

Basic Knowledge First

Jonathan Jenkins Ichikawa

October 25, 2016

1 Knowledge First

A controversial but attractive twenty-first-century approach to knowledge takes knowledge to be somehow “first”.¹ The broad idea of the “knowledge first” program is that we should treat knowledge as in some way more fundamental than other epistemic notions like belief and justification. In just what this fundamentality consists is not always made explicit; in our (2016), Carrie Jenkins and I argued that there are two rather distinct families of views that travel under the banner “knowledge first”. According to *representational* versions of the knowledge first project, saying that knowledge is more fundamental than belief amounts to saying something about how we think and talk about epistemic notions. For example, a knowledge-first claim in the representational tradition might have it that the concept KNOWS doesn’t have BELIEVES as a constituent.² (A stronger priority claim might endorse the converse constitutional claim.) Or perhaps what is at issue has to do with whether, as a psychological matter, humans tend to determine whether someone has knowledge by thinking about beliefs, or vice versa.

Jenkins and I argued that these kinds of representational claims are to a considerable degree independent of the more *metaphysical* interpretations of the knowledge-first idea. These interpretations would make claims about the state of knowledge itself—whether, for example, it is a mental state in its own right, or whether it is some sort of composite of a mental state and a non-mental state like truth. One idea in this latter family concerns the metaphysical fundamentality of knowledge. Is the category of knowledge a contingent and artificial one, useful only as a convenient shorthand or an approximation for epistemic states of greater metaphysical significance?³ Or, on the other hand, would one’s metaphysical

¹The seminal work in this program is Williamson (2000).

²For example, Nagel (2013, p. 292) characterizes the question as that of “whether the concept of knowledge is in some sense prior to the concept of belief, or whether it is composed from that concept and further conditions.”

³Compare David Lewis’s remarks:

What is it all for? Why have a notion of knowledge that works in the way I described? (Not a compulsory question. Enough to observe that we do have it.) But I venture the guess that it is one of the messy short-cuts—like satisficing, like having indeterminate degrees of belief—that we resort to because we are not smart enough to live up to really high, perfectly Bayesian, standards of rationality. You cannot maintain a record of exactly which possibilities you have eliminated so far, much as you might like to. It is easier to keep track of which possibilities you have eliminated if you—Psst!—ignore many of all the possibilities there are. And besides, it is easier to list some of the propositions that are true in *all* the uneliminated, unignored possibilities than it is to find propositions that are true in *all and only* the uneliminated, unignored possibilities.

picture of the world miss out on something important for its failure to categorise states as knowledge? I agree with recent metaphysicians who have argued that there is a kind of hierarchy of metaphysical fundamentality; the category *green* carves better at nature's joints than does the category *grue*.⁴ Although it is possible to characterise either in terms of the other, describing *grue* in terms of *green* more accurately characterises the structure of the world than would describing *green* in terms of *grue*.

Grue is a highly artificial category; one might be suspicious of it on general grounds. But not all less fundamental categories are highly artificial; livers are less fundamental than cells, but they are metaphysically significant entities in their own rights. Medical theories should not eschew livers and theorise in terms of cells instead, even if it is possible to explain the former in terms of the latter. The fundamentality in question comes in degrees.

Consequently, the fundamentality interpretation of the knowledge first program can also be interpreted in stronger or weaker ways. A very strong fundamentality interpretation of the knowledge first program would have it that knowledge is among the *wholly* fundamental categories in the world. It is inconsistent with the metaphysical reduction of epistemology to physics (or to anything else). A weaker—and more plausible—interpretation would have it simply that knowledge is *more* metaphysically fundamental than one might have thought. Perhaps, for instance, it's more fundamental than are belief and justification. Even if not strictly *first*, it holds that knowledge comes *pretty early* in putting together one's picture of the world.

It is therefore consistent with this interpretation of the knowledge first thesis that, even though knowledge is rather fundamental—much more fundamental than was assumed by the post-Gettier literature, for instance—there might be other epistemic states of even greater metaphysical fundamentality, in terms of which knowledge itself might be explained. In this paper, I will try to make the case that this is so. I'll suggest that ought to recognise a distinctive and more fundamental category of *basic knowledge*, that will help to illuminate the category of knowledge itself.

On the whole, I take my project to be a development of the knowledge first project, rather than a competitor to it. This interpretation is certainly debatable; for example, I will argue below that some of the roles sometimes attributed to knowledge, for example by Timothy Williamson, ought to be attributed to basic knowledge instead. This does count against the metaphysical priority of knowledge. I think of my approach generally friendly to a 'knowledge first' program because I agree with people like Williamson who think that knowledge comes first compared to notions like belief and justification. But I think

If you doubt that the word 'know' bears any real load in science or in metaphysics, I partly agree. The serious business of science has to do not with knowledge *per se*; but rather, with the elimination of possibilities through the evidence of perception, memory, etc., and with the changes that one's belief system would (or might or should) undergo under the impact of such eliminations. Ascriptions of knowledge to yourself or others are a very sloppy way of conveying very incomplete information about the elimination of possibilities. . . .

The only excuse for giving information about what really matters in such a sloppy way is that at least it is easy and quick! But it *is* easy and quick; whereas giving full and precise information about which possibilities have been eliminated seems to be extremely difficult, as witness the futile search for a 'pure observation language'. If I am right about how ascriptions of knowledge work, they are a handy but humble approximation. (Lewis, 1996, p. 563)

⁴See e.g. Sider (2011), Koslicki (2012). *Green* itself is not perfectly fundamental—it is not obvious that we'd be missing out on structural facts about the world if we drew the colour boundaries in different places—but unlike *grue*, it at least is definable in terms of a continuous range of light frequencies.

basic knowledge comes first compared to knowledge. Whether this should be considered a ‘knowledge first’ view at the end of the day seems largely a terminological matter.

2 Knowledge and Action

I’ll begin the case by following Williamson in discussing an analogy between knowledge and action. I agree that the analogy is instructive, and that it motivates the primacy of knowledge over belief. But unlike Williamson, I will argue that it also supports a distinct category of *basic* knowledge that enjoys primacy over knowledge generally. *Knowledge and Its Limits* begins thus:

Knowledge and action are the central relations between mind and world. In action, world is adapted to mind. In knowledge, mind is adapted to world. When world is maladapted to mind, there is a residue of desire. When mind is maladapted to world, there is a residue of belief. Desire aspires to action; belief aspires to knowledge. The point of desire is action; the point of belief is knowledge.⁵

In more recent work,⁶ Williamson has extended and clarified this analogy. Belief is something like an attempt at, or gesture towards, knowledge—what it is to believe is to do the sort of thing where, if all goes well, one ends up knowing. The same relationship holds between acting and intending—action is a central way of engaging the world with a world-to-mind direction of fit; intending is roughly a commitment towards acting.

Consider an example. Cheryl is near the dog, who has a tangled coat. A world-to-mind relation betwixt these two is Cheryl’s *knowledge that the dog has a tangled coat*. It is her successful receptive engagement with the world. Her *belief that the dog has a tangled coat* is involved in her knowing it, but on Williamson’s picture, the knowledge is more fundamentally explanatory.

Perhaps Cheryl wants to remove the tangle, so she performs an action: she *brushes the dog*. This is an active, mind-to-world, relation between Cheryl and the dog. She has an idea about how she wants things to be—namely, such that the dog is brushed—and so she performs an action. She brushes the dog. This action, too, has a kind of internal correlate: her intention to brush the dog. Just as it’s possible, for example in a skeptical scenario, to believe that the dog has a tangled coat when she really doesn’t, so too is it possible, for example in a difficult practical environment, to intend to brush the dog, even though one doesn’t brush the dog. Belief doesn’t guarantee knowledge, and intention doesn’t guarantee action.

Williamson writes:

It is obviously wrong-headed to try to understand intention without reference to action. Given the connection between belief and intention, it is also wrong-headed to try to understand belief without at least indirect reference to action. Furthermore, if knowledge is to belief as action is to intention, it is wrong-headed to try to understand belief without reference to knowledge, or to try to understand intention without at least indirect reference to knowledge.⁷

⁵Williamson (2000, p. 1)

⁶Williamson (2016a)

⁷Williamson (2016a, p. 14)

There is an affinity between the set-up here and G. E. M. Anscombe’s approach to action in *Intentions*. Anscombe resisted the factorisation of action into events combined with a kind of intentionality. ‘We do not add anything,’ she wrote, by ‘attaching to the action at the time it is done by describing it as intentional.’⁸ Compare Wittgenstein’s famous question, ‘what is left over if I subtract the fact that my arm goes up from the fact that I raise my arm?’⁹ One might try to ask, ‘what is left over if I subtract the fact that the dog is hungry from the fact that I know that the dog is hungry?’—but this way lies the Gettier program. Better, perhaps, to take knowledge and action as the primitives.

Ernest Sosa has also often invoked an analogy between knowledge and action.¹⁰ His point of emphasis, however, tends to be more about the characteristic *success* of a particular action than in the metaphysical status of the action itself. He invokes, for example, the idea of an archer who takes a shot at the target, as analogous to the formation of a belief. A belief is successful just in case it is true; a shot is successful just in case it hits its target. Knowledge, Sosa says, is a particular kind of successful belief—it is a belief that is successful in a way expressive of its competent formation. This picture is compatible with the remarks above; perhaps the normative relation between belief and truth is similar to that between shooting at a target and hitting it.¹¹ The question at issue for the present knowledge-first claim is whether the status of belief is metaphysically independent of that of knowledge. One might think not. (Similarly, one might think that what it is to *be a shot at a target* is metaphysically dependent on a more fundamental notion of an apt hitting of the target.)¹²

I propose simply to take on board Williamson’s central analogy between knowledge and action, and his argument that it is suggestive of a knowledge first (before belief) approach to epistemology. As we shall see, the analogy, pressed further, also points to a basic knowledge first (before knowledge) approach to epistemology.

3 Basic Action and Basic Knowledge

I drove my dog to the beach this afternoon. I didn’t do it by accident; it was a deliberate action. Here’s another thing I did this afternoon: I started the ignition in my car. That wasn’t an accident either; it was an intentional action too. There are various senses in which these two actions are not on a par. For example, starting the ignition was *part of* the process whereby I drove my dog to the beach. I drove her to the beach *by* starting the ignition, and doing some other things. There is a sense, then, in which the starting of the ignition was a *more fundamental action* than was driving my dog to the beach.

Some of the actions we perform, we perform by performing other actions. On pain of infinite regress, however, it seems that there must be some *basic actions*—actions that we perform, not by virtue of performing any other actions. If everything I did, I had to do by doing something else, then, assuming circular chains of action dependence are inadmissible, I’d have to perform an infinite series of increasingly basic actions to do anything at all.¹³

⁸Anscombe (1957, p. 28). Ford (2011) draws out some of the parallels between Anscombe’s approach to action and Williamson’s approach to knowledge.

⁹Wittgenstein (1998, §621).

¹⁰For example, in Sosa (2007).

¹¹But see Miracchi (2015a) and Miracchi (2015b) for an opposing view in a similar spirit.

¹²My commitments here are also neutral on the question of whether setting things out in this way implies that normativity in the domains in question need be thought of in terms of success; see e.g. McKinnon (2015, ch. 8) for a denial of this suggestion.

¹³Roughly this argument is given in Danto (1979). Sneddon (2001, p. 509) accepts that the argument shows that some actions are done directly, but denies that this establishes ‘a special category of’ basic actions. But I define basic actions as those done directly. Lavin (2013) argues that two similar regress

The issue isn't merely that in performing any action I would thereby need to be doing an infinite number of things—that would be an argument along the lines of Xeno's Paradox.¹⁴ The objectionable consequence is that, any time I perform an action, there is an infinite series of increasingly basic actions I perform, and I do each *by* doing the next.

I don't wish to assume too much of substance about just what basic actions consist in. I think the argument above shows that there are basic actions, but it doesn't tell us what they are.¹⁵ There is some temptation to suppose that the basic actions are attempts, or formations of intentions, or the kind of mental performance that guarantees success when an agent intends them.¹⁶ A more moderate restriction in the same direction would limit basic actions to the manipulations of local bodily regions.¹⁷ But I'll suggest below that there may be good reason to resist that temptation—maybe taking a sip of beer can be a basic action.

At an intuitive level, it is very plausible that starting the ignition in a car is a more basic action than is driving a dog to the beach. Perhaps one starts the ignition, too, by performing a more basic action, such as turning a key. Perhaps turning a key is something one does directly, or perhaps one turns a key by moving one's hand in a certain way, or sending a kind of electrical signal down one's arm, or activating a certain neural pattern. We can leave this open, and be guided by our best theory of basic action—though I'll register at the outset that it doesn't seem very plausible to me that activating certain neural patterns is an action I perform at all, let alone basically. Vivian Weil and Irving Thalberg assume that basic actions cannot depend on social norms and knowledge¹⁸—I do not. The only point I wish to rest on here is that there is such a thing as basic action. We do some things by doing other things, but some things, we do directly. I do not even assume that there is a general way to categorise the kinds of actions we do or can perform directly; this may vary according to circumstance or level of expertise. Maybe sometimes some photographers *focus on the eye* as a basic action, while in other cases, one can only focus on the eye by performing the action of manipulating the focus ring in a particular way. But for any token action, it is either basic or non-basic, and if it is non-basic, there is a basic action underwriting it.

If this is right, then the analogy between knowledge and action is suggestive of a special category of basic *knowledge* as well. Some things that we know, we know directly; others, we know by knowing more basic things. This is plausible independently of the analogy. Reasoning parallel to that given above about basic action is a standard motivation for *foundationalism*, the view that some of our knowledge is basic, i.e., not grounded in other knowledge. If everything I knew, I had to know by knowing some other things, then, assuming circular chains of epistemic dependence are inadmissible, I'd have to know an infinite series of increasingly basic items of knowledge to know anything at all.

It's also plausible, when one looks at examples, that some knowledge is more basic than other knowledge. I know that South Africa was governed under apartheid for much of the twentieth century, but my knowledge of this fact is indirect. I wasn't there. Part of the way I know about apartheid is by knowing what various historical records show, and what

arguments in favour of basic action are inconclusive, but he doesn't consider one of this form. Lavin may be assuming a tight connection between basic action and the analysability of action into the action's result and something 'purely mental'. See e.g. his p. 282. I agree that we need not admit basic action in any such strong sense. But it seems hard to square the idea that everything we do, we do by doing other things.

¹⁴Thanks to David Sosa for helping clarify this issue for me.

¹⁵Sneddon (2001, pp. 505–6) and Lavin (2013, pp. 276–7) challenge the intuitive identification of basic actions. I agree that this is not a matter that is straightforwardly open to reflection.

¹⁶See e.g. McCann (1974) and Hornsby (1980).

¹⁷See e.g. Weil and Thalberg (1974).

¹⁸Weil and Thalberg (1974, pp. 116–7)

various trustworthy people have told me, etc. As in the case of basic action, we don't need to decide at this point just what basic knowledge consists in.¹⁹ There is some pressure to suppose that our basic knowledge concerns only how things seem, or that which is totally indubitable, or that about which we're guaranteed to be right. But again, there may well be reasons to resist this picture; I'll discuss this in the next section. For now, I wish to assume only that there is a category of basic knowledge. There is such a thing as knowing something *directly*.

So there is action, and basic action, and knowledge, and basic knowledge. The idea that basic action is more fundamental than action simpliciter is a natural and appealing one; so too is the idea that, although knowledge is a somewhat fundamental mental state—more so than most twentieth-century epistemologists thought—*basic* knowledge is even *more* fundamental than knowledge simpliciter. In the remainder of the paper, I'll consider its prospects.

4 Foundationalism

As I said in the last section, there is good reason to be a foundationalist. Not all of one's knowledge is grounded in independently established knowledge. Foundationalism has a negative connotation in some philosophical circles, due to its associations with Cartesian internalism. But Cartesian internalism is the combination of two ideas: foundationalism and a particular criterion for foundations, namely the idea that epistemic foundations are phenomenal features that are in some sense indubitable. Descartes sought foundations upon which not even the possibility of an evil demon of a particular description could cast doubt. These two Cartesian commitments are separable. This section will sketch a few details of my own preferred approach to basic knowledge. Strictly speaking, it is independent from the main project of the rest of the paper—the idea that basic knowledge comes first is consistent with a variety of approaches to basic knowledge. But I wish to be very clear that in committing to basic knowledge, I am *not* committing to a Cartesian understanding of it. It will also be useful, in making the case for basic knowledge first, to be able to write less abstractly about some of the examples.

One reason one may wish to resist the Cartesian approach to foundationalism is that it is poorly motivated. For example, if it is motivated by the idea that there is an epistemically special internal realm to which one has infallible access, rejecting that idea amounts to rejecting the motivation. There are many contemporary points of pressure against this idea. One is Timothy Williamson's 'anti-luminosity' argument.²⁰ Another, more direct, challenge concerns empirical investigation of the epistemic difficulties of knowledge of the internal. For example, Eric Schwitzgebel's (2011) *Perplexities of Consciousness* offers several case studies that suggest that it's very difficult for many of us to know many things about our own phenomenal experiences. If Descartes was motivated by the idea that, since it's possible to be wrong about external matters, external matters can't be foundational, then just the same argument seems to show that *internal* matters can't be foundational either.

I reject the phenomenalist commitment of Cartesian foundationalism. But this is not to reject foundationalism itself—I accept an important epistemic difference between things that are known directly, and things that are only known by virtue of their support from other

¹⁹The example just given assumes that testimonial knowledge is not typically basic; this is a widely accepted view, but it is not without dissenters; see e.g. Audi (2006) and Cleve (2006), who attribute it to Thomas Reid. (Audi agrees with Reid on this matter; van Cleve does not.) I do not find the Reidian line compelling, but if one did, one could retain my view with different examples.

²⁰Williamson (2000, Ch. 4). Williamson argues that, given a modest safety requirement on knowledge, there can be no non-trivial state such that anyone in it must be in a position to know that they're in it.

things known. A full accounting and catalog of basic knowledge is well beyond my present scope, but I think there's good reason to think that factive perceptual states sometimes provide direct knowledge of the external world. I know I have hands; this knowledge is direct. I see them for myself. I don't, I think, know that I have hands by knowing that I have the appearance of hands. So I propose to let the objects of factive perceptual states play a foundational role. (I don't by any means assume these are the only foundations.)

In fact, I think it's a mistake to suppose that Cartesian reflections ever provided an *argument* for his phenomenal conception of basic knowledge. Although Descartes invokes a thought experiment involving an evil demon, his assumptions about which thought experiments were relevant in effect brought in Cartesian internalism as an undefended assumption. As I wrote in Ichikawa (2017):

Remember again that the precarious situation of adopting anti-skeptical strategies to some worries but not to others is ultimately one shared by *all* foundationalists, including the more familiar Cartesian ones. After all, it's not impossible to construct something that at least looks like a skeptical scenario with respect to introspective judgments too. For example, suppose that Edwin hasn't shaved since yesterday, and so has some stubble. Angelina looks at Edwin, and sees that he has stubble. Moreover, when she introspects, she notices that she has the perceptual experience as of seeing someone with stubble. One can challenge her perceptual knowledge by invoking a Cartesian skeptical scenario like this one:

Demon. Edwin doesn't really exist; Angelina isn't really looking at anyone with stubble. Her experience as of seeing someone with stubble is caused by an evil demon who is directly manipulating her perceptual experiences.

This Cartesian scenario challenges her perceptual knowledge, but takes her introspective knowledge for granted. But one could challenge the latter with an even *more* radical skeptical scenario:

Superdemon. Not only is Angelina not really looking at anyone with stubble; she doesn't really even have experiences *as if* she is looking at someone with stubble. Her introspective judgment of having such experiences is caused by an evil superdemon who is directly manipulating her introspective judgments.

The Superdemon scenario is a challenge to the nonskeptical verdict that Angelina knows that it appears to her as if Edwin has stubble. It is a possible case where it doesn't so appear, and it is one which, according to the challenge, Angelina's evidence is unable to rule out. The Cartesian considers this challenge to be illegitimate, perhaps along lines such as these: the scenario in question, while possible, is one in which things seem different to Angelina; so her evidence—the way things seem to her—*does* rule out the Superdemon case. This seems like a reasonable response to the Superdemon case. But notice that it is exactly parallel to the Moorean response to the standard Demon case: the scenario in question, while possible, is one in which Angelina doesn't see the same things she actually sees; so her evidence—the contents of her factive perceptual states—*does* rule out the Demon case. So the Cartesian has to explain away certain

skeptical scenarios too; he posits a difference between Demon and Superdemon, where I posit the relevant difference instead between Demon and the challenge to Angelina’s inference that Edwin hasn’t shaved today that we might call ‘Tonic’:

Tonic. Edwin did shave today, but he used a special tonic that made him grow new stubble very quickly, in an attempt to deceive Angelina into believing he hadn’t shaved today.

Tonic challenges Angelina’s inferential knowledge that Edwin didn’t shave today. On my [view], Tonic represents a legitimate skeptical challenge in the way Demon and Superdemon do not.

I see no reason to suppose it better to draw the line between Demon and Superdemon, as the Cartesian does, than it is to draw it between Tonic and Demon, as on the kind of Moorean line I’m considering. So I’ll assume there is basic knowledge of the external. Many more questions remain—just which external things are known directly, and which are inferred? My approach (though not necessarily my examples) will be compatible with a variety of answers to that question. As in the case of basic action, I do not assume that basic knowledge can be characterised in terms of its subject matter—some tokens of perceptual knowledge that there is an apple before one are basic, while others will not. (Imagine someone less familiar with apples, who needs to deliberate before identifying the object as an apple.)

5 Non-Basic Knowledge

Basic knowledge is knowledge that isn’t known by virtue of other knowledge. What should we think about the category of *non-basic* knowledge?

The first choice is to deny that there is any such thing, holding that all knowledge is basic knowledge. One might take this line with a more or a less skeptical stance. The more skeptical stance says that we have much less knowledge than we often think we do. If things we see for ourselves are basically known, then many people know they have hands; but perhaps no one, or nearly no one, knows whether they will be alive next week, or whether Abraham Lincoln existed, or where their car is now parked.²¹ An advantage of this view is that it explains the sense in which knowledge is *conclusive*—to know something is to have it settled; on this approach, what is known is only that which one’s basic knowledge guarantees to be true.

(I am assuming here, as I do throughout the paper, that testimonial knowledge is not basic. If e.g. Audi (2006) is right that testimony is a source of basic knowledge—if, when you tell me something, I end up knowing it, even without basing this fact on knowledge of e.g. what you’ve said—then limiting knowledge to be basic would be less stingy. But it would still be rather stingy—it would still imply, for instance, that there is no abductive knowledge.)

A disadvantage of this skeptical stance, of course, is that like skeptical stances generally, it is revisionary.²² Many of the things we take ourselves to know are not plausibly conclusively established by things we know directly. So this sort of story might have to be supplemented

²¹One might wish to supplement basic knowledge with a closure principle, allowing that which one recognises deductively to follow from one’s basic knowledge also to be known. The result will still be a very stingy approach to knowledge.

²²Dodd (2007) critiques Williamson’s approach to evidence, saying it has this implication.

by some kind of error theory that explains why we think we have so much more knowledge than we do.²³ Depending on what kinds of theoretical principles involving knowledge we sign up to, this sort of approach might also imply a need for a radical change in our practices more generally.²⁴

One needn't be a skeptic, however, to hold that only basic knowledge is known—one might think instead that more of what we take ourselves to know is basic than we'd previously thought.²⁵ At some points in *Knowledge and Its Limits*, Williamson himself seems to favour a view of this sort. He writes of his approach to evidence that it 'suggests a very modest kind of foundationalism, on which all one's knowledge serves as the foundation for all one's justified beliefs.'²⁶

Read literally, the idea that all of one's knowledge is foundational straightforwardly implies that none of one's knowledge is based on other knowledge. As we have seen, this is prima facie a surprising claim; it is more natural to suppose that some knowledge depends on other knowledge. Although I'd never explicitly considered the question before this moment, I think I know that Timothy Williamson wrote a doctoral dissertation; this is indirect knowledge that rests on more basic knowledge, such as my knowledge that Williamson has a Ph.D. Although the idea that all knowledge is foundational doesn't play a central role in Williamson (2000), I'll suggest below that it fits in an interesting way with some of Williamson's more recent work in epistemology.

If we do suppose that there is non-basic knowledge, we should consider its extent, and how to theorise about it. There are many possibilities; in the interest of specificity and concision, I'll just assume a relevant alternatives approach. I suspect that much of what I will go on to say is translatable into other frameworks, but I won't undertake that here.

Think of knowledge generally as matter of eliminating counter-possibilities. To know that p requires ruling out all relevant not- p cases. Is there an unchanging standard that determines which possibilities are relevant, or does it vary according to the conversational context?²⁷ Here are two reasons to prefer a contextualist implementation.

First, if there's a single privileged epistemic standard, we might inquire into the basis for its privileged status. Why think there is a 'joint in nature' at the line indicated? We may assume that there is a joint in nature corresponding to basic knowledge—the question now is whether there's *another* corresponding to knowledge in general.

I raised a version of this challenge in Ichikawa (2017). I wrote:

²³Davis (2007) and Davis (2017) offers a story of this sort in defence of a view along these lines.

²⁴Compare Williamson: "In some ways, perceiving is an even closer analogue than knowing to acting. Just as acting typically involves direct output from the agent to the external environment, so perceiving typically involves direct input from the external environment to the agent. Just as acting involves a causal flow from the agent, which constrains what can be intentionally done, so perceiving involves a causal flow to the agent, which constrains what can be perceived. For instance, if it is necessary that P, only in very limited cases can one intentionally be such that P or sensorily perceive that P; there is no similar limitation on knowing that P. But to replace knowing by perceiving as the analogue of acting would also incur significant theoretical costs, for the contents of perception play no distinctive role in practical reasoning; any content of knowledge will do." Williamson (2016a, p. 21)

²⁵A reasonably close analog here is the 'Holmesian inference' of Bird (2005); Bird argues that there is no knowledge on the basis of non-entailing evidence, but argues that much of what we would have placed into that category actually *is* based on entailing evidence.

²⁶Williamson (2000, p. 186)

²⁷As I'm using the term, 'standards' are going to be something like functions from situations in which a subject is in to sets of relevant alternatives. So views like 'interest-relative invariantism', according to which a subject's practical stakes make a difference for whether one has to rule out certain possibilities, will count as views with an unchanging standard. The standard itself encodes sensitivity to the stakes. See Ichikawa (2017, ch. 1). Another possibility would have it that the context of the *evaluator*, rather than the speaker, fixes the standards.

Consider, from the point of view of an invariantist, a case of ‘barely-knowledge’ and one of ‘almost-knowledge’—maybe the one and only standard for knowledge is barely met by my belief about the current location of my car, but barely missed by my belief about what I’m going to cover in class next Friday. It ... seems at least potentially ad hoc for the invariantist to draw the line at a particular point. One might wonder what theoretical justification there is for highlighting that particular standard as knowledge-defining.

Second, insofar as one is motivated even in part by capturing ordinary intuitions about who knows what, an important fact with which one must reckon is that many of these intuitions seem to vary in systematic ways. If the contextualists are right, they shift in ways indicative of context-sensitive discourse in general. Non-basic knowledge is just the sort of thing we tend to want to call ‘knowledge’ in less skeptical contexts, but not in more skeptical ones.

So my preference is for the contextualist implementation of the idea. Non-basic knowledge is described as ‘known’ in some contexts, but not in others. This doesn’t mean knowledge isn’t somewhat joint-carving, especially if the epistemic standards in question themselves have significant theoretical roles to play. (One of the central ideas in Ichikawa (2017) is that they do.)

With this framework for basic and non-basic knowledge in place, the stage is set for an investigation into the degree to the theoretical roles for the latter.

6 Evidence

One theoretical role often attributed to knowledge in a knowledge first framework is an explanation of evidence. Williamson famously suggested that a subject’s evidence is all and only that which they know: $E=K$. The proposal was motivated in part by the thought that, in the propositional sense of ‘evidence’, to have some evidence is to stand in an epistemic relationship to it; given a knowledge-first framework, knowledge is an attractive candidate, particularly in light of the apparent failure of explanation of evidence in terms of states like justified belief.²⁸ But $E=K$ faces challenges too—particularly when one focuses on cases of non-basic knowledge.

Consider, for example, this objection from Brian Weatherson.²⁹ Take a case of *abductive knowledge*, defined as knowledge on the basis of non-entailing evidence. Here is Weatherson’s example:

On Thursday night Trinity watches the new hit movie *The Moroccan Falcon*. She sees that the hero dies in the final scene. Her friend Magdalen is going to see the movie Friday night. If Trinity tells Magdalen that when she saw the movie the hero died in the final scene, Magdalen might well complain that the movie has been spoiled because she now know how it ends. It wouldn’t really do for Trinity to argue that she is only reporting how the movie ended on the Thursday screening, and that for all anyone knows it might end differently on Friday. The mere possibility that the movie has alternate endings hardly defeats Magdalen’s claim that she now know how the movie will end.

²⁸See Williamson (2000, pp. 200–1).

²⁹Weatherson (2009), discussed in Ichikawa (2013). The present discussion rehearses some of my thoughts from that paper.

Let p be the proposition that the hero dies in the final scene of the Thursday screening of *The Moroccan Falcon*. Let q be the proposition that the hero dies in the final scene of the Friday screening of *The Moroccan Falcon*. Consider Trinity's relationship to the propositions $p \mathcal{E} \sim q$ and $q \mathcal{E} \sim p$. I've argued that she knows that each of these are false. So according to E=K the falsity of each conjunction is part of her evidence. Intuitively, $q \mathcal{E} \sim p$ is inconsistent with Trinity's evidence since she saw that p is true, while $p \mathcal{E} \sim q$ is consistent with her evidence, though not with what she knows. So q is not part of her evidence even though she knows it, so E \neq K.

It is an important part of Weatherson's case that Trinity's knowledge about the Friday showing is based on non-entailing evidence.³⁰ One could also consider more obviously inductive cases—as Weatherson points out, a case very similar to one of Williamson's own also generates intuitions inconsistent with E=K. Suppose one draws balls from an urn with replacement, and observes that each ball drawn is black. After a long enough series, for a small enough urn, one may eventually come to know inductively that every ball in the urn is black. Nevertheless, Weatherson says, it's intuitively *consistent with the subject's evidence* that there is a red ball that coincidentally has not yet been drawn.

One possible response to Weatherson's argument would be to strengthen Williamson's approach to evidence—maybe Williamson is right that all evidence is known, but wrong that *all* knowledge is evidence. Maybe only the basic knowledge is evidence. This would straightforwardly handle Weatherson's intuitions.

Alvin Goldman's proposed alternative to E=K has it that evidence is *non-inferentially justified belief*. Since this notion is probably not factive,³¹ but evidence plausibly is,³² I don't prefer Goldman's view. But what of the idea that evidence is non-inferential *knowledge*—i.e., basic knowledge?

I certainly feel some pull towards identifying evidence with basic knowledge. Nevertheless, I don't ultimately do so. The pull I feel towards E=BK is very similar to the skeptical pull discussed in the previous section, denying that non-basic knowledge is knowledge at all. Relative to modest conversational projects, non-basic knowledge still counts as definitive in a way suggestive of its status as evidence. There are low-standards evidence ascriptions that seem correct, just as there are knowledge ascriptions. If I consider what evidence I have that bears on the hypothesis that my dog has a cold, one of the relevant facts—a piece of evidence that I have—is that she was sneezing yesterday. This, even though my knowledge that she was sneezing yesterday isn't basic. (The dog-sitter told me.)

So I'm attracted to a contextualist approach to evidence ascriptions alongside the approach to knowledge ascriptions. What is 'evidence' relative to a given conversational context is the same as what is 'known' relative to that context. Just as there is a privileged category of basic knowledge, there is an identical category of basic *evidence*. Again, many different approaches to basic evidence are available; my preference is a relatively rich one, that includes the objects of factive perceptual states.

As we move into more skeptical conversational contexts, however, it becomes harder to satisfy 'knows'. Just what it requires may be characterised in terms of basic evidence—to know that p , a subject's basic evidence must conclusively rule out every relevant alterna-

³⁰So Weatherson is assuming that Trinity's evidence doesn't include the fact that this movie theatre wouldn't switch out alternate endings of movies on different days without her having heard about it.

³¹See the next section.

³²See Williamson (2000, p. 201), Littlejohn (2012, p. 101), or Littlejohn (2013).

tive to p . In non-skeptical contexts, this is often easily done; in more skeptical contexts, ‘knowledge’ may be unavailable.

As a matter of natural language, then, the categories of knowledge and evidence end up playing semantic roles. But when it comes to the nature of evidence and the way that evidence rules out possibilities, basic knowledge seems to cut closer to the epistemological joints in nature.³³

7 Justification

Justified belief is a natural candidate for a knowledge-first explication. Characterising knowledge in terms of justified belief seems to have been a failure; might we explain justified belief in terms of knowledge instead?

One way to explain justification in terms of knowledge, following Williamson’s model about evidence, would simply identify the two. One may adopt ‘ $J=K$ ’, holding that a subject has a justified belief just in case it is a case of knowledge.³⁴ $J=K$ is often motivated by the thought that knowledge is a norm for belief—a belief is appropriate, in a sense, just in case it is knowledge. The thought behind $J=K$ is that justification measures the appropriateness of a belief in this sense.

I am inclined to agree that there is a sense in which a belief is proper just in case it is knowledge.³⁵ But I am reluctant to follow the $J=K$ theorists in supposing that this is the only normative measurement available, or the best candidate for the notion of justification. Epistemological orthodoxy recognises a distinction between justified belief and knowledge; to identify the two is to collapse a distinction. I am not convinced that the distinction is without value.

In Ichikawa (2014), I defended a different approach to justification in terms of knowledge. I suggested that we explain justification as a kind of internal compatibility with knowledge—to have a justified belief, on that view, was to be in an intrinsic state consistent with having knowledge. The thought behind the view was that there are two importantly different ways for a belief to fail to be knowledge—it’s one thing for it to fail because the external world doesn’t line up the way it’d need to for you to know; it’s another for it to fail because what you did on the inside meant you never even had the chance to know. So I agreed with these remarks from Martin Smith:

Justification, I take it, is what makes a belief a good candidate for knowledge.
I’m happy to leave this suggestion somewhat imprecise for the time being, but

³³The ideas of this section are developed further in Ichikawa (2013) and Ichikawa (2017, ch. 3).

³⁴This view has been defended by Sutton (2007), Littlejohn (2014), and Williamson (2016b).

³⁵This is not uncontroversial; one source of dissent is the observation that beliefs are sometimes useful in ways having nothing to do with whether they amount to knowledge; for example, it can be pragmatically useful to believe, contrary the evidence, that one will succeed at a difficult task. See Marušić (2011). And as Kathryn Pogin reminded me, paradigmatically epistemically irrational beliefs—certain kinds of delusions, for instance—may be critical survival tactics for sufferers of trauma. I think these cases show that there is more to the overall value of beliefs than whether they amount to knowledge. But this needn’t necessarily be at odds with a knowledge norm of belief—even if such cases are very widespread. The knowledge norm says roughly that there is a distinctive kind of epistemic normativity to belief; it is in belief’s nature that it has an epistemic aim—to be knowledge. Consequently, any belief that isn’t knowledge is thereby worse *in this distinctive epistemic way*. Still, there may be good reasons of other sorts why it might be beneficial to have a belief that is defective *qua* belief. Some unjustified beliefs might help one have a better life, for instance. As a rough analogy, consider Sosa’s case of an archer shooting at a target; by its nature, a shot at a target is something that fails if it misses. But a missed shot may still be useful in many ways—it may provide good practice, or it may impress spectators, or it may lighten the mood. For more on the distinctively epistemic teleology of belief, see e.g. Miracchi (2015a).

the basic idea is this: My belief is justified just in case I have *done my epistemic bit*—the rest, as it were, is up to fate. My belief will qualify as knowledge provided that the world obliges or cooperates—but *I* am not required to do anything further. Smith (2010, p. 12)

Since publishing that paper, I’ve reconsidered the motivation for casting the intrinsic in such a central role. Given my preferred externalist approach to basic knowledge, it’s not obvious why one should demarcate justification as consistency of the *intrinsic* with knowledge. I still think, as I did then, that it is useful to separate failures to know that from the subject from failures to know that arise from the broader world, but in my newer work on this topic,³⁶ I characterise the difference in terms of basic knowledge and cognitive processing:

JPK S’s belief is justified iff there is a possible individual, alike to S with respect to all relevant basic evidence and cognitive processing, whose corresponding belief is knowledge.

If one adopted the Cartesian assumption that basic knowledge just *is* knowledge of the intrinsic, then the present, more general, formulation of JPK will be very close to the version I defended before. But if, as I do now, one supposes that perception affords us with basic knowledge of the external world, then JPK will have it that for my belief to be justified, it must be possible for someone with the same basic knowledge, *including knowledge of the external things I can see for myself*, to know.

On the more generalised picture of JPK, we suppose there is a particular realm over which we have a special epistemic responsibility. Fulfilling this responsibility is not in any sense trivial—I do not claim, for example, that basic evidence is luminous or anything like it. Williamson is, I think, right to say that there isn’t a “cognitive home” in the sense of a domain about which one cannot be mistaken, or to which one is guaranteed epistemic access. But I do think there is a different, more fully normative, notion for which we might repurpose the label “cognitive home”. The cognitive home is the domain over which one has the kind of distinctive epistemic responsibility at issue in justification. The justification norm is a requirement that one keep one’s cognitive home in order.

Given the contextualist approach gestured at above, when it comes to the non-basic, things might be ‘known’ relative to lax conversational contexts but ‘unknown’ relative to stricter ones. The contextualist framework is straightforwardly applied to JPK as well. I left the dog asleep downstairs half an hour ago; my non-basic belief that she’s still in the house is knowledge, at least by low standards. Relative to those same standards, it’s a justified belief, even if it happens to be false—if, for example, she has uncharacteristically jumped out a window. My cognitive home, in a broader world in which things are as they seem, would amount to knowledge, even if it doesn’t actually.

I mentioned already the idea, suggested by a brief remark in *Knowledge and Its Limits*, that *all* knowledge is basic. This idea too could be implemented into the JPK framework—the result would be the view that all of one’s knowledge is part of one’s cognitive home. Since JPK has it that for a belief to be justified, one’s cognitive home must be consistent with one’s having knowledge, the result, on this view, is that a belief is justified just in case it is knowledge. So J=K theorists needn’t object to my framework. On the contrary, my

³⁶Ichikawa (2017, ch. 4)

framework explains why $J=K$ fits naturally with the idea that all knowledge is foundational. One's approach to basic knowledge may dictate one's approach to justification.

8 Assertion

Given the approach to basic knowledge and the kind of contextualism I've signed up for, how should one think about the knowledge norm of assertion? Keith DeRose (2002) argued that the latter, combined with the apparent fact that how well-positioned one must be in order properly to make an assertion can vary according to one's conversational context, provides an argument for contextualism. But this argument is fallacious—contextualism is not the view that whether one knows depends on one's conversational situation; it's the view that knowledge ascriptions vary in content according to the speaker's conversational context.³⁷

A strengthening of the knowledge norm of assertion into a basic knowledge norm of assertion seems rather clearly to be a non-starter, at least if we assume that much of our knowledge is non-basic. It implies that we violate the assertion norm any time we say something we know indirectly. For example, I've made various assertions in this paper about which philosophers have offered which arguments; in at least many of those cases, although I didn't say anything I didn't know, I almost certainly said things I didn't know *basically*. For example, I don't know *directly* that Jonathan Sutton defends the view that $J=K$. I think for example that my knowledge of his views depends on my knowledge of what his book says.³⁸ A basic knowledge norm is too demanding.

Nevertheless, I think there is a way to use the idea that assertability depends on one's own conversational context, with an emphasis on the role of basic knowledge, to articulate an interesting and plausible competitor to the more familiar knowledge norm of assertion. Such is the project of this section.

My starting point is a Stalnakerian approach to conversational contexts. According to this approach, contexts are modelled as sets of possible worlds—one can think of them as the worlds compatible with the presuppositions of the context. A world's inclusion in the context signifies that it is treated as an open possibility in the context. Assertions have the typical conversational effect of shrinking the context—if someone asserts that p , and the assertion is accepted into the conversation, conversational participants will go on to presuppose that p —they will operate with a revised context characterised by the subset of the original worlds in which p .

Call the conversational *increment* of an assertion its intended marginal contribution. When one asserts p , that assertion aims to reduce the context by removing those not- p worlds consistent with it. The increment associated with a given assertoric content will vary from context to context. In a conversation in which less of relevance is being presupposed, the increment will generally be larger than it will in conversations that presuppose more. My idea is to consider this incremental contribution, rather than the propositional content of the assertion itself, to be the locus for the assertion norm. One shouldn't modify the context in a way that goes beyond one's epistemic position.

³⁷Blackson (2004) makes this complaint against DeRose's argument. DeRose (2009) recognises this issue and modifies the argument for contextualism.

³⁸I've read his book. I don't know whether to think that as a result I know directly what his book says. Does that knowledge depend on my knowledge of what my copy of the book says? Similarly, I'm not sure whether to think I ever have direct knowledge of who has what views, in cases when I do not rely on having read a book. I know some of Timothy Williamson's views because I've seen him state them. This may or may not suffice for basic knowledge. Many details remain unresolved.

Suppose I was at dinner with Sarah, who had the lasagna. Now I'm in a conversation about that evening, and someone says: 'Allison had the lasagna'. (Prior to the conversation I had no idea what Allison had.) We all update accordingly, and now we're in a context that presupposes that Allison had the lasagna. No worlds in which Allison failed to have the lasagna are compatible with the context. Now I say: 'Sarah had the same thing Allison did.' My assertion is also taken up, and so the context is updated by removing worlds in which Sarah and Allison had different meals. Since only worlds where Allison had lasagna were there prior to my assertion, this amounts to the same thing as removing worlds in which Sarah didn't have the lasagna. The thought behind my *incremental norm of assertion* is that I must stand in the proper epistemic relationship to that increment. What matters, for whether my assertion is permitted, is my epistemic relationship to those worlds—the worlds where Sarah didn't have the lasagna.

It's an implication of this framework that the propositional content of the sentence I uttered—that Sarah and Allison ate the same thing—needn't necessarily be known for my assertion to be a proper one. If my conversational partner is mistaken, or lying, about Allison, my proper assertion might even be false. Nevertheless, it *updates* the conversation in a way accurately reflective of my epistemic position.

What epistemic state must I hold towards the increment? Elimination by basic knowledge seems a strong candidate. If my conversational contribution is to eliminate from the context a certain set of worlds, my basic knowledge had better rule out those worlds. One might worry that this requirement is too restrictive, for the same reasons as those I discussed in the context of a basic knowledge norm of assertion above. But, since most conversational contexts presuppose that many skeptical possibilities don't obtain, the increment will often be ruled out by a speaker's basic evidence. For example, even though I don't have basic knowledge that Jonathan Sutton defends $J=K$, I do plausibly have basic knowledge that, *combined with the presuppositions of ordinary conversations*, entails that he does. Suppose I have basic knowledge that his book includes a defence of $J=K$.³⁹ (His book is in front of me right now.) If I'm speaking in a conversational context that includes the presupposition that his book describes his view, then by asserting 'Sutton defends $J=K$ ', I eliminate from the context only worlds in which his book expresses his view, and doesn't defend $J=K$. And these worlds, my basic knowledge rules out.

This approach differs from the more familiar knowledge norm in key ways. One is that it encodes much more direct sensitivity to the cooperative, social goals of communication. What I am permitted by the norm to assert doesn't just depend on my epistemic position; it depends on what kind of context I'm in. It also classifies as licensed some assertions that would violate the knowledge norm. This will happen in conversational contexts that are themselves inconsistent with the truth. While I don't think it's obvious what to say about such cases, I consider the result a welcome one. I think there is something intuitively proper about the kinds of assertions in question. This does *not* seem to me a mere matter of excusable norm violation. The speaker has updated in a context in a way that is genuinely reflective of their basic knowledge.

Notice that our practice of *reporting* assertions is very flexible. We tend to feel free to choose among a wide variety of non-equivalent descriptions of people's utterances, when we're not trying to quote them. If I say 'Sarah had the same thing Allison did' under the

³⁹This is a simplifying assumption that may not be true; perhaps my basic knowledge is only that this book with his name on it defends the view. The point still goes through, since we typically speak in conversational contexts that presuppose that academic monographs are written by the people named on their covers.

circumstances described above, I think people on the whole would feel free to describe my utterance in any of the following ways:

- Jonathan said that Sarah had the same thing Allison did.
- Jonathan said that Sarah had lasagna.
- Jonathan said that someone else had lasagna too.
- Jonathan said that Sarah and Allison both had lasagna.
- Jonathan said that Sarah was at this restaurant.

All of these are accurate ways of characterising my incremental impact on the conversational context. This further reinforces the thought that this dynamic influence is a primary locus of normative evaluation. I think there are advantages to focusing one’s assertion norms here, instead of on the contents literally asserted.

9 Conclusion

I do not pretend to have made anything like a definitive case for a basic knowledge first approach to epistemology. I’ve merely sketched some ways in which basic knowledge might be well-suited to playing interesting epistemic roles. At a minimum, I hope I’ve made the case that theorising in terms of basic knowledge is a promising strategy, and one that ought to be considered alongside various knowledge-emphasising approaches.

In addition to the questions about the applications of basic knowledge to theorising in general, many questions also remain about the nature and extent of basic knowledge. As I’ve suggested above, I am sympathetic to the view that at least often, when one perceives that p , one has basic knowledge that p . One important question is, under what circumstances does this occur? Not everyone has the ability to perceive directly whether a Bateleur eagle is present, but some experts may. What distinguishes such experts from people who *can* know via sight that something is a Bateleur, but who can do so only by inferring this fact from its superficial physical characteristics? Can there be basic knowledge of controversial matters, such as whether a comment is racist? These are important questions—doubly so if I am right that the category of basic knowledge is an important one.

On the whole, I hope I’ve made at least a reasonably plausible case in favour of the idea basic knowledge should be recognised as a distinctive category of knowledge, and that in at least some instances, it is a better candidate than knowledge itself for epistemological theorising. So I am inclined to treat it is rather metaphysically fundamental—perhaps even more so than knowledge itself.

One might interpret this as an argument against the knowledge first program—certainly the name ‘basic knowledge first’ suggests that knowledge is not first. If one assumes that the idea that knowledge is first commits one to the idea that knowledge is a ‘black box’, to be used in theorising less basic phenomena, but never to be peeked into and illuminated itself, what I have said is inconsistent with the knowledge first stance. As I indicated in §1 however, I do not subscribe to this assumption; I think the sense in which knowledge is ‘first’ is a relative one. (Perhaps a more accurate, if less catchy, name for it would have been ‘knowledge *earlier* epistemology’.) I don’t think, for instance, that the knowledge first stance commits one to denying a reduction of the psychological to the biological or physical.

But as I indicated above, I consider this a terminological matter about how exactly we should use the phrase ‘knowledge first’.⁴⁰

Bibliography

- Anscombe, G. E. M. 1957. *Intention*. Harvard University Press.
- Audi, Robert. 2006. “Testimony, Credulity, and Veracity.” In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*, 25–49. Oxford University Press.
- Bird, Alexander. 2005. “Abductive Knowledge and Holmesian Inference.” In Tamar Szabo Gendler and John Hawthorne (eds.), *Oxford Studies in Epistemology*, volume 1, 1–31. Oxford University Press.
- Blackson, Thomas A. 2004. “An Invalid Argument for Contextualism.” *Philosophy and Phenomenological Research* 68:344–345.
- Cleve, James Van. 2006. “Reid on the Credit of Human Testimony.” In Jennifer Lackey and Ernest Sosa (eds.), *The Epistemology of Testimony*. Oxford University Press.
- Danto, Arthur C. 1979. “Basic Actions and Basic Concepts.” *Review of Metaphysics* 32:471–485.
- Davis, Wayne A. 2007. “Knowledge Claims and Context: Loose Use.” *Philosophical Studies* 132:395–438.
- . 2017. “Loose Use and Belief Variation.” In Jonathan Jenkins Ichikawa (ed.), *Routledge Handbook of Epistemic Contextualism*. Taylor & Francis.
- DeRose, Keith. 2002. “Assertion, Knowledge, and Context.” *The Philosophical Review* 111:167–203.
- . 2009. *The Case for Contextualism*. Oxford University Press.
- Dodd, Dylan. 2007. “Why Williamson Should Be a Sceptic.” *Philosophical Quarterly* 57:635–649.
- Ford, Anton. 2011. “Action and Generality.” In Anton Ford, Jennifer Hornsby, and Frederick Stoutland (eds.), *Essays on Anscombe’s Intention*. Harvard University Press.
- Hornsby, Jennifer. 1980. *Actions*. Routledge and Kegan Paul.
- Ichikawa, Jonathan Jenkins. 2013. “Basic Knowledge and Contextualist “E = K”.” *Thought: A Journal of Philosophy* 2:282–292.
- . 2014. “Justification is Potential Knowledge.” *Canadian Journal of Philosophy* 44:184–206.

⁴⁰I presented versions of this paper to audiences at an Epistemology workshop at the University of Helsinki in 2015 and the Episteme Conference in South Africa in 2016. Thanks to audiences at both events for helpful feedback and discussion. Among others, I benefitted from conversations with Jessica Brown, Branden Fitelson, Michael Hannon, Jennifer Lackey, Maria Lasonen-Aarnio, Nick Leonard, Lauren Leydon-Hardy, Alan Millar, Lisa Miracchi, Ram Neta, Miriam Schoenfield, Susanna Schellenberg, Baron Reed, Blake Roeber, David Sosa, Timothy Williamson, and Sarah Wright. I am especially grateful to Kathryn Pogin, who provided comments at the Episteme conference and gave detailed written feedback.

- . 2017. *Contextualising Knowledge: Epistemology and Semantics*. Oxford University Press.
- Ichikawa, Jonathan Jenkins and Jenkins, C. S. I. 2016. “On Putting Knowledge ‘First.’” In J. Adam Carter, Emma Gordon, and Benjamin W. Jarvis (eds.), *Knowledge First*. Oxford University Press. In Press. Draft online (2 June, 2016) from <http://jichikawa.net/research/>.
- Koslicki, Kathrin. 2012. “Varieties of Ontological Dependence.” In Fabrice Correia and Benjamin Schnieder (eds.), *Metaphysical Grounding: Understanding the Structure of Reality*, 186. Cambridge University Press.
- Lavin, Douglas. 2013. “Must There Be Basic Action?” *Noûs* 47:273–301.
- Lewis, David K. 1996. “Elusive Knowledge.” *Australasian Journal of Philosophy* 74:549–567.
- Littlejohn, Clayton. 2012. *Justification and the Truth-Connection*. Cambridge University Press.
- . 2013. “No Evidence is False.” *Acta Analytica* 28:145–159.
- . 2014. “The Unity of Reason.” In Clayton Littlejohn and John Turri (eds.), *Epistemic Norms: New Essays on Action, Belief and Assertion*, 135–154. Oxford University Press.
- Marušić, Berislav. 2011. “The Ethics of Belief.” *Philosophy Compass* 6:33–43.
- McCann, Hugh J. 1974. “Volition and Basic Action.” *Philosophical Review* 83:451–473.
- McKinnon, Rachel. 2015. *The Norms of Assertion: Truth, Lies, and Warrant*. Palgrave Macmillan.
- Miracchi, Lisa. 2015a. “Competence to Know.” *Philosophical Studies* 172:29–56.
- . 2015b. “Knowledge Is All You Need.” *Philosophical Issues* 25:353–378.
- Nagel, Jennifer. 2013. “Knowledge as a Mental State.” *Oxford Studies in Epistemology* 4:275–310.
- Schwitzgebel, Eric. 2011. *Perplexities of Consciousness*. The MIT Press.
- Sider, Theodore. 2011. *Writing the Book of the World*. Oxford University Press.
- Smith, Martin. 2010. “What Else Justification Could Be.” *Noûs* 44:10–31.
- Sneddon, Andrew. 2001. “Does Philosophy of Action Rest on a Mistake?” *Metaphilosophy* 32:502–522.
- Sosa, Ernest. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University Press.
- Sutton, Jonathan. 2007. *Without Justification*. The MIT Press.
- Weatherson, Brian. 2009. “E ≠ K.” Online (2 June, 2016) from <http://brian.weatherson.org/EK.pdf>.

- Weil, Vivian M. and Thalberg, Irving. 1974. "The Elements of Basic Action." *Philosophia* 4:111–138.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press. ISBN 0198250436 (alk. paper).
- . 2016a. "Acting on Knowledge." In *Knowledge First Approaches in Epistemology and Mind*. Oxford University Press. In Press. Draft online (1 June, 2016 from http://http://www.philosophy.ox.ac.uk/members/philosophy_panel/tim_williamson).
- . 2016b. "Justifications, Excuses, and Sceptical Scenarios." In Julien Dutant and Fabian Dorsch (eds.), *The New Evil Demon*. Oxford University Press. In Press. Draft online (2 June, 2016) from http://www.philosophy.ox.ac.uk/members/philosophy_panel/tim_williamson.
- Wittgenstein, Ludwig. 1998. *Philosophical Investigations*. Wiley-Blackwell.