, but it is simply a confusion, historical and philosophical, to think that Aristotle is offering us a theory of science.

3. Descartes's ideal theory

When the goal of theoretical inquiry is pitched as high as Aristotle suggests, it becomes natural to wonder whether it can be reached at all. Ptolemy had remarked in the *Almagest*, for instance, that of the various theoretical disciplines 'only mathematics can provide sure and unshakable knowledge (*eidēsin*) to its devotees'. As for physics and metaphysics, they 'should be called guesswork rather than knowledge (*katalēpsin epistēmonikēn*)' and 'there is no hope that philosophers will ever be agreed about them'.[18] The situation looked more or less the same all the way into the Renaissance, when Pietro Pomponazzi remarked that 'philosophy would be beautiful, if it were as certain as mathematics. For metaphysics and philosophy are conjectural, and on

[17] Not of course that early English-language philosophers *invented* the word 'knowledge', which has its origins in Old English. But it would have been perfectly possible, given linguistic usages in the seventeenth century, for English-language philosophers to use 'science' as the standard noun for ordinary knowledge, just as '*scientia*' was so used in Latin. One sees such a usage in Joseph Glanvill, for instance, in 1661, who finds it quite natural to switch back and forth between 'science' and 'knowledge', as when he remarks that 'He is the greatest ignorant, that knows not that he is so: for 'tis a good degree of science, to be sensible that we want it' (*Vanity of Dogmatizing*, in Glanvill 1661, Ch. 23, p. 225). Locke's *Essay* of 1689, however, aiming to rid philosophy of its Latinate Aristotelianisms, always prefers 'knowledge' in such contexts.

[18] *Almagest* I.1 (Ptolemy 1984, p. 36).

almost any subject one may find different opinions, so that it is like playing with toys'.[19]

One response to such sceptical tendencies would be to dismiss these 'philosophical' domains as promising candidates for systematic theoretical knowledge. Only in mathematics, one might think, can anything approaching the human ideal be found. Another response would be to revise one's ideal theory. The project of idealization, after all, is not meant to yield standards that only a god could achieve. The point instead is to define what sort of knowledge *we* might be able to achieve, given the world we live in. This sort of downward revision culminated in the notoriously pessimistic conclusions of Hume and Kant in the eighteenth century, but might be said to have begun back in the seventeenth, when figures like Galileo and Newton rejected the Aristotelian framework and turned toward a model of science focused less on ultimate explanations and more on identifying lawlike, quantitative rules that capture the observable phenomena.[20]

Still another path would be to embrace the Aristotelian challenge, and insist that philosophy and science can meet it. This was the task that René Descartes set himself, most famously in his *Meditations*, where he attempts to give our knowledge of God and soul the sort of ideal epistemic status that would make it worthy of being called *scientia*. To be sure, Descartes rejects much of the old *Posterior Analytics* framework, such as the syllogism and the requirement of universality. But Descartes holds on to the general framework of an epistemic ideal at which theoretical inquiry should aim. One can indeed find this assumption in place from his very earliest work. According to Rule 2 of his *Rules for the Direction of the Mind*, 'We should attend only to those objects of which our minds seem capable of having certain and indubitable cognition' (X:362).[21] He then immediately remarks, in discussing this rule, that 'All *scientia* is certain and evident cognition'. A few lines below, he adds this: 'So, in accordance

[19] As quoted in Perfetti 2008, Sect. 5, from an unpublished manuscript. This is one of two obstacles Pomponazzi describes as plaguing philosophy, the other being that it does not pay well.

[20] The extent to which either Galileo or Newton rejected the Aristotelian tradition is extremely controversial, and here I mean to advert only to one familiar way of thinking about what happened. For varying perspectives, see e.g. Westfall 1977, McMullin 1978, Cohen 1980. For a recent collection of papers on *scientia* across the seventeenth century, see Sorell, Rogers, and Kraye 2010.

[21] All Cartesian references are to the standard edition, Descartes 1897, which also serves to identify passages in the standard translation, Cottingham et al. 1991.

with this Rule, we reject all such merely probable cognitions and resolve to believe only what is perfectly cognized and what cannot be doubted'. What these passages suggest is that *scientia* is a kind of cognition (*cognitio*), the kind that is certain, evident, and indubitable. He goes on in this same section to characterize this and other rules as ones that 'will help us ascend to the peak of human cognition … ' (X:364). It seems, then, that *scientia* is perfect cognition — or, at least, as perfect as a human being can achieve.

This impression receives confirmation at the start of the *Meditations*, when Descartes makes this famous pronouncement: 'I realized it was necessary, once in the course of my life, to demolish everything completely and start again right from the foundations *if I ever wanted to establish anything in the sciences that was stable and likely to last*' (VII:17). I emphasize the last phrase to call attention to Descartes's goal: attaining stable and lasting results in the sciences. Of course he does not here have in mind 'the sciences' in our modern sense of the term; instead, he is referring to the acquisition of *scientia*, and his ambition is to show how we can acquire *scientia* even in the hitherto murky domain of God and soul. His view, indeed, turns out to be that *scientia* is possible especially in these domains. The standard he holds himself to is remarkably high. In a letter to his disciple Regius written in May of 1640, just after he had completed the *Meditations*, Descartes distinguishes between *scientia* and the conviction (*persuasio*) possessed by someone who cannot help but assent to the clear and distinct perception of some self-evident truth. 'I distinguish the two as follows: there is [mere] conviction when there remains some reason which could lead us to doubt; *scientia* is conviction based on a reason so strong that it can never be shaken by any stronger reason' (III:65). Conviction is a state in which one cannot refrain from assenting to a proposition — in a purely subjective sense, the proposition is indubitable. Presumably, such conviction will be based on *some* reason. But it will not count as *scientia* unless it is based on *unshakable* reasons. This would seem to be indubitability in a stronger sense: it is not just that one is not presently able to doubt the proposition, but that there is no way in which one will ever be able to doubt the proposition, given the reasons one has for it. As he puts it in the Second Replies, 'no cognition that can be rendered doubtful seems fit to be called *scientia*' (VII:141).

These passages point toward the two best-known characteristics of Descartes's conception of *scientia*: his *infallibilism*, that such beliefs must be certain and indubitable; and his *foundationalism*, that such

beliefs must be supported by a foundation of beliefs that are themselves certain and indubitable. To these two characteristics we might add a third, his *internalism*, which is to say that the possession of certain, foundationally structured beliefs does not yield *scientia* unless the believer grasps the reasons that show beyond doubt the truth of the beliefs. This last feature is particularly clear in the following passage, where the requisite adequacy of reasons gets expressed in interpersonal terms.

> Whenever two persons make opposite judgments about the same thing, it is certain that at least one of them is deceived, and it seems that neither has *scientia*. For if the reasoning of one of them were certain and evident, he would be able to lay it before the other in such a way as eventually to convince the other's intellect as well. (*Rules* 2; X:363)

As before, the evidence that Descartes requires for *scientia* must be not just subjectively persuasive, but objectively good, one test of which is whether these reasons would be able 'eventually to convince' others. What we can now also say is that those reasons must be possessed by the agent — they must be *internal* to her. It is not enough that she be able to acquire reasons in principle, or even that she be *able* to understand them *if* they were shown to her. Instead, she must *have* these justifying reasons in the sense that she must actually at some point have held in her mind the reasoning that shows why her current belief is indubitable in the strongest sense. She need not presently be considering all of that reasoning — the requirement is not quite that strong[22] — but it must be the case that she can 'lay it before the other' whenever necessary. Descartes regularly expresses the utmost confidence that his views will pass just this test, once they are understood. Thus he tells Regius in July 1645 that, 'I consider my opinions to be so certain and evident that whoever rightly understands them will have no occasion to dispute them' (IV:248).

Unlike Aristotle's case, it is not customary to treat Descartes's conception of *scientia* as a theory of science. Instead, though with no more plausibility, Descartes is routinely read as advancing a theory of knowledge, and hence he is commonly regarded as the archetypical proponent of foundationalism, infallibilism, and internalism in epistemology.

[22] Thus, after establishing the existence of a non-deceiving God in Meditation 5, and concluding that 'everything that I clearly and distinctly perceive is of necessity true', Descartes adds that 'even if I am no longer attending to the arguments that led me to judge that this is true, as long as I remember that I clearly and distinctly perceived it, there are no counter-arguments that can be adduced to make me doubt it, but on the contrary I have true and certain *scientia* of it' (VII:70).

All of this is quite true, but only with respect to Descartes's ideal theory. As for what is commonly said about Descartes's theory of knowledge, most of this is entirely wrong. If our subject is 'knowledge' as that word is used today, then Descartes is not a foundationalist, not an infallibilist, and not an internalist. Indeed, if epistemology is conceived in its usual modern guise, then Descartes cannot be said to have a theory of knowledge at all. What he has is an idealized epistemology, a theory of *scientia*.

Two sorts of considerations make this quite clear. First, the theory of *scientia* is so demanding that virtually no one can be said to have achieved it, other than Descartes and his followers. Descartes in fact claims that, up until his time, the only *scientia* possessed by anyone has been mathematical *scientia*. For instance, in explaining why we should attend only to what we can cognize with certainty, he remarks that 'if my reckoning is correct, out of all the sciences so far devised, we are restricted to just arithmetic and geometry if we stick to this Rule' (*Rules* 2; X:363). Years later, in *The Search for Truth*, he speaks of 'the slight progress we have made in the sciences whose first principles are certain and known to all' and then adds:

> In the other sciences, whose principles are obscure and uncertain, those who are willing to state their view honestly must admit that, for all the time they have spent reading many a vast tome, they have ended up realizing that they have *scientia* of nothing and have learned nothing. (X:526)

Of course, Descartes thinks that he has managed to push the bounds of *scientia* quite a bit farther. But did he really believe that, up until the middle of the seventeenth century, no one had *knowledge* of anything other than a few claims in mathematics? A sceptic might be happy with this result, but Descartes was no sceptic, or so it is always supposed.

In fact, Descartes goes even farther. He famously holds that 'the certainty and truth of all *scientia* depends on the one cognition of the true God, to such an extent that I was incapable of perfect *scientia* about anything else until I recognized him' (*Med.* 5; VII:71). If *scientia* were knowledge, then this would entail that the atheist lacks knowledge — as, apparently, would anyone who believes in the wrong sort of God. Thus in the Second Replies, in a passage quoted in part already, he remarks that 'I do not dispute that an atheist can clearly *cognize* that the three angles of a triangle are equal to two right angles. I maintain only that his cognition is not true *scientia*, since no cognition that can be rendered doubtful seems fit to be called *scientia*'

(VII:141). Even in geometry, then, *scientia* is available only to someone who has the right religious beliefs, and who uses those religious beliefs in just the right way to ground that *scientia*. All things considered, it can begin to look doubtful that anyone other than Descartes (and his most devoted followers) has ever had *scientia* about anything. If we think of Cartesian *scientia* as knowledge, then we should think of Descartes as a fairly radical kind of sceptic. In so far as that result seems obviously wrong, we should stop thinking of *scientia* as knowledge.

The second reason for denying that Cartesian *scientia* is knowledge is that Descartes disavows any connection between *scientia* and what we should believe. As we ordinarily conceive of knowledge, what we believe ought to reflect what we know. In cases where we lack knowledge, we should therefore either lack belief or at least should believe with hesitation, doubtingly. According to Descartes, in contrast, *scientia* has no connection with what we ought to believe. In the First Meditation he remarks that his habitual opinions 'are doubtful in a way, but are nevertheless highly probable, and are such that it is much more reasonable to believe than to deny them' (VII:22). This passage compares only belief and denial: he does not say that it is more reasonable to believe than to suspend judgement. But a passage from the Synopsis to the *Meditations* goes farther. There he remarks that 'the great benefit of these arguments is not, in my view, that they prove what they establish — namely, that there really is a world, that human beings have bodies, and so on — things that no sane person has ever seriously doubted' (VII:15–16). It would be insane to doubt these matters; yet it is also the case, according to Descartes, that no one has ever had *scientia* of these things. Apparently, then, our lack of *scientia* concerning some proposition has no bearing on whether we ought to believe it. This is quite alien to our modern conception of knowledge.

This aspect of Descartes's account is closely related to the first aspect. Sceptics had traditionally maintained that we ought to withhold assent: if we must act, we should act only as if we have beliefs about the world. Because Descartes sees no such connection between *scientia* and belief, his form of scepticism (if it should be called that at all) is of a purely theoretical sort. Hardly anyone has *scientia*, but this makes no real difference to their lives. For Descartes as a philosopher and scientist, the rarity of *scientia* is a depressing result, and one that he wants to change. But he has no expectation that his methods will lead everyone to acquire *scientia*, and no real interest in seeing that

happen. Indeed, the preface to the *Meditations* explains that he wrote in Latin rather than French 'lest weaker intellects might believe that they too ought to set out on this path' (VII:7). Now Descartes does say — from as early as the *Rules*, as we have seen already — that those of us who seek the truth should believe only what we can grasp with certainty. But this advice applies strictly to those who are pursuing *scientia*.[23] Those who have no interest in that project, or lack the ability to pursue it, or simply have not yet found the time to do so, are certainly not supposed to give up all their beliefs. That, as Descartes says, would be insanity. Even in the absence of *scientia*, ordinary folk *ought* to go on believing what they do. Whether such beliefs should be counted as knowledge, in some more ordinary sense of the word, is not something that Descartes shows any interest in.[24]

As in Aristotle's case, it should seem natural and even obvious to think of Cartesian *scientia* in this way, as soon as one conceives of the project of an idealized epistemology. The problem has been that our modern epistemological framework contains no such conceptual space, leaving scholars without any good way to think of Descartes's position, except as a theory of knowledge.[25] Of the many such efforts

[23] See e.g. *Principles* I.1–3, Fifth Replies (VII:350–51), *Rules* 2–3, *Discourse* 2 'first rule' (VI:18), and esp. *Discourse* 4 (VI:31): 'Since I now wished to devote myself solely to the search for truth, I thought it necessary … '. Broughton (2002), though not concerned with the topic at hand, compares the method of doubt to a game in which one 'winds up suspending judgment about things it would be quite reasonable to believe'. Why? 'The answer must be that the meditator wants what he thinks the method of doubt can give him: a way to achieve sturdy and lasting results in the sciences' (pp. 49–50). Here Broughton is simply paraphrasing the start of the *Meditations* (VII.17, as quoted earlier in the main text) — but why speak of the 'sciences', when Descartes's immediate aim in the *Meditations* is declaredly philosophical? The problem, as in the literature on Aristotle, is that without the notion of an idealized epistemology we lack any better conceptual framework in which to locate Descartes's project.

[24] For a very different reading of Descartes here, see Owens 2008, which contends that once we distinguish the lower standards applicable to 'practical affairs' from the higher standards applicable to belief, we find that for Descartes 'belief or judgment is … governed in all contexts by the rule of certainty' (p. 165). There are passages that might bear this construal, where Descartes distinguishes between what we should believe when searching for the truth and how we should conduct our lives (see e.g. *Principles* I.3; Fifth Replies [VII:350–51], and a 1641 letter to Hyperaspistes [III:422–23]). But to read Descartes in this way is, in effect, to make him into a kind of neo-sceptic, someone who thinks the vast majority of people know nothing or almost nothing and should suspend their beliefs until acquainted with the Cartesian method. It seems to me that the main current of the texts runs strongly against this conclusion.

[25] Newman (2010) is notably sensitive to whether Descartes's epistemology is a theory of knowledge in our sense of the term. He settles on the convention of referring to 'Knowledge' with a capital K, thereby marking, but not settling, the question of what such a thing is supposed to be. Hookway (2003) remarks of the *Meditations* that 'the modern concern with

to frame some epistemic space for Descartes, Ernest Sosa's stands out. Sosa finds in Descartes two levels of knowledge, an ordinary level of mere *cognitio* and a higher grade of *scientia*, which he describes as 'reflective, enlightened knowledge'.[26] The result is that, for Descartes, not all knowledge rises to the level of *scientia* and, accordingly, much of what ordinarily gets said about Cartesian epistemology applies in fact only to one particular, high-grade kind of knowledge. Sosa sees Descartes as the forerunner to his own project, which is to identify two levels of epistemic goodness that might be worthy of the term 'knowledge', depending on how strictly one understands that term. Although I am broadly sympathetic to all of this, I understand these two levels differently. Descartes is not giving us two ways to understand what it is to know something, nor is he marking two particularly salient boundaries of epistemic excellence.[27] One finds in Descartes various casual assumptions, almost entirely undeveloped, about what it is to know something in the ordinary sense. Then one finds in Descartes an ideal theory, which he thinks he can attain in certain domains, as described in the *Meditations* and elsewhere. This is, if you like, a kind of knowledge, but it has little to do with 'knowledge' as we now use that term.

One is liable to misunderstand Descartes's project unless one stresses that what is at stake is the *human* ideal. Perhaps the most common way of understanding Descartes's epistemology is as a radical form of infallibilism, according to which true knowledge (whatever that means) requires absolute certainty.[28] The worry then naturally arises

"knows" and its cognates is almost entirely absent,' but offers no alternative picture of Descartes's project other than remarking that 'his aim seems to be to show that he can contribute to scientific inquiry successfully' (p. 196). For Descartes's continuity with the Aristotelian conception of *scientia*, see e.g. Wolterstorff 1996, though Wolterstorff would seem to overstate the situation when he remarks that 'Descartes took over intact the traditional medieval tripartite scheme of knowledge, faith, and opinion, offering no substantial innovation in how these are to be understood. His attention fell almost entirely on that species of knowledge which is *scientia*' (p. 182).

[26] Sosa 1997, p. 240. See also Sosa 2007, pp. 126–33.

[27] Sosa (2009) makes it clear that he does not suppose there are exactly two well-defined conceptions of knowledge, low and high. Rather, there is a spectrum of worse and better cases, and the question of where to apply the label 'knowledge' is, for Sosa, 'largely verbal' (p. 430). This is appealing in its attempt to escape the lexicology trap, but does not bring us any closer to my main point here, which is that Descartes is describing not one or another threshold along a spectrum, but the ideal limit of that spectrum.

[28] Indeed, Descartes is often thought to be emblematic of a longer tradition. Thus Williams 2004 remarks that 'For much of its history, our epistemological tradition tended to insist that

that this makes scepticism virtually inevitable. In response to this, Descartes scholars are wont to highlight those passages where Descartes admits that the certainty he claims for himself is merely 'human certainty' (Second Replies, VII:144), not that of God or an angel: 'What is it to us if someone imagines that something whose truth we are so firmly convinced of appears false to God or to an angel, and so *is* false, absolutely speaking? What do we care about this "absolute falsity", since we neither believe it nor have even the smallest suspicion of it?' (VII:145). Such a response is liable to seem ad hoc, particularly coming from the author of the notorious sceptical arguments of the First Meditation. It seems one should either insist on certainty or not, and if one cannot achieve it, one should simply say so, rather than waffling over the distinction between absolute certainty and human certainty. Once one sees Descartes's project as a version of ideal theory, however, such worries disappear. It is, as remarked already, no part of an idealized epistemology to engage in a quixotic search for the divine ideal. The point is to understand what *we* are capable of, and then consider the circumstances and domains in which that might profitably be achieved. From Descartes's point of view, then, it is just tedious and misguided to make objections that turn on requiring a sort of certainty that is not within our power. The kind of supernatural certainty that is impossible for us is quite irrelevant to his project.

4. Ideal theory applied

What were philosophers trying to accomplish, in formulating an epistemology of ideal conditions? An answer to this question may suggest something of what we could accomplish through an idealized epistemology today.

Idealized epistemology can be thought of as coming in two stages: the first obvious, the second less so. The obvious initial stage is *idealization*. As I have stressed already, this does not involve abstracting away from all human cognitive limitations, but instead focuses on what would count as perfection for beings such as us, in a world

knowledge properly so-called requires absolute certainty' (p. 123). It would be hard, I think, to find *anyone* who has supposed that *any* human cognitive state can be *absolutely* certain. It is true that, within the later Aristotelian tradition inherited by Descartes, a fairly high degree of certainty was typically demanded for *scientia*. But this is not because of some notion of how we should 'properly' speak of 'knowledge', but because of the historically prominent focus on idealized epistemology.

such as ours. This is the sort of account we have seen Aristotle and Descartes develop. I take it that no one will dispute the interest of this first stage, though some will be more pessimistic than others about the extent to which its realization is possible, and many will take issue with the details of Aristotle's or Descartes's approach. However the details are worked out, this is a project we care about not because any of us aspires to perfection at every moment of our lives, but because some of us make it our concern to seek such perfection at least with respect to some small corner of reality.

An idealized epistemology might do more, however, than simply set out a rulebook for professional inquirers. The second stage of the approach aims at shedding light on when ordinary agents, in ordinary circumstances, are entitled to believe the things they believe. This is the *application* stage — the stage that considers how much of the human epistemic ideal might reasonably be applied to ordinary cognitive agents in everyday life. Descartes does not show much interest in this second stage of an idealized epistemology. His concerns were almost entirely theoretical: he had the ambition of founding an entire metaphysics and natural philosophy on grounds that were as close to ideal as is humanly possible. Normative questions in general were of only occasional interest to him, and normative questions about belief, under non-ideal circumstances, arise only at the margins.

Aristotle, in contrast, developed this second stage in considerable detail. His theory of dialectic is, among other things, an epistemology of non-ideal conditions. The *Topics*, his treatise devoted to this subject, describes what one should do in cases where one or another component of demonstrative reasoning is not available, and considers what merit is left to arguments that fall short of the rigour of *epistēmē*.[29] Although dialectic does not rise to the level of the ideal, it is appropriately deployed by certain kinds of people in certain kinds of situations, and is a worthy subject of philosophical investigation. For the next two millennia, epistemology largely wrapped itself around these two frameworks, demonstrative and dialectical, and subsequent Aristotelians devoted

[29] The theory of dialectic is clearly not only a watered-down epistemology. It does play that role, inasmuch as it provides guidance on forming beliefs in non-ideal circumstances. But it also plays a more practical role in various arenas that Aristotle was concerned with, such as in the art of examining others (see *Soph. Ref.* 11, 172a22). Beyond dialectic lies rhetoric, the focus of which is even more practical and less oriented toward the acquisition of truth and the production of knowledge. What dialectic and rhetoric share, in contrast to the demonstrative method of the *Analytics*, is that they 'are concerned with such things as come, more or less, within the general ken of all men and belong to no one *epistēmē*' (*Rhetoric* I.1, 1354a2).

considerable effort to investigating the conditions under which one or another method was most appropriately deployed.

Aristotle places ideal theory at the head of his epistemology, and then considers how those requirements need to be relaxed in everyday life, both to account for the various intermediate stages on the path toward the ideal, and to account for the many cases where even pursuing the ideal is unrealistic. This suggests that we too might draw on ideal theory to address a question that lies at the heart of epistemology: the question of when ordinary agents, in ordinary circumstances, are entitled to believe the things they believe. The idea behind this suggestion is that a normative account of our epistemic position, non-ideal as it is, presupposes some conception of the ideal. This sort of methodology — ideal theory as foundational for one's broader theory — is familiar enough in other normative domains. In political philosophy, for instance, it is common to frame a theory of the just state around an account of what an ideally just state would look like. The project is not, of course, to describe a form of government suitable only for the gods. Nor is it supposed that the only just state would be one that perfectly satisfies the ideal theory. The goal is an understanding of what a just state would be for beings such as us, in a world such as this one. With some such picture of the ideal, calibrated against what is actually possible, we are able to think about what sort of political structures we might reasonably demand. This is precisely what I am recommending in the epistemic domain: that one take as foundational a conception of the human cognitive ideal, and then apply that to the question of what we ought to believe.

Epistemology today goes about things quite differently. Rather than describe an ideal and then consider how close we might come to achieving it, the modern epistemologist is often concerned with questions of threshold: exactly what divides knowledge from mere true belief. It would be as if political philosophers spent most of their time trying to define exactly where the borderline falls between the just and the unjust state, or as if ethicists focused on just precisely how good an act must be in order to count as praiseworthy. To be sure, there will be cases in the moral or political domain where such questions of threshold have practical relevance. But it would be odd to expect clear lines of demarcation, and odd to think that the principal task of normative theory is to discover those lines. The heart of my argument regarding epistemology, then, is that it should take seriously its normative dimension, and therefore should learn from the strategies found in other such domains.

Admittedly, there may be special reasons why epistemology pays special attention to boundary conditions. For it is often claimed that the difference between having and not having knowledge marks the divide between when one should and should not form beliefs, assert those beliefs, and act on them.[30] In that case determining the precise threshold of knowledge becomes something of immediate and constant importance to all of us. Now, one might well question whether the concept of knowledge really is what marks these divides, but suppose that it does, and suppose that this gives epistemic boundary conditions a special salience missing in other normative domains. Yet even if all of this is so, there is still reason to want an ideal epistemic theory, because — as in other normative cases — it seems plausible that decisions about such boundary conditions will turn at least in part on an account of the ideal at which we should aim. Consider, for instance, how much evidence an agent should have before forming a belief on the basis of the senses, or the testimony of others. These are notoriously difficult questions, and have admitted of many different kinds of answers, on different kinds of grounds. My suggestion is that we ought to ask in such cases what the epistemic ideal would be, for creatures such as us, in a world such as this. If we understood, for instance, what it would look like for us to be ideally situated, epistemically, with regard to sensory experience, then this will contribute to a clearer sense of what would count as adequate justification for ordinary perceivers, in ordinary cases.

As an example of how the project of idealization can help adjudicate such boundary disputes, consider the debate between externalists and internalists. Are we justified in believing propositions about the external physical world just in virtue of naively relying on our senses? The externalist says yes: cognitive agents need not have access to the factors that justify their beliefs; all that matters is that those factors be in place — that the process simply be, for instance, *reliable*. Internalists insist, in contrast, that agents have some kind of access to what justifies their beliefs.[31] This is a dispute about the boundary between knowledge and unjustified belief. It is natural to think that the state of affairs licensed by the externalist is farther from the cognitive ideal.

[30] For the link between knowledge and assertion see, most prominently, Williamson 2000, Ch. 11. For the thesis that what we know determines how we should act, see Hawthorne and Stanley 2008.

[31] For a collection of important papers on this debate, see Kornblith 2001. For a useful recent overview, see Madison 2010.

But is it *too* far to count as knowledge? How should we decide? If we focus on ideal theory, we might conclude that such questions do not need an answer at all — that such boundary disputes are *merely* lexicological questions, and that in reality there is nothing more than a spectrum of cases ranging from less to more ideal. But suppose, again, that we think the boundary matters a great deal, because we think that much in our practical lives rests on when we can be said to have knowledge. Ideal theory then seems poised to help, because it offers an account of what is ideally possible for human beings. If it turns out that there is no non-circular path toward establishing the reliability of the senses, then the externalist's version of knowledge had better be good enough, on pain of intellectual and practical paralysis. If, however, the senses can be given some sort of solid footing, outside of themselves, then perhaps we would be right to insist on the internalist's picture of knowledge. Of course, there is no agreement about such matters, which means that an idealized approach would hardly settle the dispute between internalists and externalists. Still, in so far as the dispute is not merely lexicological, but raises normative questions about the conditions under which it is right for us to form beliefs, reflection on the epistemic ideal would at least help clarify what the debate is about.

This last paragraph points toward an important disanalogy between epistemology and other normative domains. Disputes over the ethical or political ideal are largely disputes over the nature of the ideal at which we should aim. Deontologists differ from consequentialists, and consequentialists differ among themselves. Liberals differ from libertarians. There are also such disagreements among epistemologists — there is, for instance, dispute over whether inquiry should aim at truth or knowledge, or perhaps at understanding.[32] But the most controversial and far-ranging questions of idealized epistemology are concerned not with the ideal aims of inquiry in the abstract, but with the extent to which the achievement of these ideals is *possible for us*. Much of what one might object to in the ideal theories of Aristotle and Descartes, after all, lies not with the character of the ends at which they aim, but with the feasibility of our achieving those ends. Hence although we might hope to reach broad agreement about the cognitive ideal for human beings in the abstract, there is considerable disagreement about how much of that ideal is actually within our reach.

[32] See e.g. the papers in Haddock, Millar, and Pritchard 2009.

Inasmuch as ought implies can, such disagreements infect the question of what our epistemic aims ought to be.

As in other normative domains, however, an idealized epistemology holds out the hope not only of clarifying our actual practices, but of putting us in a position to critique those practices. Locke's political theories paved the way to the modern liberal state. Mill's utilitarianism has played a similar role in ethics. Reflection on the epistemic ideal might likewise lead us to reconsider what we regard as responsible and irresponsible in the doxastic domain. A very demanding conception of the ideal — according to which we are capable of providing and ideally ought to provide a thorough justification for everything we believe — might lead us toward higher expectations in everyday life. Believing without sufficient evidence, as we so often do, might come to seem increasingly irresponsible. On the other hand, a very pessimistic conception of the ideal — according to which non-circular arguments are rarely to be had — might encourage a more tolerant attitude toward belief. Instead of scorning those who hold various religious, ethical, and political views upon insubstantial evidence, we might indulgently regard such naive individuals as on more or less the same footing as everyone else.[33] Conceiving of epistemology as a branch of normative theory, rather than a branch of lexicology, might thus have the effect of making the field prescriptive rather than merely descriptive.

5. Ramifications of ideal theory

An idealized epistemology has the potential to illuminate many areas of current epistemological research. The previous section discussed the debate between internalists and externalists. Here, briefly, I survey some other domains.

Foundationalism versus coherentism: Debates over the justificatory structure of knowledge, somewhat peripheral in the current literature, move back to centre stage when epistemology is idealized. The beliefs of ordinary agents may form a hopeless tangle, hardly amenable to either a foundationalist or a coherentist analysis. Even so, ideal theory can ask the question of what would be best, and what would be possible for us. Aristotle and Descartes both took foundationalism as their ideal, but there is considerable room to wonder whether this is right. And once one focuses the discussion squarely on the ideal, it becomes

[33] See e.g. the argument in Alston 1991 for the evidential value of mystical experience, on the grounds that it is no worse, evidentially, than ordinary sensory experience.

possible to consider various more complex and less obvious possibilities, judging some varieties of coherence better than others, for instance, or perhaps even finding solutions other than the most familiar alternatives.[34]

Formal epistemology: It is an oddity of epistemology today that it divides sharply between those who pursue traditional problems of epistemology and those who pursue the subject in quantitative form, relying on the probability calculus and other such formal methods. This is not odd in itself, inasmuch as almost every area of philosophy gets pursued in various more or less formal ways, but it is odd here because it is very hard to see how these two sides of epistemology connect. Formal epistemology not only seems to have no bearing on the traditional approach, but also seems to have little relevance to ordinary epistemic agents at all. Hence it is a familiar criticism of Bayesianism, and other such technical approaches, that they describe a methodology that ordinary agents could not possibly pursue, because of the overwhelming mathematical complexity that would attend its application to any real world case.[35] Once epistemology becomes idealized, such criticisms can be answered, because we can see formal epistemologists as the descendants of Aristotle and Descartes, articulating a conception of the epistemic ideal without supposing that the method is readily practicable. This project is an interesting one in its own right, and it points toward the next step: taking the ideal theory, and considering what sorts of applications it might reasonably have for ordinary cognitive agents.

Virtue epistemology: Idealized epistemology shares with virtue epistemology the conviction that epistemology should be conceived of as a normative enterprise. If the best way to conceptualize our doxastic obligations is in terms of our cognitive and volitional dispositions, then an idealized epistemology will be a virtue epistemology — not surprisingly, given that the virtues traditionally have been thought

[34] For variations on coherentism, and an assessment of their tenability, see e.g. Huemer 2010. Klein (1999) has notoriously proposed one unfamiliar alternative: that justification might go on forever. For a view of justificatory structure at the opposite, pessimistic extreme, see Harman 2003, which argues for a general foundationalism according to which all one's beliefs are foundational, in the sense of being prima facie justified.

[35] For objections to Bayesianism along these lines, see e.g. Harman 1986, Ch. 3 and Kitcher 2002, pp. 396–97, who himself suggests that Bayesianism is best conceived of as a potentially useful idealization. Christensen 2004, Ch. 6 and 2007 explores in some detail the prospects for idealization.

of as *perfections* of an agent.[36] What the idealized framework might add to current discussions is a normative benchmark, a shared agreement about what human beings might ideally hope for. Without such a benchmark, virtue epistemology runs the risk of trafficking in mere platitudes: a reliable memory is good, and so is creativity, so is intellectual modesty, and so is self-confidence.[37] It would be not at all platitudinous, however, to describe the cognitive ideal for beings such as us, and to determine just how far toward that ideal one might reasonably be expected to go. The result might be a picture of cognitive excellence that looks very different — or simply much more nuanced — than what we ordinarily suppose.

Naturalized epistemology: If naturalizing epistemology means detaching it from normative questions, then its aims are quite distinct from those of idealized epistemology. But there is good reason to think that a purely descriptive, scientific investigation of human reasoning will shed light on normative issues. Focusing on the human ideal means focusing on what is ideal for beings with cognitive powers such as ours. Understanding these powers would require that we go beyond familiar platitudes about human reasoning such as that one should be open-minded, careful, modest, and yet have trust in oneself. What is needed is serious empirical research into the way perception and thought yield belief. In this way cognitive science might interact with epistemology by showing us what human beings are and are not capable of, and thereby helping to shape our understanding of the epistemic ideal. This in turn has the potential to influence normative expectations on us all, influencing our ideas of when cognitive agents can be said to be justified in their beliefs.[38]

Epistemology and metaphysics: Epistemology in its usual modern form stands rather far from the abstract concerns of the metaphysician. An idealized epistemology, in contrast, raises some of the central questions of metaphysics, inasmuch as reflection on the human epistemic ideal requires reflection not just on our cognitive capacities but also on the kind of world we live in. To be in an ideal epistemic state

[36] See Roberts and Wood 2007, pp. 65–9.

[37] Zagzebski 1996 proposes grounding an account of the intellectual virtues on 'the motivation for knowledge' (p. 167), although she acknowledges that this is not wholly satisfactory in the context of her theory, since she ultimately seeks to define knowledge in terms of those same virtues. A grounding in ideal theory seems more likely to capture the notion of an intellectual virtue, though it will make it less straightforward to define knowledge in terms of the virtues.

[38] For an interesting recent example of this sort of project, see Bishop and Trout 2005.

requires not only the perfect operation of one's cognitive powers, but also that those powers be directed to the right sorts of objects. When Plato offered his theory of Forms, he regarded them as the ideal objects of inquiry, and the *Republic*'s image of a divided line depicts the gap between those who have ideal knowledge and those who are merely lovers of sights and sounds. If Plato were right about the Forms, then an idealized epistemology would need to register that fact, and restrict ideal knowledge in the way that Plato recommended. Part of the project of Aristotle's *Posterior Analytics*, however, is to argue against the Platonic view, and to put in its place a different, but still worthy, object of ideal knowledge, the intrinsic essences of things. Perhaps this too is not right. Perhaps, as Locke thought, the real essences of things are not knowable to us, or perhaps there are no real essences at all. Are there instead laws of nature? Are there metaphysically necessary truths? Such questions should matter to any well-developed epistemology, because we can scarcely arrive at a view about what we are responsible for understanding without an account of what the proper candidates are. Idealized epistemology thus requires an understanding not just of what human beings are capable of, but of the kind of world we live in, at the deepest metaphysical level.

Philosophy of education: Recent philosophy, prizing the shallow over the deep, has held in great esteem the project of defining knowledge, while ignoring almost entirely the question of how we ought to educate people. When Plato, in contrast, identified the Form of the Good as the ultimate goal of human inquiry, he immediately turned to a lengthy discussion of how we ought to frame an ideal educational system.[39] That the philosophy of education ought to be a part of epistemology is alien to philosophy today. But when the topic is reconceived along the present, idealized lines, it becomes natural to ask how we might best bring people toward this ideal. The question of course demands the insights of a wide range of disciplines, but at its foundations requires an answer to the philosophical question of what the goal of education ought to be.

Epistemology and science: It is not clear that science has a particularly close relationship with the English word 'knowledge' — not any closer, at any rate, than does any other domain of human inquiry. Scientists seek knowledge, no doubt, but often they seem to want something more

[39] *Republic* VII. (I owe these remarks to the suggestion of an anonymous reader for *Mind*.) The philosophy of education has not been completely neglected by recent philosophers; for an impressive collection of work, see Siegel 2009.

than knowledge, and often they seem reconciled to achieving far less. Science does, however, seem to have especially close ties to idealized epistemology. This is not just because of the historical connection that runs from Aristotle's *epistēmē* through Descartes's *scientia* and on to modern science, but because science does — at least on its face — seem to have as its aim something like an epistemically ideal grasp of reality.[40] Perhaps this is a naive conception of how science works, or even of how it ought to work. But the project of understanding what science is and ought to be might profit from considering the scientific enterprise in the context of an idealized epistemology.

Variantism and invariantism: Many recent epistemologists have wondered whether the requirements on knowledge might somehow vary according to context.[41] The question, as it is usually pursued, seems mainly one of lexicology, but it need not be pursued in that way. Ideal theory begins by abstracting away from contextual effects, and imagines human beings with nothing better to do than become cognitively perfect. From ideal theory, we can go on to ask about what kind of evidence should be required in different kinds of situations, and in different domains of discourse. At this level, variantism gets taken for granted, and the question is exactly how different contexts should affect our normative expectations.

Scepticism: When epistemology focuses on the boundary conditions between knowledge and mere true belief, scepticism looks like a central issue. From the perspective of an idealized epistemology, however, sceptical worries are much harder to motivate. There will be the humanly ideal cognitive state, of course, which one may or may not ever achieve in a given domain, but beneath that there will simply be varying degrees of distance from that ideal. One may, within this framework, be more or less pessimistic about the strength of our position, but the Yes/No question of whether we have knowledge need not be asked.[42] If one

[40] See e.g. William Whewell's mid-nineteenth-century remark: 'Now there do exist among us doctrines of solid and acknowledged certainty, and truths of which the discovery has been received with universal applause. These constitute what we commonly term Science' (as quoted in van Fraassen 2002, p. 147). Accounts of science like these are what Kitcher (1993) calls the Legend.

[41] For influential statements of the view, see e.g. Cohen 1986 and DeRose 2002.

[42] Compare Quine 1981, p. 180: 'There is an obstacle in the verb "know". Must it imply certainty, infallibility? Then the answer is that we cannot. But if we ask rather how we are better warranted in believing one theory than another, our question is a substantial one'. BonJour 2010, in contrast, takes the seamless continuum in degrees of justification to show that the only tenable conception of knowledge is infallibilist. This seems no more plausible as a

persists in thinking that knowledge marks the crucial divide between when we should and should not form beliefs, then we can understand what the sceptics must claim: they must claim, as did the sceptics of old,[43] that we are wrong to believe the things we do — that our beliefs are too far from the ideal for us to be justified in maintaining them. Yet once we identify the sceptical question not as lexical but as normative — are we carrying out our doxastic practices in the right way, or at any rate a good-enough way? — then the negative sceptical answer looks just preposterous. Who could think that what epistemic analysis reveals is that our cognitive practices are so far from the ideal that we should simply stop forming beliefs? From a lexicological point of view, it is conceivable that systematic inquiry into the patterns of our knowledge ascriptions might arrive at a meaning for 'knowledge' that can never be satisfied.[44] But once one frames the conception of an idealized epistemology, and then distinguishes between the ideal itself and the real-world application of the ideal, it becomes clear that the burden of argument weighs heavily on the sceptic. The sceptic would have to establish not just that our beliefs fall sort of some lofty ideal (to which the label 'knowledge' has somehow been affixed), but that our beliefs fall short of even the minimal justification decently required for us to maintain them at all. Maybe so. Maybe everything we believe is more like faith than like knowledge. Ideally, we could show that this is not so, either by proving our beliefs true, or at least by proving them to be highly probable. We do not have to be in that ideal position, however, to think that the sceptic's position is incredible. Accordingly, it should be no surprise that scepticism is the least believed of major philosophical doctrines.[45]

theory of knowledge than it would be plausible, in other normative domains, to insist that only the ideal condition is praiseworthy.

[43] The classic text is Sextus Empiricus, *Outlines of Scepticism*. There is scholarly debate, however, over whether even Pyrrhonian sceptics such as Sextus want us to give up our beliefs — see, in particular, the papers collected in Burnyeat and Frede 1997.

[44] See, most famously in modern times, Unger 1975.

[45] *PhilPapers* 2009 found that a mere 1.8% of philosophy professors surveyed accept scepticism about the external world, with only 3% more leaning toward scepticism. Over 80% of respondents accepted or leaned toward 'non-sceptical realism' about the external world — a greater degree of consensus than was found on any other question in the survey.

6. Conclusion

The history of philosophy sheds light on the meta-epistemological question of how to go about studying knowledge. It shows that the modern preoccupation with the threshold separating knowledge from mere true belief is anomalous, in that epistemology has traditionally focused more on the ideal case than on defining boundaries. An idealized epistemology points toward a normative, prescriptive rather than descriptive enterprise. The ultimate aim of epistemology, so conceived, would be an ethics of belief — a theory of the conditions under which it is right to believe and assert the things we believe. Such a theory can be grounded, as in other normative domains, on an account of the ideal for beings such as us, in a world such as this. An account of the human ideal is interesting in its own right, and also would help us evaluate what sort of justification is appropriate for ordinary agents, across the different domains of belief, in varying circumstances. This part of epistemology has been little developed, but might have the kind of impact on our ordinary epistemic practices that moral theory has had on the ethical domain. Perhaps the results would turn out to match fairly closely with the linguistic intuitions of English speakers about the word 'knowledge'. Or maybe not, and that might turn out to be even more interesting. Hitherto, epistemologists have only interpreted the word in various ways. The point, perhaps, should be to change it.[46]

References

Albert the Great 1651: *Opera omnia*, ed. P. Jammy. Lyon.

Alston, William P. 1991: *Perceiving God: A Study in the Epistemology of Religious Experience*. Ithaca, NY: Cornell University Press.

—— 2005: *Beyond 'Justification': Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.

Aristotle 1993: *Posterior Analytics*, tr. J. Barnes, second edition. Oxford: Clarendon Press.

Barnes, Jonathan 1969: 'Aristotle's Theory of Demonstration'. *Phronesis*, 14, pp. 123–52.

—— 1993: 'Introduction'. In Aristotle 1993, pp. xi–xxii.

[46] This paper has benefited greatly from feedback from Yuval Avnur, Dominic Bailey, David Barnett, Robert Hanna, Thomas Metcalf, Jason Potter, Michael Sechman, Christopher Shields, and an audience at CU/Boulder.

Benson, Hugh H. 2000: *Socratic Wisdom: The Model of Knowledge in Plato's Early Dialogues*. New York: Oxford University Press.
Berti, Enrico (ed.) 1981: *Aristotle on Science: The Posterior Analytics*. Padua: Edizioni Antenori.
Bishop, Michael A. and J. D. Trout 2005: *Epistemology and the Psychology of Human Judgment*. Oxford: Oxford University Press.
BonJour, Laurence 2010: 'The Myth of Knowledge'. *Philosophical Perspectives*, 24, pp. 57–83.
Broughton, Janet 2002: *Descartes's Method of Doubt*. Princeton: Princeton University Press.
Broughton, Janet and John Carriero (eds) 2008: *A Companion to Descartes*. Oxford: Blackwell.
Burnyeat, Myles 1981: 'Aristotle on Understanding Knowledge'. In Berti 1981, pp. 97–139.
Burnyeat, Myles and Michael Frede 1997: *The Original Sceptics: A Controversy*. Indianapolis: Hackett.
Butts, Robert E. and Joseph C. Pitt (eds) 1978: *New Perspectives on Galileo*. Dordrecht: Reidel.
Christensen, David 2004: *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford: Clarendon Press.
—— 2007: 'Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals'. *Oxford Studies in Epistemology*, 2, pp. 3–31.
Cohen, I. Bernard 1980: *The Newtonian Revolution*. Cambridge: Cambridge University Press.
Cohen, Stewart 1986: 'Knowledge and Context'. *Journal of Philosophy*, 83, pp. 574–83.
Craig, Edward 1990: *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
Dancy, Jonathan, Ernest Sosa, and Matthias Steup (eds) 2010: *A Companion to Epistemology*, second edition. Oxford: Blackwell.
de Pace, Anna 1993: *Le matematiche e il mondo: ricerche su un dibattito in Italia nella seconda metà del Cinquecento*. Milan: Francoangeli.
DePaul, Michael and Linda Zagzebski (eds) 2003: *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Clarendon Press.
DeRose, Keith 2002: 'Assertion, Knowledge and Context'. *Philosophical Review*, 111, pp. 167–203.
Descartes, René 1897: *Oeuvres de Descartes*, ed. C. Adam and P. Tannery. Paris: Cerf.

——1991: *The Philosophical Writings of Descartes*, eds. J. Cottingham, R. Stoothoff, D. Murdoch, and A. Kenny. Cambridge: Cambridge University Press.
Everson, Stephen (ed.) 1990: *Epistemology*. Cambridge: Cambridge University Press.
Fine, Gail 1990: 'Knowledge and Belief in *Republic* V–VII'. In Everson 1990, pp. 85–115.
Gilbert, Neal Ward 1960: *Renaissance Concepts of Method*. New York: Columbia University Press.
Glanvill, Joseph 1661: *The Vanity of Dogmatizing*. London.
Greco, John and Ernest Sosa (eds) 1999: *The Blackwell Guide to Epistemology*. Oxford: Blackwell.
Haddock, Adrian, Alan Millar, and Duncan Pritchard (eds) 2009: *Epistemic Value*. Oxford: Oxford University Press.
Harman, Gilbert 1986: *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.
——2003: 'Skepticism and Foundations'. In Luper 2003, pp. 1–11.
Hawthorne, John and Jason Stanley 2008: 'Knowledge and Action'. *Journal of Philosophy*, 105, pp. 571–90.
Hendricks, Vincent F. and Duncan Pritchard (eds) 2006: *New Waves in Epistemology*. Aldershot: Ashgate.
Hookway, Christopher 2003: 'How to be a Virtue Epistemologist'. In DePaul and Zagzebski 2003, pp. 183–202.
Huemer, Michael 2010: 'Foundations and Coherence'. In Dancy, Sosa, and Steup 2010, pp. 22–33.
Irwin, Terence 1988: *Aristotle's First Principles*. Oxford: Clarendon Press.
John Duns Scotus 1639: *Opera Omnia*, ed. L. Wadding. Lyon.
Kaplan, Mark 1985: 'It's Not What You Know that Counts'. *Journal of Philosophy*, 82, pp. 350–63.
Kitcher, Philip 1993: *The Advancement of Science: Science without Legend, Objectivity without Illusions*. New York: Oxford University Press.
——2002: 'Scientific Knowledge'. In Moser 2002, pp. 385–407.
Klein, Peter D. 1999: 'Human Knowledge and the Infinite Regress of Reasons'. *Philosophical Perspectives*, 13, pp. 297–325.
Kornblith, Hilary 1999: 'In Defense of a Naturalized Epistemology'. In Greco and Sosa 1999, pp. 158–69.
——(ed.) 2001: *Epistemology: Internalism and Externalism*. Cambridge, MA: MIT Press.

Lesher, J. H. 2001: 'On Aristotelian 'Επιστήμη as "Understanding"'. *Ancient Philosophy*, 21, pp. 45–55.
Locke, John 1975: *An Essay concerning Human Understanding*, ed. P. H. Nidditch. Oxford: Clarendon Press.
Ludlow, Peter 2005: 'Contextualism and the New Linguistic Turn in Epistemology'. In Preyer and Peter 2005, pp. 11–50.
Luper, Steven (ed.) 2003: *The Skeptics: Contemporary Essays*. Aldershot: Ashgate.
McMullin, Ernan 1978: 'The Conception of Science in Galileo's Work'. In Butts and Pitt 1978, pp. 209–58.
Madison, B. J. C. 2010: 'Epistemic Internalism'. *Philosophy Compass*, 5, pp. 840–53.
Moser, Paul K. (ed.) 2002: *The Oxford Handbook of Epistemology*. Oxford: Oxford University Press.
Newman, Lex 2010: 'Descartes' Epistemology'. The Stanford Encyclopedia of Philosophy [online encyclopedia], Fall 2010 Edition, ed. Edward N. Zalta, <http://plato.stanford.edu/archives/fall2010/entries/descartes-epistemology/>.
Owens, David 2000: *Reason without Freedom: The Problem of Epistemic Normativity*. London: Routledge.
——2008: 'Descartes's Use of Doubt'. In Broughton and Carriero 2008, pp. 164–78.
Pasnau, Robert 2010: 'Medieval Social Epistemology: *Scientia* for Mere Mortals'. *Episteme*, 7, pp. 25–41.
Perfetti, Stefano 2008: 'Pietro Pomponazzi'. The Stanford Encyclopedia of Philosophy [online encyclopedia], Fall 2008 Edition, ed. Edward N. Zalta, <http://plato.stanford.edu/archives/fall2008/entries/pomponazzi/>.
PhilPapers 2009: 'The PhilPapers Survey'. <http://philpapers.org/surveys>, accessed 26 July 2012.
Preyer, Gerhard and Georg Peter (eds) 2005: *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. Oxford: Clarendon Press.
Ptolemy 1984: *Almagest*, tr. G. J. Toomer. New York: Springer.
Quine, W. V. O. 1969a: 'Epistemology Naturalized', in his 1969b, pp. 69–90.
——1969b: *Ontological Relativity and Other Essays*. New York: Columbia University Press.
——1981: *Theories and Things*. Cambridge, MA: Harvard University Press.

——1990: *Pursuit of Truth*. Cambridge, MA: Harvard University Press.
Reynolds, Philip L. 2006: 'The Infants of Eden: Scholastic Theologians on Early Childhood and Cognitive Development'. *Mediaeval Studies*, 68, pp. 89–132.
Riggs, Wayne 2006: 'The Value Turn in Epistemology'. In Hendricks and Pritchard 2006, pp. 300–23.
Roberts, Robert C. and W. Jay Wood 2007: *Intellectual Virtues: An Essay in Regulative Epistemology*. Oxford: Clarendon Press.
Sanches, Francisco 1988: *That Nothing Is Known (Quod Nihil Scitur)*, ed. and tr. E. Limbrick and D. F. S. Thomson. Cambridge: Cambridge University Press.
Sextus Empricus 2000: *Outlines of Scepticism*, tr. J. Annas and J. Barnes. Cambridge: Cambridge University Press.
Siegel, Harvey (ed.) 2009: *The Oxford Handbook of Philosophy of Education*. Oxford: Oxford University Press.
Sinnott-Armstrong, Walter (ed.) 2004: *Pyrrhonian Skepticism*. New York: Oxford University Press.
Sorell, Tom, G. A. J. Rogers, and Jill Kraye (eds) 2010: *Scientia in Early Modern Philosophy: Seventeenth-Century Thinkers on Demonstrative Knowledge from First Principles*. Dordrecht: Springer.
Sosa, Ernest 1997: 'How to Resolve the Pyrrhonian Problematic: A Lesson from Descartes'. *Philosophical Studies*, 85, pp. 229–49.
——2007: *A Virtue Epistemology, vol. 1: Apt Belief and Reflective Knowledge*. Oxford: Clarendon Press.
——2009: 'Replies to Brown, Prichard and Conee'. *Philosophical Studies*, 143, pp. 427–40.
Stich, Stephen 1990: *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.
Taylor, C. C. W. 1990: 'Aristotle's Epistemology'. In Everson 1990, pp. 116–42.
Thayer, H. S. 1979: 'Aristotle on the Meaning of Science'. *Philosophical Inquiry*, 1, pp. 87–104.
Thomas Aquinas 1882–: *Opera Omnia*. Rome: Commissio Leonina.
Unger, Peter 1975: *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
van Fraassen, Bas 2002: *The Empirical Stance*. New Haven: Yale University Press.

Weatherson, Brian 2008: 'Deontology and Descartes' Demon'. *Journal of Philosophy*, 105, pp. 540–69.
Westfall, Richard S. 1977: *The Construction of Modern Science: Mechanisms and Mechanics*. Cambridge: Cambridge University Press.
Williams, Michael 2004: 'The Agrippan Argument and Two Forms of Skepticism'. In Sinnott-Armstrong 2004, pp. 121–46.
Williamson, Timothy 2000: *Knowledge and Its Limits*. Oxford: Oxford University Press.
Wolterstorff, Nicholas 1996: *John Locke and the Ethics of Belief*. Cambridge: Cambridge University Press.
Zagzebski, Linda 1996: *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge*. New York: Cambridge University Press.

