Knowledge

Bryan Frances
Allan Hazlett

Forthcoming in Howard Robinson and Barry Dainton (eds.), *A Companion to Analytic Philosophy*, Continuum Press.

This chapter recommends four topics in the theory of knowledge: skepticism, knowledge attributions, the nature of knowledge, and the value of knowledge. These topics (among others) have taken center stage in contemporary analytic epistemology.

SKEPTICISM

You can't get far on a theory of truth unless you tackle the liar paradox. The reason is that that paradox seriously threatens the truth of some of the most fundamental assumptions one is naturally inclined to make regarding truth and meaning. For analogous reasons, one can't get far in epistemology unless one deals with the skeptical paradox.

Skeptical arguments come in various forms, but in recent philosophy almost all of them have this one:

1. In order to know P, you have to meet condition C.
2. You don't meet condition C.
3. So, you don't know P.

Here are just two candidates for C:

> C1: You know, or are in a position to know, or justifiably believe, that it's not the case that almost all your ordinary beliefs aren't true because you're a brain in a vat being fed electrical signals so as to generate the mere illusion of having a normal body and living a normal life.
>
> C2: You know, or are disposed to know, or justifiably believe, that it's not the case that almost all your ordinary beliefs aren't true because there are no trees, shoes, or other ordinary objects, as the paradoxes about material objects prove (there are no shoes, although there are "particles arranged shoe-wise", and your beliefs are about the former instead of the latter).

P is chosen to be a substantive, contingent, empirical claim almost all of us think is pretty well known, such as 'I own three pairs of shoes'. The reason these arguments merit close attention – and are considered paradoxical – is that, with P and C specified in this way:

- There are good reasons for (1).
- There are good reasons for (2).
- There are good reasons against (3).
- There are good reasons for the validity of the argument (1)-(3), so there are good reasons to think that it's impossible for (1) and (2) to be true while (3) is false.

We tend to think that perception and elementary reasoning are pretty good at generating evidence or justification fully adequate for knowledge (when combined with true belief). We do seem to know at least *some* substantive, contingent, empirical claims, by using perception and ordinary reasoning, like 'I own three pairs of shoes'. So, we are highly confident in the falsehood of (3).

The reason (1) is so plausible is that it is so easy to *reliably reason* from 'I own three pairs of shoes' to 'It is not the case that C holds', for either C1 or C2. After all, if you know (or justifiably believe) that you own three pairs of *shoes*, then it's the easiest thing in the world to see that it simply has to be the case that *there are shoes*, which is obviously inconsistent with both C1 and C2. Often this is put in terms of the *Closure Principle*: if you know P (e.g. that you own three pairs of shoes) and you know that P entails Q (e.g. that your owning three pairs of shoes entails that there are shoes), then you know (or are in a position to know) Q (e.g. that there are shoes). If we can reliably reason from 'There are black holes in the Milky Way galaxy' to 'There are black holes', then surely we can reliably reason from 'I own three pairs of shoes' to 'There are shoes'. Reliably reasoning from 'There are shoes' to the falsity of C1 and C2 looks just as elementary.

The reason (2) is so plausible is that it seems like we don't satisfy C1 or C2. Take C1 first. What could you possibly *do* to show that you're not a brain in a vat (BIV)—with or without three pairs of shoes? Well, someone might say, mere brains don't have feet, and you can clearly *see* your feet. But that so-called "seeing" might be an illusion manufactured by the mad scientists who are controlling your brain. You might say that it's impossible for scientists or anyone else to create this illusion. But that "fact" about science might be an illusory bit of "information" fed to you by just those scientists. It doesn't take a great deal of reflection before you see the plausibility of the idea that you can't do anything to show that you're not a BIV—at least not in any ordinary way of doing something to prove a point. In addition, in all probability there is little you can do to show that the paradoxes about material composition fail to show that there are no shoes in our universe. However, C1 and C2 seem a bit different on this score. Whereas it seems that it was *impossible* for you to do anything to rule out the BIV possibility, we have the feeling that it's possible for you to do something to rule out the no-shoes possibility indicated in C2, as you could "just" give a knockdown argument that refutes the philosophers and physicists who say that there are no composite objects, including shoes. But if you don't have the ability to do anything to rule out those metaphysical theories, then it looks like you can't do anything to rule out the no-shoes possibility mentioned in C2.

So is there some other way you can know that you're not a BIV and that the paradoxes of material composition don't show that there are no shoes? You might appeal to the commonsensical point that knowing something and proving it are two different things – that someone can know something even if she cannot prove it. Sometimes we have the ability to tell that something is the case, but not the ability to prove that it's the case – think of how you have the ability to recognize people by their faces, but would be hard pressed to prove that your identifications were correct. However, this won't help with C1 and C2. Perhaps an elite professional metaphysician has the resources to tackle the paradoxes of material composition, and so has the ability to tell whether they do or do not show that there are no shoes, but most of us don't have that ability. And the situation is even worse when it comes to C1, for it doesn't seem possible that *anybody* could have the ability to tell whether she was a brain in a vat. So (2) certainly seems true.

The argument (1)-(3) certainly appears to fit *modus tollens* perfectly, and since that's a deductively valid form of reasoning, it seems impossible for (1) and (2) to be true while (3) is false.

Hence, we have a paradox: a collection of four claims that are individually plausible yet collectively inconsistent. Here are the four claims:

- *(1) is true*
- *(2) is true*
- *(3) is false*
- *It's impossible for (1) and (2) to be true while (3) is false*.

So there must have been some mistake in our reasoning. The reasoning offered above in favor of those four claims is profoundly misleading in some way that is hard to detect. In fact, it is so hard to detect that despite centuries of dedicated investigation epistemologists have come to no agreement regarding which of the four claims is false. Our task is to discover which claim is false, and to explain why.

Since there are four claims in our inconsistent set, the four ways of avoiding the inconsistency involve denying one of the four individual claims. That is, we have *just four* options in responding to the paradox: hold that (1) is false, hold that (2) is false, hold that (3) is true, or hold that the argument isn't valid. One can hardly claim to have a thorough response to the paradox if one doesn't indicate which option one is adopting. *Skepticism* is the view that (3) is true. Let's consider the other possible thorough responses.

The *Neo-Moorean* says that (2) is false. The position gets inspiration from G. E. Moore, who would have said that (2) is false. On this view, we do know that we're not BIVs. Different philosophers give different reasons for rejecting (2). Moore (1939) would argue that we can know we're not BIVs by deducing that conclusion from some ordinary item of knowledge, such as ‘I own three pairs of shoes’; versions of this idea have been defended by Ernest Sosa (1991, 1999b, 1999b) and James Pryor (2000). For *reliabilist* Neo-Mooreans (like Sosa), the mere fact that my belief-forming faculties are reliable – the mere fact that they usually yield mostly true beliefs – is enough to ensure that most of my beliefs amount to knowledge, including my belief (should I come to form it) that I’m not a BIV. On this view, we do not need to do anything to show that our faculties are reliable, to ensure that most of our beliefs amount to knowledge; for reliabilists it is sufficient that our faculties *are* reliable. For criticism of this approach, see Stroud 1994 and Fumerton 1995. Alternatively, Jonathan Vogel (1990) argues that we can know we aren’t BIVs via an inference to the best explanation; Berent Enç (1990) offers criticism of that view.

The *closure denier* says that (1) is false. For example, Robert Nozick (1981) argues that we can know that we have shoes even though we don’t and can’t know that we aren’t shoeless BIVs. In so doing, these philosophers reject the Closure Principle (mentioned above): if you know P and know that P entails Q, then you know (or are in a position to know) Q. On the one hand, this is an attractive position: we get to say, with common sense, that *of course* we know that we own three pairs of shoes, but we also acknowledge the skeptic’s achievement by admitting that we can’t know that we aren’t shoeless BIVs. But then again, isn’t it incredibly easy to deduce, from your knowledge that you have three pairs of shoes that fit your feet quite well, that you have shoes, and so come to know that have you shoes? What inference could be easier? And isn’t it equally easy to then deduce that you’re not a shoeless BIV, which by definition has no shoes? Such considerations make the rejection of (1) and the Closure Principle look awkward at best. The places to start studying: Dretske 1970, Nozick 1981, Kvanvig 2006, and Kripke 2011.

The last way of dealing with the skeptical paradox, rejecting the validity of the (1)-(3) argument, introduces the topic of the next section: epistemological attributions. By studying how we *talk about* knowledge we may not only get a new way to respond to the skeptical argument but a better understanding of what knowledge really is.

But before we get to that issue it is worth mentioning that epistemologists have recently been studying another kind of skeptical problem. The argument goes something like this: you start out believing P, you then acquire some information that gives you powerful reason to think that your belief is false or your basis for your belief is seriously flawed; you don’t have any evidence that can “counter” that worrisome information; so, you don’t know P (even if you started out knowing P). For instance, P might be ‘We have free will’, ‘God doesn’t exist’, or ‘Truth is a consistent notion’, but you know that there are loads of people smarter and much better informed than you are who think your belief is false. Alternatively, the worrisome information you acquire might be the claim that most people’s opinions on P are typically greatly exaggerated, and there is no reason to think you are an exception. Or P might be some political opinion of yours, and the problem is that people almost always reach their political opinions through non-evidential means, and there is no reason to think you are an exception. On this kind of skeptical problem, see Feldman 2005, Frances 2010 and forthcoming, Kornblith 2010, and White 2010.

KNOWLEDGE ATTRIBUTIONS

A story condensed from DeRose 2009 is an excellent starting point into the literature on knowledge attributions:

> Thelma, Louise, and Lena are friends who all work in the same office. Today is their day off, but, before getting an early dinner together, they decide to walk up to the office to pick up their paychecks. Thelma and Lena are also interested in finding out whether their colleague John is at work, as they are involved in a small office bet with a couple of other workers over whether John would show up today. As they pass the door to John's personal office, they see his hat hanging on the hook in hallway, which, in their long experience, has been a sure-fire sign that John is in fact at work. Satisfied that John is at work and that Thelma and Lena, who bet that he would be, are in a position to collect their winnings from some other office workers, the three friends pick up their checks, go out to dinner together, and then part company, Thelma going to a local tavern to meet other friends, and Louise and Lena each heading in different directions to go home.
>
> Thelma meets a friend at the tavern and says to him, 'Hey, John was at work today. Pay up!' When her friend asks, 'How do you know?' Thelma describes the above facts. Then Thelma says, 'Lena knows that John was there, too, as she was with me.'
>
> At the same time Louise is stopped by the police on her way home. They are conducting an extremely important investigation and are seeking to determine whether John was at work that day. When the police ask her whether she could testify that John was at work, Louise replies, 'Well, no, I never saw him. I could testify that I saw his hat hanging in the hall, which is a very reliable sign that he's at work. But I suppose John could have left his hat on the hook when he went home some previous day.' When the police ask Louise whether Lena might know whether John was in, Louise replies, 'No. She has the same reasons I have for thinking John was there, but, like me, Lena doesn't know that John was there.'

It seems that when Thelma says at the tavern 'Lena knows that John was there', her claim is true. But it also seems that when Louise tells the police 'Lena doesn't know that John was there', her claim is true. But that looks impossible, as they seem to be asserting contradictory claims. Theorists of knowledge attributions have the task of explaining why both Thelma's 'Lena knows that John was at work' and Louise's 'Lena doesn't know that John was at work' are both conversationally appropriate even though they certainly appear to be inconsistent with one another.

Contextualists adopt the most straightforward explanation: the reason the two assertions are conversationally appropriate is that they are both *true* (Cohen 1988, 1999, DeRose 1995, 2009, Lewis 1996). And that means that for at least some sentences of the form 'S knows P', there are two tokens of the sentence that concern the same S and P (and time) where one token is true and another is false. This can happen, according to contextualism, when the two tokens are uttered in different contexts. So

'S knows P' is akin to 'S is tall', which can have two tokens with different truth-values depending on the standards that are operative in two different contexts (e.g., a single person can, at one time, be tall for a high school student but not tall for a basketball player).

Contextualists need an explanation of *how* tokens of 'S knows p' get different truth-values, and this explanation will have to pay careful attention to semantic issues; this is one reason why the philosophy of language has become important to epistemology. One idea (mentioned above) is that in different conversational contexts there are different epistemic *standards* a belief has to meet in order to be called 'knowledge': in Louise's police context the standard is high, in Thelma's tavern context it is low, and Lena's belief meets the low but not the high standard. Again, compare this to how 'Chris is tall' can be true relative to one context (in which we are discussing high school students in general) but false relative to another (in which we are discussing only the students on the high school basketball team). But there are alternative ways of understanding the context-sensitivity of 'knows', including by appeal to salient alternatives (Lewis 1996) and salient contrast (Schaffer 2005, 2007).

According to the contextualist, whether or not *my* use of 'S knows P' is true will depend not only on S (e.g., whether she believes p, what evidence she has, whether her belief was formed reliably) and on P (whether P is true), but will depend on the context surrounding *me*. It isn't hard to see how these ideas might apply to the skeptical paradox. When we are strongly inclined to reject (3)—that is, to reject the sentence 'So, you don't know P'—contextualists say we are in a low-standards, "everyday" context in which (3) is false (or that we are imagining ourselves in such a context). But when we are attracted to the skeptic's premise (2) – 'You don't meet condition C' – contextualists say we are in a high-standards, "skeptical" context in which (2) is true (or we are imagining ourselves in such a context). (Alternatively, (2) might be without truth-value in that context because of some kind of indeterminacy; see DeRose 2004). In the low-standards context I speak the truth when I utter the sentence 'I know I have three pairs of shoes', even though this sentence would be false (or indeterminate) were I to utter it in a high-standards context. The skeptical conclusion, in a high-standards context, is compatible with the truth, in a low-standards context, of our ordinary claims to knowledge. For criticism of this response to the skeptical paradox, see Sosa 2000.

Contextualism is a theory about the meaning and use of 'knows', and it has several rivals in the theory of knowledge attributions. According to *subject-sensitive invariantism* (Hawthorne 2004, Stanley 2005, Fantl and McGrath 2009), whether a use of 'S knows P' is true depends not only on such "epistemic" factors as S's evidence or the reliability of S's methods of belief-formation but also on S's *practical circumstances*. (Subject-sensitive invariantists thus endorse "pragmatic encroachment," which will be discussed in the next section.) Contextualism says that the details of the *attributor* of 'S knows P' make a difference as to its truth conditions; subject sensitive invariantism says that the practical circumstances of the *subject* S make a difference as to whether S knows. However, a contextualist could adopt (something like) both theses by saying that the truth-value of 'S knows P' depends on both subject and attributor factors (Lewis 1996, DeRose 2009). Alternatively, *relativism* says that the truth-value of 'S knows P' depends on the person *evaluating* the sentence (MacFarlane 2005). Finally, *Griceans* (Rysiew 2001, Hazlett 2009) deny that the truth condition for 'S knows P' vary with context, and account for the difference in conversational appropriateness by appeal to Grice's theory of conversational implicature. For instance, perhaps when Louise says to the police that Lena doesn't know that John was at work she speaks falsely—Lena really does know—but correctly conveys the information that the police will not be able to rely on Lena's testimony in court. Louise is trying to be informative for the police: although they asked whether Lena knows, what they are really after is whether Lena *saw John at work*, and were Louise to say that Lena knows, it would imply something false: that Lena saw John at work.

The task for the theory of knowledge attributions is to offer a detailed response to the various cases in a convincing way. In particular, one needs a well developed response to as many of the cases as possible (for a good portion of them, see Hawthorne 2004, Stanley 2005, and Rysiew 2011).

THE NATURE OF KNOWLEDGE

The skeptical problem concerns the *scope* of our knowledge: it concerns *what* we can know, or *how much* we can know. A distinct, though related, problem concerns the *nature* of knowledge; this problem concerns what knowledge *is*. (Note that this is also distinct from, though related to, a question discussed above, about the meaning and use of the word 'knows'.) Analytic philosophers have explored the nature of knowledge primarily by attempting to articulate necessary and sufficient conditions for knowledge. A minority look elsewhere, to the genealogy of the concept of knowledge (Craig 1991, Kusch 2009, Dogramaci forthcoming), for illumination.

According to the *tripartite analysis of knowledge*, first articulated by Socrates in the *Meno*,

> S knows that p iff (i) p is true, (ii) S believes that p, and (iii) S is justified in believing p.

We'll comment on each of the three conditions. Knowledge, according to condition (i), is *factive*: being in such a factive state with respect of some proposition P entails that P is true. (Note well that the view that knowledge is factive does not entail that the truth is transparent to us, for it is compatible with this view that we rarely, if ever, know whether we know.) Condition (i) is sometimes defended by appeal to a linguistic thesis about 'knows' (cf. Hazlett 2010). We should keep in mind that cultural or historical relativists about truth can embrace the necessity of condition (i): that people once knew that the earth was flat is no counterexample to the necessity of the truth condition, if at the time they knew that proposition, it was true.

According to condition (ii), knowledge is a species of belief. This view can be challenged by appeal to the following kind of case (Radford 1966, Williams 1973): an intellectually cautious student correctly answers a set of exam questions, whilst suspending judgment on the contents of her answers; the student knows the answers to the exam questions, but she does not believe those answers. This suggests an alternative conception of knowledge, on which to know is to possess (without necessarily believing) true information (Williams 1973), or perhaps to have the ability to access true information in action (Ryle 1949). The issues at play here are as much issues about the nature of belief as about the nature of knowledge; compare a conception of belief on which it is that propositional attitude that combines with desire to produce action, on which the student in the above example does believe her answers, with a conception of belief on which it entails conscious, affirmative judgment, and on which the student does not believe her answers. It is agreed that knowledge requires the possession, in some sense, of true information. The tripartite analysis understands "possession" as belief; others (Radford and Williams) think this requires too much; still others think this requires too little, since knowledge requires certainty (Ayer 1956, Unger 1975).

Condition (iii) is controversial in two ways. First, some argue that true belief is sufficient for knowledge (Sartwell 1991; cf. Hazlett 2010). A jealous husband knows about his wife's affair when he truly believes that she is having an affair, regardless of whether his belief is justified (Williams 1973). Alternatively, you might think that 'knows' is polysemous such that on one sense of "knows" it is synonymous with "true belief". Second, some theories of knowledge, while not endorsing the thesis that true belief is sufficient for knowledge, do not require justification for knowledge (Goldman 1967, Nozick 1981).

The fate of condition (iii) depends on the nature of justification. Some suggest a *dialectical conception of justification*, on which being justified in believing P is a matter of being able to give reasons for P, or of being able to justify one's belief that P to others, or of being able to avoid criticism for believing P (Ayer 1956, Wittgenstein 1969). Two sorts of considerations speak against appealing to this conception of justification in our analysis of knowledge. First, we may have occasion to attribute knowledge to creatures who fall outside the dialectical "space of reasons," e.g. young children and other non-linguistic animals. Second, we may find ourselves unable to justify our ordinary beliefs against the skeptic's challenge, thus anti-skeptical epistemology will require a non-dialectical conception of justification (Sosa 1991).

We might then opt for an *evidentialist conception of justification*, on which being justified in believing P is a matter of having sufficient evidence for P. According to a prominent evidentialist account of knowledge, knowledge is true belief that is well-founded, where a belief in P is *well-founded* iff it is properly based on evidence that justifies P (Conee and Feldman 2004).

Reliabilists (see above) reject this account of knowledge, either in favor of the view that justification is not necessary for knowledge (see above) or in favor of a *reliabilist conception of justification* (Goldman 1979, Plantinga 1993, Sosa 1991), on which being justified in believing P is a matter of one's belief in P being formed (and sustained) in a reliable way. This leaves open the nature of reliability; one possibility is a modal account (Nozick 1981, Sosa 2001) that defines reliability in terms of sensitivity (S's belief that p is sensitive iff were it the case that ~P, then S wouldn't believe P) or safety (S's belief that P is safe iff S would not easily believe P when P is false).

At the heart of the dispute between evidentialists and reliabilists is a disagreement about whether it is possible for someone to know P without having reasons for believing P. There is an important sense of "having reasons" on which reliabilists deny that knowing requires having reasons. Epistemologists have discussed cases in which someone's belief in p is reliably formed and sustained, but where (in the aforementioned sense) she does not have reasons for believing P, e.g. a savant endowed with a highly reliable faculty of mathematical intuition, whose mathematical beliefs are by hypothesis not based on evidence. The beliefs of the savant are reliably formed and sustained, no matter how you understand these notions, but not well-founded. Shall we say that she knows? If yes, then our sympathies lie with some form of reliabilism. If no, then with some view more akin to evidentialism.

The tripartite analysis of knowledge is rejected by most contemporary epistemologists, on the basis of the following kind of counterexample (Gettier 1963):

> Smith and Jones are the only applicants for a certain job. Smith has strong evidence (testimony from the president of the company they have applied to) that Jones will get the job, and also (perceptual evidence) that Jones has ten coins in his pocket. From these two propositions Smith deduces that the man who will get the job has ten coins in his pocket. However, it turns out both that Smith will get the job (the president was mistaken) and that Smith happens to have ten coins his pocket (although he doesn't know this).

Smith's belief (that the man who will get the job has ten coins in his pocket) is true and justified; however, intuitively, Smith does not know that the man who will get the job has ten coins in his pocket. The tripartite analysis is therefore false. Coming up with a satisfactory account of the nature of knowledge that is not subject to this kind of counterexample is known as the *Gettier problem*.

Why does Smith fail to know? You might think that he fails to know because his belief is deduced from a falsehood; this thought forms the basis of the "no false lemmas" approach to the Gettier problem, on which S knows that p only if S's belief that p is not deduced from any falsehoods, but many have thought that this diagnosis fails to capture the essence of Gettier's case. Consider another case (Goldman 1976):

> Henry is unknowingly driving though "fake barn country," where most of structures in the fields are barn facades, which from a distance look like real barns; the rest of the apparent barns are real. Driving past one of the rare real barns, Henry looks out the window, sees the barn in the distance, and forms the belief that there is a barn in the field.

Henry's belief (that there is a barn in the field) is true and justified; however, intuitively, Henry does not know that there is a barn in the field. The tripartite analysis is therefore false, and in this case Henry does not deduce his belief from any falsehood: his belief is simply based on (or is a constituent of) his perception of the barn in the field.

The case of Henry suggests that whether someone knows depends crucially on her *environment*. Had Henry formed the belief that there's a barn in the field while driving through a more normal environment, he would have known that there was a barn in the field. And both the case of Smith and the case of Henry suggest that knowledge is incompatible with certain kinds of *luck* (Nagel 1979, Pritchard 2005). The problem with the tripartite analysis is that it fails to take any of these facts into account.

Reliabilist virtue epistemologists (Sosa 2004, Greco 2010) have articulated a solution to the Gettier problem based on the idea that knowledge is a cognitive achievement. In general, we can distinguish between (mere) success and achievement: an archer is successful when her arrow hits her target, but her success is an achievement when her arrow hits her target as a result of her archery skills, if her successful shot manifests her competence at archery. We are asked to think of cognition as analogous to such an activity, and as admitting of the same distinction: my (mere) true belief is a success, but my knowledge is an achievement – a success that manifests my competence at forming true beliefs. Given this conception of knowledge, we can see why Henry's true belief falls short of knowledge: it is not an achievement of Henry's. (The same, mutatis mutandis, when it comes to Smith.) His belief is analogous to a competent archer's shot which reaches its target by luck, e.g. by being blown off course, and then back on course, by two chance gusts of wind. Such lucky success is to be distinguished from genuine achievement, where success is *due to* or a *manifestation of* the abilities of the successful agent.

Timothy Williamson (2001) proposes a more radical approach to the Gettier problem: a rejection of the project of analyzing knowledge in terms of belief and evidence, in favor of *"knowledge first" epistemology*, on which notions like belief and evidence are to be analyized in terms of knowledge.

The Gettier problem suggests that whether someone knows depends crucially on her environment; other cases suggest that whether someone knows depends crucially on her *practical circumstances*, including the decisions she has to make and her preferences and concerns:

> Train 1: John is waiting on a train platform as a train pulls up and stops. He has taken this particular train about twice a month for a couple years. Someone standing next to him asks him 'Do you know if this train stops at Melrose?' John has never taken the train to Melrose but he is certain he has heard the train personnel say 'Melrose is the next stop' a great many times while riding this train; his memory is ordinary. So, he says 'Yes, it stops at Melrose'.
>
> Train 2: everything is the same as in case 1 except that the people asking John are police officers with 'Bomb Squad' on their uniforms who have run up on the waiting platform and announced that it's imperative that they get to the Melrose station as soon as possible. They turn to John and ask him 'Do you know if this train stop at Melrose?' John understands the stakes at hand. He says 'I am pretty sure it does but you'd better ask the conductor'.

A natural explanation of why it's fine, in Train 1, for John to say that the train stops as Melrose is that he knows that the train stops at Melrose, but an equally naturally explanation of why it would be wrong, in Train 2, for John to say that the train stops at Melrose, is that he *doesn't know* that the train stops at Melrose. If he *knew*, in Train 2, why wouldn't be able to appropriately tell the police that the train stops there?

According to the thesis of *pragmatic encroachment* (Hawthorne 2004, Stanley 2005, Fantl and McGrath 2009), whether someone knows P depends on her practical circumstances: in Train 2, more evidence is needed for John to know that the train stops at Melrose, given that something serious seems to ride on the reliability of his testimony. Defenders of pragmatic encroachment have understood their view as a radical departure from epistemological orthodoxy, although this picture has been challenged (Grimm

2011). Alternatively, you might favor a Gricean account of the train cases (Hazlett 2009), on which John needs more than knowledge to answer the police affirmatively in Train 2.

Inquiry into the nature of knowledge has received criticism recently on the ground that it is illegitimately based on appeals to philosophical intuition, and in particular that it is illegitimately based on appeals to philosophical intuitions about the extension of the concept of knowledge. "Experimental philosophers" have argued that people's intuitions about Gettier cases vary with presumably irrelevant factors such as ethnic background and gender (Weinberg, Nichols, and Stich 2001, Buckwalter and Stich 2011), but these experimental results could not be replicated (Nagel forthcoming). Alternatively, the legitimacy of appeals to intuition can be problematized a priori by asking whether knowledge is a natural kind (Kornblith 2006). For if knowledge is not a natural kind, we might then question the philosophical, as opposed to anthropological or sociological, interest in limning the extension of some socially constructed concept of ordinary language.

An alternative to the method based on appeals to philosophical intuitions about the extension of the concept of knowledge is one based on appeals to evaluative premises, where the value of knowledge is a constraint on theories of the nature of knowledge (Kvanvig 2003). This project promises more than a limning of the extension of a socially constructed concept of ordinary language – it promises an account, at least, of something that we care about (in as much as there is a connection between what we care about and what is valuable). And when we return to the "traditional" methods employed in the theory of knowledge, we find the evaluative method has been with us all along: the tripartite analysis is articulated by Socrates in connection with the question of how knowledge is more valuable than (mere) true belief (see below), and regardless of whether the cognitive achievement conception of knowledge is well-motivated by intuitions about Gettier cases, which are the sort of intuitions to which we have just considered challenges, it seems well motivated by the idea that knowledge is more valuable than (mere) true belief (again, see below).

THE VALUE OF KNOWLEDGE

Many epistemologists maintain that knowledge is valuable. This claim, however, admits of numerous disambiguations.

First, we can distinguish between the claim that knowledge is *always* valuable and the claim that knowledge is *sometimes* valuable. Second, consider the claim that knowledge *per se* is valuable, i.e. that knowledge is valuable "as such": being F is per se valuable when anything that is F is valuable just in virtue of being F. (Therefore, that knowledge per se is valuable entails that knowledge is always valuable.) Third, consider the claim that knowledge is *prima facie* valuable: X is prima facie valuable when X is valuable "at first glance" or "on its face." Fourth, consider the claim that knowledge is *pro tanto* valuable: X is pro tanto valuable when X has a value that can be trumped by other values. Fourth, consider the claim that knowledge is *intrinsically* valuable: X is intrinsically valuable iff the fact that X is valuable supervenes on the intrinsic properties of X. (Intrinsic value contrasts with *extrinsic* value.) Fifth, consider the claim that knowledge is *finally* valuable: X is finally valuable iff X is valuable for its own sake. Final value contrasts with *instrumental* value: X is instrumentally valuable iff X is valuable for the sake of something else (wholly distinct from X). (Intrinsic value is therefore a species of final value, and instrumental value a species of extrinsic value.) Sixth, you might think that knowledge has (what we can call) *eudaimonic* value: X is eudaimonically valuable when X is valuable vis-à-vis wellbeing. So, for example, you might think that knowledge is *good for* the knower; if you think this, then you think that knowledge has eudaimonic value. Seventh, you might argue that knowledge has *constitutive* value: X has constitutive value iff X is valuable in virtue of being part of a valuable whole, such as a good human life. (Constitutive value could naturally be understood as a species of final eudaimonic value, to be contrasted with instrumental eudaimonic value, which could be called "practical" or "prudential" value.) Eighth, you might want to say that knowledge has *epistemic* value, although the notion of "epistemic value" is obscure, as is its relationship to eudaimonic and constitutive value.

In contemporary epistemology, the most frequently discussed question concerning the value of knowledge is the *Meno question*: why, and in what way, is knowledge more valuable than true belief (if it is)? In the *Meno*, Socrates presents an answer based on the idea that knowledge is more secure than (mere) true belief (97e-98a). On its face, this suggests that the instrumental eudaimonic value of knowledge is always greater than the instrumental eudaimonic value of (mere) true belief, in virtue of the fact that knowledge is resistant to revision in a way that (mere) true belief isn't. But this seems wrong. Imagine that a credulous nitwit is absolutely convinced of some proposition, on the basis of the testimony of his obviously unreliable guru. The proposition, as it turns out, is true. Suppose that a cautious thinker is reasonably confident of the truth of the same proposition, on the basis of careful consideration of the evidence that bears on the proposition. The (mere) true belief of the credulous nitwit might be much more secure, much more resistant to revision, than the knowledge of the cautious thinker: nothing will make the nitwit give up her conviction, while the cautious thinker will be moved to revise if evidence against the proposition emerges.

Contemporary epistemologists (Kvanvig 2003, Pritchard 2010) have articulated the Meno question in terms of epistemic value: why, and in what way, is knowledge more *epistemically* valuable than (mere) true belief (if it is)? As Duncan Pritchard (2010) argues, the value problem arises when we assume:

> *Truth value monism***:** True belief (and avoiding false belief), and only true belief (and avoiding false belief), has fundamental epistemic value.

Truth value monism does not entail that only true belief has epistemic value. It allows for the possibility of other epistemic goods, but only when the epistemic value of said goods can be explained in terms of the epistemic value of true belief. Thus the truth value monist might explain the epistemic goodness of the virtue of intellectual humility by appeal to the fact that (i) the intellectual humble person will be a more reliable believer, and (ii) a more reliable believer is one more likely to acquire true beliefs. But truth value monism seems to imply that any true belief, whether amounting to knowledge or not, will be doing maximally well in terms of the fundamental epistemic good. Compare two beliefs of mine, one an instance of knowledge, the other not an instance of knowledge, but merely a true belief. Both are equally good in terms of the only fundamental epistemic good, according to truth value monism. My knowledge differs from my (mere) true belief, though: perhaps the former is based on good evidence, while the latter is wishful thinking. But if truth value monism is right, the epistemic value of basing your beliefs on good evidence must be non-fundamental; we must say that basing your beliefs on good evidence is epistemically valuable because this makes it more likely that your beliefs will be true (or something like that). But now we cannot make sense of how my knowledge is more valuable than my (mere) true belief, in virtue of the former being based on good evidence: for my (mere) true belief already enjoys the epistemic good that basing one's beliefs on good evidence is said to bring. Non-fundamental epistemic value is "swamped" by fundamental epistemic value: when the latter is present, the former loses its appeal.

Compare a gold miner, who (like the Bond villain) loves only gold (cf. Zagzebski 1996). Gold, for him, is the only thing that has fundamental value. Much, of course, has non-fundamental value: pickaxes and dynamite, for example, which are reliable tools for mining for gold. Suppose now we ask the miner what he prefers: an ounce of gold, brought out of the ground with pickaxe and dynamite, or an ounce of gold, won by blind chance at the casino? If gold really is the only thing that has fundamental value for him, then he must be indifferent between these two treasures. Both are equivalent in terms of fundamental value; they only differ in non-fundamental value. If the miner prefers one of these prizes over the other, it is because he loves something other than gold.

Defenders of the cognitive achievement conception of knowledge (Sosa 2004, Greco 2010) have offered an answer to the Meno question. They argue that achievement is more valuable than (mere) success, and from this, along with the view that knowledge is an achievement and true belief a (mere) success, it follows that knowledge is more valuable than (mere) true belief. It is obscure whether the value of knowledge is epistemic value, on this view, but that is because the notion of "epistemic value" is obscure. Perhaps we should reject truth value monism: like a gold miner who values gold obtained through a hard day in the mine over gold won at the casino, perhaps we value truth acquired through intellectual virtue to truth acquired by luck. This will appeal to those already antipathetic to truth value monism, e.g. those who urge us to consider the epistemic value of understanding and wisdom, or those who urge us to consider not only the epistemic value of states such as knowledge and true belief, but also that of intellectual virtues such as openmindedness, consistency, and intellectual integrity (cf. Zagzebski 1996).

A final possibility would be to reject the presupposition of the Meno question, and deny that knowledge is more valuable than (mere) true belief. If a solution to the value problem eludes us, we must return to this presupposition and critically examine it. Whence the intuition that knowledge is more valuable than (mere) true belief? It is sometimes suggested (Kvanvig 2003) that a defense of the distinctive value of knowledge is needed to justify epistemologists' concern with knowledge. However, this argument is unsound. What might be relevant to justify epistemological inquiry would be the value of *knowing* about knowledge, not the value of knowledge. For there could easily be valuable knowledge about a worthless object: it is worth knowing about the molecular nature of rhodium, even if rhodium is worthless. (Moreover, it seems to us, no defense of the value of knowing about crustacean sexual behavior is needed to vindicate biological curiosity about it. Curiosity itself is vindication enough, we might say, although we do better to remember a liberal principle: no one needs to vindicate their preferred harmless (intellectual) activities, by showing that they are valuable.) The premise that epistemologists are wrong to study knowledge, unless knowledge is distinctively valuable, is false.

The Meno question concerns the value of knowledge, as opposed to true belief; a more basic question concerns the value of knowledge, as opposed to ignorance, in the form of a lack of belief, or in the form of false belief. Some philosophers defend the constitutive eudaimonic value of knowledge (Zagzebski 2003, Sosa 2003, Lynch 2004, Greco 2010, Baril 2010), others the instrumental eudaimonic value of knowledge (Kornblith 1993). It seems like some knowledge is worthless: I open the Wichita phonebook to a random page and memorize a name and phone number – an instance of knowledge, if ever there was one, but it is not plausible that this knowledge contributes to my wellbeing, either in an instrumental or a constitutive way. One might reply that such knowledge does have constitutive value, but that such value is trumped by something else, e.g. the time and effort it takes to memorize the number (Lynch 2004). But things do not *seem* that way: it is not like a case in which I must choose between pastrami and corned beef; even when I settle on pastrami, the appeal of the corned beef remains. In the case of the phonebook, knowledge seems to have *no* appeal.

Which knowledge has constitutive value? We might appeal to those truths that are objectively important or significant (Baril 2010), or to the answers to those questions that interest us or pique our curiosity. Perhaps we might be satisfied with our explanations coming to an end here: some items of knowledge have constitutive value, and others don't, and there's no further explanation of where the boundary between those species of knowledge lies.

An alternative to this would be to argue that knowledge is always *epistemically* valuable, even if instances of knowledge differ in their eudaimonic value. But this leads us back to the obscurity of the notion of "epistemic value." Truth value monism seems to imply that all instances of knowledge are created equal, when it comes to their epistemic value. This would-be implication seems wrong: knowledge of particle physics seems not only eudaimonically but epistemically more valuable than knowledge of random phonebook entries. We can capture this intuition by saying something about "epistemic value": epistemic value accrues when there is *mind-to-world fit*, i.e. when there is accurate representation of the world. Knowledge of particle physics involves greater mind-to-world fit than does knowledge of random phonebook entries – as Nick Treanor (forthcoming a) puts it, when I know the truths of physics I *know more* than when I know trivial truths. We might explain this by appeal to the idea that some true propositions contain more truth than others (Treanor forthcoming b), or to the idea that belief in some true propositions "carves nature at the joints," while belief in others doesn't (Sider 2011).

Suppose we were satisfied with some version of the claim that knowledge is valuable. What explains this? *Why* is knowledge valuable? Epistemologists sometimes appeal, in connection with the value of knowledge, to the first sentence of Aristotle's *Metaphysics* on which "all men by nature desire knowledge." Appealing to our curiosity, in connection with the value of knowledge, is puzzling, if one assumes *realism about the value of knowledge*, on which whether something is valuable is independent of whether anyone values it. If realism is true, that everyone desires knowledge neither entails, nor is entailed by, that knowledge is valuable. What Aristotle has in mind, of course, is that virtuous people desire knowledge, that people are curious if they are functioning properly, in accord with their human nature. Aristotle is no anti-realist about the value of knowledge; but he's also not appealing to our mere curiosity in defense of the value of knowledge. And although in his famous sentence he says that everyone naturally wants knowledge (*to eidenai*), what he's really defending (in the passages that follow) is an interest in understanding of causes and principles (*episteme*) and in particular an interest in understanding of fundamental causes and principles (*sophia*).

That said, perhaps *anti-realism about the value of knowledge*, on which the value of knowledge is in some sense exhausted by the fact that people value knowledge, is the way to go. Recently some philosophers have defended and criticized *epistemic expressivism* (see Chrisman forthcoming), on which evaluative or normative epistemic discourse should be understood as (among other things) expressing the speaker's commitment to certain norms of belief formation. We might understand talk of the value of knowledge along similar lines.

Others have adopted a *hypothetical imperative approach* to the value of knowledge; on this view knowledge is good because it is a case of us successfully achieving something we desire (Papineau 1999, Dretske 2000; for criticism see Kelly 2003, Grimm 2008). This is a kind of realism (since it assumes the goodness of desire-satisfaction), but may not offend anti-realist intuitions.

Many epistemologists sympathetic with realism about epistemic value have turned to the idea that "belief aims at truth." Where Aristotle saw a teleological link between people and knowledge; these contemporary thinkers see a teleological link between belief and truth. This might provide the explanation we seek: if "belief aims at truth," then true belief (and therefore knowledge) is good, in as much as it is belief that has fulfilled its "aim." Put another way, if "belief aims at truth," then true belief (and therefore knowledge) is good, qua belief. But what does it mean to say that "belief aims at truth"? This metaphorical slogan must be given literal content. There are three possible approaches here. The first is a (relatively) *literal approach* that takes belief to require a psychologically real desire or intention (to believe the truth and nothing but the truth) on the part of any believer. This view is problematic in as much as this psychological claim is implausible. The second is a *Darwinian approach* that takes truth to be a biological function of belief; just as it is the biological "aim" of the heart to pump blood, it is the biological "aim" of belief to be true (Millikan 1984, Papineau 1993). This view is problematic in as much as it rests on empirically uninformed speculation about the natural history of cognition. The third is a *normative approach* that takes truth to be a "constitutive standard of correctness" for belief; on this view the goodness of true belief is a conceptual truth that flows a priori from the essential nature of belief (Wedgwood 2002, Shah 2003). This view is problematic in as much as it seems incompatible with philosophical naturalism; to escape this consequence, some of those who defend it favor epistemic expressivism (Shah and Velleman 2005).

ANNOTATED BIBLIOGRAPHY

DeRose, Keith. 1995. "Solving the Skeptical Problem", *The Philosophical Review*, 104: 1-52. Stewart Cohen and DeRose developed the contextualist theory of knowledge attributions. DeRose's 1995 article has become a standard in the field.

Shope, Robert. 2002. "Conditions and Analyses of Knowing," in P. Moser (ed.) *The Oxford Handbook of Epistemology*. Oxford University Press. Overview of issues arises in connection with the analysis of knowledge.

Sosa, Ernest. 1991. *Knowledge in Perspective*. Cambridge University Press. Influential essays from a leading contemporary philosopher, on topics ranging from the Gettier problem to virtue epistemology.

Stroud, Barry. 1984. *The Significance of Philosophical Skepticism*. Oxford University Press. An influential treatise on skepticism but before the contemporary emphasis on knowledge attributions. A powerful defense of the skeptic's reasoning.

Zagzebski, Linda. 1996. *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Wellbeing Foundations of Knowledge*. Cambridge University Press. Groundbreaking work on intellectual virtue, articulating and addressing some basic questions about the value of knowledge.

OTHER WORKS CITED

Ayer, A.J. 1956. *The Problem of Knowledge*. London: Macmillan.

Baril, A. 2010. "A Eudaimonist Approach to the Problem of Significance", *Acta Analytica*, 25: 215-41.

Buckwalter, Wesley, and Stich, Stephen. 2011. "Gender and the Philosophy Club", *The Philosopher's Magazine*, 52: 60-5.

Chrisman, Matthew. Forthcoming. "Epistemic Expressivism", *Philosophy Compass*.

Cohen, Stewart. 1988. "How to be a Fallibilist", *Philosophical Perspectives*, 2: 91-123.

Cohen, Stewart. 1999., "Contextualism, Skepticism, and the Structure of Reasons", *Philosophical Perspectives 13: Epistemology*: 57-89.

Craig, Edward. 1990. *Knowledge and the State of Nature*. Oxford University Press.

Cuneo, Terence. 2007. *The Normative Web*. Oxford University Press.

DeRose, Keith. 2004. "Single Scoreboard Semantics", *Philosophical Studies*, 119: 1-21.

DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism and Context, Volume I*. Oxford University Press.

Dogramaci, S. Forthcoming. "Reverse Engineering Epistemic Evaluations", *Philosophy and Phenomenological Research*.

Dretske, Fred. 2000. "Norms, History, and the Constitution of the Mental", in his *Perception, Knowledge, and Belief*. Cambridge University Press.

Dretske, Fred. 1970. "Epistemic Operators', *Journal of Philosophy*, 67: 1007-23.

Enç, Berent. 1990. "Is Realism Really the Best Hypothesis?" *Journal of Philosophy*, 87: 667-8.

Fantl, Jeremy and Matthew McGrath. 2009. *Knowledge in an Uncertain World.* Oxford University Press.

Feldman, Richard. 2005. "Epistemological Puzzles about Disagreement", in Stephen Hetherington (ed.) *Epistemology Futures*. Oxford: Oxford University Press.

Field, Hartry. 2009. "Epistemology without Metaphysics", *Philosophical Studies*, 143: 249-90.

Frances, Bryan. 2010. "The Reflective Epistemic Renegade", *Philosophy and Phenomenological Research*, 81: 419-63.

Frances, Bryan. Forthcoming. "Philosophical Renegades", in Jennifer Lackey and David Christensen (eds.) *New Essays on Disagreement*. Oxford: Oxford University Press.

Fumerton, Richard. 1995. *Metaepistemology and Skepticism.* Rowman & Littlefield.

Goldman, Alvin. 1967. "A Causal Theory of Knowing", *Journal of Philosophy*, 64: 357-72.

Grimm, Stephen. 2008. "Epistemic Goals and Epistemic Values", *Philosophy and Phenomenological Research*, 77: 725-44.

Hawthorne, John. 2004. *Knowledge and Lotteries*. New York and Oxford: Oxford University Press.

Hazlett, Allan. 2009. "Knowledge and Conversation", *Philosophy and Phenomenological Research*, 78: 591-620.

Hazlett, Allan. 2010. "The Myth of Factive Verbs", *Philosophy and Phenomenological Research*, 80: 497-522.

Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique", *Philosophy and Phenomenological Research*, 66: 612-40.

Kornblith, Hilary. 2010. "Belief in the Face of Controversy", in Richard Feldman and Ted Warfield (eds.) *Disagreement*. Oxford: Oxford University Press.

Kornblith, Hilary. 2006. "Appeals to Intuition and the Ambitions of Epistemology", in S. Heatherington (ed.), *Epistemology Futures* (Oxford University Press).

Kornblith, Hilary. 1993. "Epistemic Normativity," *Synthese* 94 (3): 357-76.

Kripke, Saul. 2011. "Nozick on Knowledge", in Saul Kripke (ed.), *Philosophical Troubles. Collected Papers Vol I*. Oxford University Press.

Kusch, Martin. 2009. "Testimony and the Value of Knowledge", in A. Haddock, A. Millar, and D. Pritchard (eds.) *Epistemic Value*. Oxford University Press.

Kvanvig, Jonathan. 2006. "Closure Principles", *Philosophy Compass*, 1: 256-67.

Kvanvig, Jonathan. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.

Lewis, David. 1996. “Elusive Knowledge”, *Australasian Journal of Philosophy*, 74: 549-67.

Lynch, Michael. 2004. *True to Life: Why Truth Matters*. MIT Press.

Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. MIT Press.

MacFarlane, John. 2005. “The Assessment Sensitivity of Knowledge Attributions”, *Oxford Studies in Epistemology*, 1: 197-233.

Moore, G.E. 1939. “Proof of an External World”, *Proceedings of the British Academy*, 25: 273-300.

Nagel, Jennifer. Forthcoming. “Intuitions and Experiments: A Defense of the Case Method in Epistemology”, *Philosophy and Phenomenological Research*.

Nagel, Thomas. 1979. “Moral Luck”, in his *Mortal Questions*. Cambridge University Press.

Nozick, Robert. 1981. *Philosophical Explanations.* Harvard University Press.

Papineau, David. 1993. *Philosophical Naturalism*. Blackwell.

Papineau, David. 1999. “Normativity and Judgment”, *Proceedings of the Aristotelian Society*, 73: 16-43.

Pritchard, Duncan. 2010. “Knowledge and Understanding”, in A. Haddock, A. Millar, and D. Pritchard, *The Nature and Value of Knowledge: Three Investigations*. Oxford University Press.

Pritchard, Duncan. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

Pryor, James. 2000. “The Skeptic and the Dogmatist”, *Noûs*, 34: 517–49.

Radford, Colin. 1966. “Knowledge—by Examples”, *Analysis*, 27: 1-11.

Ryle, Gilbert. 1949. *The Concept of Mind.* University of Chicago Press.

Rysiew, Patrick. 2011. “Epistemic Contextualism”, *The Stanford Encyclopedia of Philosophy (Winter 2011 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/win2011/entries/contextualism-epistemology/>.

Rysiew, Patrick. 2001. “The Context-Sensitivity of Knowledge Attributions”, *Noûs*, 35: 477-514.

Schaffer, Jonathan. 2005. “Contrastive Knowledge”, in Tamar Szabo Gendler & John Hawthorne (eds.) *Oxford Studies in Epistemology 1*. Oxford University Press.

Schaffer, Jonathan. 2007. “Knowing the Answer”, *Philosophy and Phenomenological Research*, 75: 383–403.

Shah, Nishi. 2003. “How Truth Governs Belief”, *Philosophical Review*, 112: 447-82.

Shah, Nishi & Velleman, David. 2005. “Doxastic Deliberation,” *Philosophical Review*, 114: 497-534.

Sider, Theodore. 2011. *Writing the Book of the World*. Oxford University Press.

Sosa, Ernest. 2003. “The Place of Truth in Epistemology”, in M. DePaul and L. Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press.

Sosa, Ernest. 2000. “Skepticism and Contextualism”, *Noûs*, 34: 1-18.

Sosa, Ernest. 1999a. “How Must Knowledge Be Modally Related to What Is Known?” *Philosophical Topics*, 26: 373-84.

Sosa, Ernest. 1999b. “How to Defeat Opposition to Moore”, *Philosophical Perspectives*, 13: 141-53.

Sosa, Ernest. 1991. *Knowledge in Perspective: Selected Essays in Epistemology.* Cambridge University Press.

Stanley, Jason. 2005. *Knowledge and Practical Interests*. New York and Oxford: Oxford University Press.

Stroud, Barry. 1994. “Scepticism, ‘Externalism’, and the Goal of Epistemology”, *Proceedings of the Aristotelian Society: Supplementary Volume*, LXVIII: 291-307.

Treanor, Nick. Forthcoming a. “The Measure of Knowledge”, *Noûs*.

Treanor, Nick. Forthcoming b. “Trivial Truths and the Aim of Inquiry”, *Philosophy and Phenomenological Research*.

Unger, Peter. 1975/2002. *Ignorance: The Case for Scepticism*. Oxford University Press.

Vogel, Jonathan. 1990. “Cartesian Skepticism and Inference to the Best Explanation”, *Journal of Philosophy*, 87: 658-66.

Wedgwood, Ralph. 2002. “The Aim of Belief,” *Philosophical Perspectives*, 16: 267-97.

Weinberg, Jonathan, Nichols, Shaun, and Stich, Stephen. 2001. “Normativity and Epistemic Intuitions”, *Philosophical Topics*, 29: 429-60.

White, Roger. 2010. “You Just Believe That Because…”, *Philosophical Perspectives*, 24: 573-615.

Williams, Bernard. 1973. “Deciding to Believe”, in his *Problems of the Self*. Cambridge University Press.

Williamson, Timothy. 2000. *Knowledge and Its Limits.* Oxford University Press.

Zagzebski, Linda. 2003. “Intellectual Motivation and the Good of Truth”, in M. DePaul and L. Zagzebski Linda. (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford University Press.