

KNOWLEDGE FIRST EPISTEMOLOGY

Timothy Williamson

Introduction

“Knowledge first” is a slogan for epistemology that takes the distinction between knowledge and ignorance as the starting point from which to explain other cognitive matters. It reverses the direction dominant in much twentieth-century epistemology, which treated belief as explanatorily prior to knowledge, attempting to analyze knowledge as belief that meets further conditions, such as truth and justification. By contrast, a knowledge first epistemologist might treat believing something as treating it as if one knew it.

The most striking difference between knowledge and belief is that knowledge entails truth while belief does not. There is false belief but no false knowledge. Some people *believe* that Africa is a single country, but since it is false that it is a single country, they do not *know* that it is a single country. They just believe falsely that they know that Africa is a single country. In this sense, all knowledge but not all belief is successful. Thus knowledge first epistemology gives explanatory priority to success. This does not mean that belief first epistemology gives priority to failure. Rather, it gives explanatory priority to conditions that are neutral between success and failure: some beliefs constitute knowledge, others are false.

Most epistemologists agree that while knowing entails believing truly, believing truly does not entail knowing. Someone does not know something he believes truly on the say-so of his guru, who invents things to tell him at random without regard to their truth or falsity. Although merely believing truly involves a sort of success—getting the answer right—it also involves, unlike knowing, a sort of cognitive malfunction. Thus knowledge is a more full-blooded success condition than true belief. Knowledge first epistemology understands cases of cognitive malfunctioning in terms of their deviation from cases of cognitive functioning, as opposed to treating the two kinds of case more symmetrically.

Similar contrasts arise in the philosophy of action, where intentionally doing something is the full-blooded success condition. It stands to the world-to-mind direction of fit as knowing stands to the mind-to-world direction of fit. Believing corresponds to trying, the neutral condition that obtains in both cases of success and cases of failure. Falsely believing corresponds to trying and failing. Merely believing truly without knowing corresponds to trying and doing the intended thing but not as intended: for instance, you hit the intended target because your intention so agitated you that your trigger

finger slipped. Knowledge first epistemology corresponds to a philosophy of action that understands cases of trying to do something without intentionally doing it in terms of their deviation from cases of intentionally doing it, again as opposed to treating the two kinds of case more symmetrically.

We cannot reasonably expect to decide between the two directions of explanation on the basis just of such abstract characterizations of the difference. Rather, we must explore the strengths and weaknesses of each approach in application to more specific problems.

The Analysis of Knowledge

One traditional aim of epistemology was to provide an “analysis of knowledge.” Sometimes this meant analyzing the *concept* of knowledge, perhaps by providing a more complex sentence synonymous with “S knows that P” and thereby breaking the concept down into its constituent concepts, expressed by the constituent phrases of the complex sentence. At other times what it meant was analyzing the underlying *nature* of knowledge, which might be done without providing such a synonym, since that nature might have to be discovered through arduous investigation, without having already been written into the structure of our concepts or the meanings of our words. Either way, an analysis of knowledge was supposed to provide a non-circular statement of necessary and sufficient conditions for S to know that P. A paradigm candidate for such an analysis of knowledge is the once-popular “JTB” account, on which S knows that P if and only if S has a justified true belief that P.

Notoriously, the JTB analysis was refuted by Edmund Gettier, who gave examples to show that having a justified true belief is not sufficient for knowing. In the following decades, many attempts were made to repair the JTB analysis, by adding further conditions or modifying the justification condition. In each case, most epistemologists regard the attempted repair as unsuccessful, because it falls to other counter-examples, although some proposed analyses still have supporters. One reaction to this depressing track record is that it only shows that we must go on looking; analyzing knowledge was never going to be an easy task. Another reaction is that the track record constitutes good inductive evidence that there is no such analysis to be found: knowledge is not a combination of belief and other conditions more basic than knowledge itself; belief is not prior to knowledge. To decide between these contrasting reactions, we must widen the terms of the debate.

Was there ever good reason to *expect* non-circular necessary and sufficient conditions for knowledge? On some theories of concepts, all concepts are built up out of a small basic stock of simple concepts. A typical motivation for such a theory is empiricist; the simple concepts are supposed to have a peculiarly direct relation to experience. The concept of knowledge is not a promising candidate to be one of those few simple building-blocks. However, such theories of concepts are very hard to sustain on current evidence; very few words in natural languages seem to have that sort of semantic complexity. Nor do epistemologists usually make explicit appeal to theories of concepts in defending the program of seeking analyses of the concept of knowledge. Although very general presuppositions about the nature of concepts may have helped to motivate the idea that conceptual analysis is the goal of philosophy, they do not now provide good reason to expect a non-circular analysis of the concept of knowledge.

Another suggested reason is this. We already have some non-trivial *necessary* conditions for knowledge, such as truth and belief. Why then can't knowledge be analyzed as true belief plus whatever must be added to true belief to get knowledge? As it stands, "whatever must be added to true belief to get knowledge" is a circular specification, since it explicitly mentions knowledge; the problem is that we have been given no reason to expect that the circularity can be eliminated. The usual analogy here is with color concepts. Although being colored is a non-trivial necessary condition for being red, that gives us no reason to expect being red to have an analysis as being colored plus whatever must be added to being colored to get being red, under a different specification of the latter that does not mention red.

A third suggested reason is this. According to *internalist* theories of mind, a purely mental state has no necessary consequences for the world outside the mind: a brain in a vat can be in the same purely mental states as someone who is perceiving and acting on the environment in the normal way. On this view, knowing that it is snowing is not a purely mental state, because it has the non-mental necessary consequence that it is snowing. By contrast, believing that it is snowing has no obvious non-mental necessary consequences, and so might yet count as a purely mental state. One might therefore expect the hybrid, impurely mental state of knowing that it is snowing to be analyzable into purely mental components (such as believing that it is snowing) and purely non-mental components (such as it being true that it is snowing). However, such internalism about the mind is ill-supported. Many states that would ordinarily be regarded as mental have necessary consequences for the external world. For instance, loving Mary and hating Mary seem to be mental states, but being in those states has the necessary consequence that your apparently normal causal connections to Mary are not totally illusory. There are strong arguments for *semantic externalism*, on which the content of intentional states such as belief typically depends on causal connections between the thinker and the environment, so that those states have necessary consequences for the external world. For instance, if all the causal connections (direct and indirect) of one's "snow" beliefs had been with masses of tiny flying moths that just look like snow, they would not really have been beliefs that it was snowing; they would have been beliefs in a different proposition about that other phenomenon instead.

Faced with such examples, internalists might retort that such states are not *purely* mental. Indeed, they might be forced to concede that hardly any of the mental states we explicitly attribute in natural language are purely mental. The purely mental states become hypothetical constructs. Even granted that someone who believes that it is snowing is also in various internal states without necessary consequences for the external world, it is unclear why any of those internal states should count as mental. In saying what is peculiarly mental in the states of loving or hating Mary and of believing that it is snowing, it is natural to focus on their *aboutness*: loving or hating Mary is about Mary (she is the one on whom the requitedness or unrequitedness of the emotion depends); believing that it is snowing is about snow (it is the stuff on which the truth or falsity of the belief depends). That is the *point* of such states. Take away their aboutness and you take away what is most mental in them. Yet their necessary aboutness was what made them non-internal. Thus the internalist strategy of postulating internal states to be the underlying purely mental states treats as inessential exactly those features that seem central to the mentality of ordinary "impure" mental states.

Far from compromising "pure" mentality, semantic externalism arguably articulates a deep insight into the nature of mentality. But if mental states are typically ways of

relating to one's environment, the truth-entailing nature of knowledge is no obstacle to its being a mental state in a quite unstretched sense. Knowing that it is snowing essentially involves awareness of one's environment in a way that mere believing that it is snowing does not: but that enhanced awareness hardly makes knowing *less* mental than believing. On this view, knowing is just as mental a state as believing. Its mental consequences are not exhausted by believing, or even by justified believing. Such externalism about mental attitudes themselves, not just their contents, is a natural development of knowledge first epistemology. If knowing is not an impure mental state, the argument lapses for expecting it to be analyzable into purely mental and purely non-mental components. Further challenges to the conception of knowing as a mental state will be discussed in later sections.

The arguments that knowledge must be analyzable all rest on highly contentious assumptions. Further inductive evidence for its unanalyzability comes from the track record of failure in programs for analyzing other philosophically central properties and relations, such as meaning and causation. Although such programs often give interesting partial results, attempts to state strictly necessary and sufficient conditions in non-circular terms typically lead to a regress of ever more complex analyses and ever more complex counter-examples. This further reduces the probability that knowledge has such an analysis in terms of belief. Of course, it also reduces the probability that belief has such an analysis in terms of knowledge.

Knowledge first epistemology challenges the project of analyzing knowledge in another way too. Analyses are supposed to be non-circular. In particular, the inclusion of the justification condition in analyses such as JTB presupposes that it is suitably independent of knowledge itself. As will appear later, knowledge first epistemology questions that independence.

Problems of Access

A salient feature of the state of knowing is that one is not always in a position to know whether one is in it. For example, if you are a victim of a clever hoaxer who makes it appear to you to be raining when it isn't, then, for all you are in a position to know, you know that it's raining. Really, you don't know that it's raining, but you are not in a position to know that you don't know that it's raining. Arguably, one can also know without being in a position to know that one knows (see below). Whether one knows is not fully open to introspection. If core mental states such as believing, desiring and feeling sensations *are* fully open to introspection, that would be a new reason to deny that knowing is a "pure" mental state.

Our limited access to whether we know might be used as an objection to knowledge first epistemology in other ways too. Some want epistemology to be *operational*, in the sense of providing rules for rational inquiry such that one is always in a position to know whether one is complying with them. A classic example of a non-operational rule is the useless advice "Believe what is true!"; if you knew what was true, you wouldn't need the advice. By contrast, the rule "Proportion your belief to the evidence!" sounds more useful; it seems to have a better chance of being operational. Non-operational epistemology has a third-personal aspect; it evaluates the agent's epistemic status as it were from the standpoint of an external observer, with access to information unavailable to the agent. Operational epistemology is first-personal in a corresponding sense. Whether one knows is a question for non-operational epistemology, whereas questions

of rationality, justification and evidence seem to belong to operational epistemology. This suggests an objection to knowledge first epistemology: the inapplicable should not have priority over the applicable.

The reply on behalf of knowledge first epistemology is that states fully open to introspection and fully operational rules are an illusion. Any non-trivial state is such that one can be in it without being in a position to know that one is in it. An argument for that conclusion goes like this. Call a state *luminous* if whenever one is in it one is in a position to know that one is in it. For any non-trivial state, one can change from being in it to not being in it through a very gradual process. Since our powers of discrimination are limited, in the last moments of the process at which one is still in the state, one cannot discriminate how one is (in the relevant respects) from how one is in the first moments at which one is no longer in the state. In the former moments, one is in the state without being in a position to know that one is in it. Consequently, the state is not luminous (this can all be made more precise). Thus any luminous state is trivial. In particular, states of believing, desiring, and feeling sensations are non-luminous because non-trivial. One can feel pain without being in a position to know that one feels pain. One can believe something without being in a position to know that one believes it. Equally, since knowing is a non-trivial state, one can know something without being in a position to know that one knows it.

In many cases, there are other independent arguments for the non-luminosity of those states. For instance, postulating unconscious beliefs and desires might help to explain someone's actions. It is not even clear that one is always in at least as good a position to determine whether one believes something as to determine whether one knows it. For instance, someone who has just slowly and painfully lost her religious faith might be in a better position to know that she does not know that there is a god than she is to know that she does not believe that there is a god. Although one is *often* in a position to know without observing oneself "from the outside" whether one believes, desires or feels something, one is also often in a position to know without such further observation whether one knows something.

If no non-trivial mental states are fully open to introspection, then the fact that knowing is not fully open to introspection does not make it less mental than believing, desiring, and feeling sensations. Similarly, if one is always in a position to know whether one is complying with an operational rule, then the state of complying with that rule is luminous, and therefore trivial, in which case the rule itself must be trivial. Thus a fully operational epistemology is an impossible ideal. Even matters of rationality, justification and evidence are non-luminous.

The widespread temptation to conceive an agent's evidence as consisting of their present subjective states might, itself, result from the assumption that agents must always be in a position to know what evidence they have. If one rejects that assumption, one can take seriously a straightforward conception of agents' total evidence as comprising the total content of their knowledge ($E = K$), a view much closer to the way scientists treat evidence—and a characteristic thesis of knowledge first epistemology.

The equation $E = K$ has controversial implications for skepticism. Skeptics often argue that we cannot know whether we are in an everyday scenario or a corresponding skeptical scenario, because we have the same evidence in each. Even some anti-skeptics concede such sameness of evidence. But a defender of common sense can contest that assumption. In the everyday scenario, I know that I have hands. In the skeptical scenario, I lack hands and therefore do not know that I have them. Thus my knowledge

differs between the two scenarios. Given $E = K$, that difference in my knowledge constitutes a difference in my evidence: the skeptical argument rests on a false premise. Part of my cognitive deprivation in the skeptical scenario is that I am not aware that I have less evidence than in the everyday scenario. Of course, that does not exhaust the skeptic's resources. The skeptical argument can be reworked in various ways. Nevertheless, once one is alert to the way in which skeptical reasoning can conceal such unargued epistemological assumptions, one is better placed to dispute the reasoning.

Questions of justification and rationality concern doing the best one can on one's limited evidence; given $E = K$, that means doing the best one can on one's limited knowledge. If so, they are not prior to questions of knowledge. What you are justified or rational in believing depends on what you know. Thus the occurrence of a justification or rationality condition in an analysis of knowledge risks circularity.

Such externalism about evidence also has implications for the epistemology of philosophy itself. Philosophers have implicitly tended to conduct their reasoning as though their evidence were restricted to data on which all the rival theories can agree. For example, if B denies that there are macroscopic physical objects, then A might cite as evidence not the "question-begging" non-psychological fact that there are rocks but only the psychological fact that we have an "intuition" that there are rocks. Of course, if A is trying to change B's mind, it is normally pointless for A to cite as evidence facts that B denies. However, if A's goal is to find the truth on some matter for herself, not to persuade B of it, why shouldn't she use her knowledge that there are rocks, rather than just her knowledge that she has an intuition that there are rocks. The non-psychological fact might be far more relevant than the psychological one to her question. That might sound like dogmatically rejecting B's point of view. However, the alternative strategy of using only universally accepted evidence is hopeless, since there is virtually no such evidence. Some philosophers will deny that anyone has an intuition that there are rocks, perhaps because they are eliminativists about such mental states, while accepting that there are rocks. To engage in philosophy is to engage in controversy. Of course, that does not mean that there are no constraints. $E=K$ says that your evidence comprises what you do know, not what you believe you know. That it is often hard to recognize the difference in practice might be an inescapable aspect of our cognitive predicament: even the epistemology of philosophy must be non-operational.

Reasons for Action

Here is yet another objection to knowledge's being a mental state. Genuine mental states play a role in causal explanations of rational action. Such explanations cite the agent's beliefs and desires. They do not cite the agent's knowledge as such, for even if some of those beliefs do in fact constitute knowledge, that is irrelevant to their role in causing the action. If you want a drink of water and believe that your glass contains water, then other things being equal you will drink from your glass, whether or not you really know that it contains water. But if believing screens out knowing in the causal explanation of action, then knowing has no serious claim to be a mental state over and above believing.

One problem with that argument is its assumption that causal explanations of action are always of action "at the next instant." If any time elapses from when the agent is in the original mental states of interest to when the action is completed, the difference between knowing and believing might matter causally. In the example above, if

a teetotaler falsely believes a glass of gin to contain water, he might spit the gin out as soon as he tastes it without drinking any; that is far less likely to happen if he knows the glass to contain water. Even true belief tends to be less persistent than knowledge, since true beliefs can depend on false ones in ways that knowledge cannot; the former are therefore more vulnerable than the latter to abandonment through discovery of such false assumptions before the action has been completed. Thus believing does not always screen out knowing in the causal explanation of action.

From a normative perspective, knowing is arguably *more* relevant to rational action than believing is. Suppose that, in deciding whether to take certain pills, I treat as a reason for taking them that they will alleviate my hay fever symptoms. If you point out that I don't know whether they will alleviate my symptoms, you reveal a defect in my decision-making with respect to that premise. By contrast, if I do know that the pills will alleviate my symptoms, my decision-making is not defective with respect to that premise. Mere belief in the premise is not what matters, since that is consistent with the belief's being both false and irrational. If I don't know the premise but have strong evidence that I do know it, that may be a good excuse for having treated it as a reason, but that does not make the decision-making non-defective, otherwise no excuse would be needed. In some cases of extreme urgency, I may rely on premises that I know myself not to know even though they are highly probable on my evidence. The urgency is another excuse for defective reasoning; it does not remove the defect. To do that, I might have to replace the premise that the medicine will alleviate my symptoms by the premise, which I might know, that it is highly probable on my evidence that the medicine will alleviate my symptoms. Such a view of reasons for action is another natural development of knowledge first epistemology.

Belief and Assertion

If one was never willing to treat "P" as a reason for doing anything, even when one believed "If P, the house is on fire," then one's commitment to "P" would be too weak to count as belief. Given the connection between knowledge and reasons for action, it follows that if one believes that P, one is willing to make decisions in a way that is non-defective only if one knows that P. This suggests a simpler cognitive norm: believe that P only if you know that P. Call that *the knowledge rule for belief* (KRB).

Versions of Moore's paradox confirm KRB. Something is wrong with believing "It's raining and I don't know that it's raining," even though the belief could very easily be true. Although the two conjuncts are logically consistent with each other, the second undermines the first. KRB explains how. It tells me to believe the conjunction only if I know the conjunction. But I cannot know the conjunction, for to know it I must know the first conjunct ("It's raining"), in which case the second conjunct is false, making the whole conjunction false and therefore unknown. Thus in believing "It's raining and I don't know that it's raining," one automatically violates KRB. If KRB is not in force, it is quite unclear what is wrong with believing such conjunctions.

If a justified belief is one that satisfies the norms for belief, then by KRB only knowledge constitutes justified belief. In particular, there are no justified false beliefs. This might rescue the letter of the JTB account of knowledge from Gettier, since his counterexamples rely on the assumption that there are justified false beliefs, which generate justified true beliefs because the deceived subject can competently deduce true conclusions from false premises and thereby extend his justification from the latter to the former.

However, KRB certainly does not rescue the spirit of JTB. Gettier treated the notion of justification as proponents of JTB had intended; they made truth an independent condition precisely because they did not think that it followed from the other two. KRB makes JTB effectively circular as an analysis of knowledge, since the relevant norm invokes knowledge. Given KRB, the sense of “justification” that Gettier shared with proponents of JTB seems to involve excuse as much as justification: a “justified” false belief is a belief for which the agent has a good excuse, which may consist in strong but misleading evidence that the belief constitutes knowledge.

In the sense in which only knowledge constitutes full justification, one can also define a graded sense of justification, using the equation $E = K$. For the probability of a proposition on one’s evidence becomes its probability conditional on the total content of one’s knowledge. That probability can be interpreted as one’s degree of justification to believe the proposition. The structure of such evidential probabilities can be explored through mathematical models of epistemic logic, sometimes with unexpected results. For example, one can know something even if the probability on one’s own evidence that one knows it is close to zero.

Assertion is an analog for speech of belief. It has an analog of KRB, the knowledge rule for assertion (KRA): assert that P only if you know that P. Just as Moore-paradoxical beliefs provide evidence for KRB, so Moore-paradoxical assertions provide evidence for KRA. Moreover, KRA can be supported by a wide range of linguistic data. As with KRB, opponents of KRA tend to object that it sets an unrealistically high standard. For proponents of KRA, such objections, too, mistake excuses for justifications. If the anti-luminosity argument is correct, one can violate *any* non-trivial norm for assertion or belief without being in a position to know that one is violating it.

Both KRA and KRB are natural developments of knowledge first epistemology, for they imply that even central non-truth-entailing cognitive attitudes such as asserting and believing depend normatively on knowing.

Perception and Memory

Knowing is not the only truth-entailing attitude. If one perceives that it is raining, it is raining. If it is not raining, one at most seems to oneself to perceive that it is raining. Similarly, if one remembers that it was raining, it was raining. If it was not raining, one at most seems to oneself to remember that it was raining. Perceiving and remembering are further instances of truth-entailing mental attitudes, if knowing is already one.

Arguably, perceiving and remembering entail not just truth but knowledge: if one perceives that it is raining, one knows that it is raining; if one remembers that it was raining, one knows that it was raining. Objections to these claims focus on cases in which one’s perceptual or memory apparatus is in fact functioning properly, but misleading evidence casts doubt on whether it is functioning properly. However, it is not clear in such cases that one perceives or remembers that P without knowing that P. One may perceive the rain without perceiving *that* it is raining, or remember the rain without remembering *that* it was raining, if one fails to recognize the rain as rain. Alternatively, if one does recognize the rain as rain, presumably one does know that it is or was raining. That the misleading evidence casts doubt on whether one knows does not mean that one does not know.

If perceiving and remembering are knowledge-entailing mental attitudes, a wider conjecture suggests itself: that knowing is the most general truth-entailing mental

attitude, the one you have to a proposition if and only if you have any truth-entailing mental attitude to it at all. The defense of such a conjecture requires interpreting “mental attitude” appropriately: for example, having forgotten that P had better not count as a mental attitude towards the proposition that P. At least, very many truth-entailing attitudes are closely related to knowing. If you are conscious or aware that P, you know that P; if you learn, discover, or recognize that P, you come to know that P. These examples illustrate how hard it is to imagine a cognitive life in which knowing plays no part. It would be a life in which one was never conscious or aware that something was the case, and never perceived, remembered, learned, discovered, or recognized that something was the case.

Knowledge first epistemology has important similarities to *disjunctivist* accounts of perception, which are sophisticated modern versions of “naive realism.” Disjunctivists often explain failure (illusion, hallucination) in terms of its relation to success (genuine perception), rather than treating the two cases on a par. Typically, they also hold that there is a mental state (in an unstretched sense) that one is in when and only when one is genuinely perceiving: one *takes in* the world. Consequently, they deny that one’s mental state can be exactly the same in cases of failure as in cases of success, with the only difference consisting in external causal relations inessential to that state. They might well agree that one has more perceptual evidence in the good case than in the bad case: it is just that sometimes, when one is in the bad case, for all one knows one is in the good case. Furthermore, disjunctivists are often sympathetic to the claim that perceiving that P entails knowing that P. To a first approximation, therefore, disjunctivism about perception is simply knowledge first epistemology as applied to perception. However, some qualifications are needed.

First, disjunctivism is usually advanced as a theory about perception in its own right, without appeal to wider claims of knowledge first epistemology.

Second, disjunctivism takes its name from ideas like this: a state of having things perceptually appear to one a certain way is a disjunction of two radically different disjuncts; *either* one is in the successful state *or* one is in the unsuccessful state. The impression is sometimes given that the two disjuncts are being characterized independently of each other, which would undermine the explanatory priority of the first disjunct over the second. However, there is little prospect of grouping together all instances of having it perceptually appear to one that P without perceiving that P except under such a negative characterization; the second disjunct is not really being characterized independently of the first and the supposedly disjunctive state. The strictly *disjunctive* aspect of disjunctivism might be its least defensible feature. But even if the letter of disjunctivism is rejected, its spirit is retained in the idea that success in perception cannot be analyzed as a combination of success-neutral perceptual appearances with external causal perceptions inessential to the subject’s mental state.

Third, some disjunctivist accounts focus on object perception (seeing a tree) rather than fact perception (seeing that it is a tree). Object perception is less “epistemic” than fact perception. A cat can look at a king without knowing that he is a king. Here the relation to knowledge first epistemology is less direct. Nevertheless, the emphasis on the explanatory priority of successful cognitive engagement with the external world remains similar.

The extensive analogies between perception and memory raise the question of disjunctivism about memory. So far, such a view has been less fully explored.

Further Themes

Knowledge first epistemology has many more applications. One issue is the nature of consciousness. While opponents often regard knowledge first epistemology as false to what conscious experience is like, as shown by arguments from illusion, proponents reply that many states of consciousness do *seem* essentially to involve external objects and facts; thus when those states are as they seem, they *do* essentially involve external objects and facts, just as states of conscious knowledge and perception do. In cases of illusion, even our states of consciousness are not always as they seem. That is just what the anti-luminosity argument predicts: conscious subjects are sometimes not in a position to know what states of consciousness they are in.

Another application of knowledge first epistemology is to questions about what determines the content of mental attitudes: for instance, what constitutes mental reference to one object rather than another? A central role is often assigned to a principle of charity that prefers assignments of content that maximize the subject's true beliefs. However, assigning true beliefs to a subject might not be an interpretive virtue when those beliefs do not constitute knowledge. This suggests using instead a principle of charity that maximizes the subject's knowledge rather than true belief.

Knowledge first epistemology is still an alien way of thinking to many philosophers, despite its roots in naive realism and common sense. Much of its detailed working out has been very recent, and all the issues raised in this chapter require extensive further investigation. Knowledge first epistemology is contributing a further stage to the development of externalist ideas over past decades. The process is likely to continue.

Further Reading

- Austin, J.L. (1964) *Sense and Sensibilia*, Oxford: Oxford University Press. (Classic defense of common sense about perception and bridge between earlier and later versions of knowledge first epistemology.)
- Gettier, E. (1963) "Is Justified True Belief Knowledge?" *Analysis* 23: 121–3. (The classic article.)
- Gibbons, J. (2001) "Knowledge in Action," *Philosophy and Phenomenological Research* 62: 579–600. (Defense of knowledge constraint on reasons for action.)
- Greenough, P. and Pritchard, D. (eds) (2009) *Williamson on Knowledge*, Oxford: Oxford University Press. (Extensive collection of critical essays on *Knowledge and its Limits*, with replies.)
- Haddock, A. and Macpherson, F. (eds) (2008) *Disjunctivism: Perception, Action, Knowledge*, Oxford: Oxford University Press. (Large collection of papers for and against.)
- Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford: Clarendon Press. (Examination of puzzles about chance and shifting standards, using knowledge first ideas.)
- and Stanley, J. (2008) "Knowledge and Action," *Journal of Philosophy* 105: 571–90. (Defense of knowledge norm.)
- Hinton, J.M. (1973) *Experiences*, Oxford: Clarendon Press. (Early version of disjunctivism about perception.)
- Martin, M.G.F. (2002) "The Transparency of Experience," *Mind and Language* 17: 376–425. (Disjunctivism about experience.)
- (2004) "The Limits of Self-Awareness," *Philosophical Studies* 120: 37–89. (Disjunctivism about perception.)
- McDowell, J. (1998). *Meaning, Knowledge and Reality*, Cambridge, Mass.: Harvard University Press. (Collection of his papers, with many themes related to knowledge first epistemology, including disjunctivism.)
- Phillips Griffiths, A. (ed.) (1967) *Knowledge and Belief*, Oxford: Oxford University Press. (Anthology with pieces by John Cook Wilson and H.A. Pritchard advocating an earlier form of knowledge first epistemology.)
- Snowdon, P. (1980–1) "Perception, Vision, and Causation," *Proceedings of the Aristotelian Society* sup. vol. 81: 175–92. (Disjunctivism about perception.)

THE ANALYSIS OF KNOWLEDGE

- Unger, P. (1975) *Ignorance: A Case for Scepticism*, Oxford: Oxford University Press. (Early source of linguistic evidence for knowledge first ideas.)
- Williamson, T. (2000) *Knowledge and its Limits*, Oxford: Oxford University Press. (Development of a knowledge first epistemology along the lines presented here.)
- (2007) *The Philosophy of Philosophy*, Oxford: Blackwell. (A knowledge first epistemology of philosophy, with discussion of knowledge maximization.)