

No Need for Excuses: Against Knowledge-First Epistemology and the Knowledge Norm of Assertion^{*}

Joshua Schechter
Brown University

1. Introduction

Since the publication of Timothy Williamson's *Knowledge and its Limits*, knowledge-first epistemology has become increasingly influential within epistemology.¹ This paper discusses the viability of the knowledge-first program. It has two main parts. In the first part, I briefly present knowledge-first epistemology (Section 2) as well as several big picture reasons for concern about this program (Section 3). While these considerations are pressing, I concede that they are not conclusive. To determine the viability of knowledge-first epistemology will require philosophers to carefully evaluate the individual theses endorsed by knowledge-first epistemologists as well as to compare it with alternative packages of views. In the second part of the paper, I contribute to this evaluation by considering a specific thesis endorsed by many knowledge-first epistemologists – the knowledge norm of assertion. According to this norm, roughly speaking, one should assert that *p* only if one knows that *p*. I present and motivate this thesis (Section 4). I then turn to a familiar concern with the norm: In many cases, it is intuitively appropriate for someone who has a strongly justified belief that *p*, but who doesn't know that *p*, to assert that *p*. Proponents of the knowledge norm of assertion typically explain away our judgments about such cases by arguing that the relevant assertion is improper but that the subject has an excuse and is therefore not blameworthy for making the assertion (Section 5). I argue that that this response

^{*} This is the penultimate draft of a paper in J. Adam Carter, Emma Gordon, and Benjamin Jarvis (eds.) (2017) *Knowledge-First: Approaches in Epistemology and Mind*, Oxford: Oxford University Press, pp. 132–159.

¹ Proponents of knowledge-first epistemology include Williamson (2000), Hawthorne and Stanley (2008), and Littlejohn (2013). See McGlynn (2014) for a comprehensive critical discussion. My discussion of knowledge-first epistemology is heavily indebted to Williamson (2000).

will not work (Section 6). In many of the problem cases, it is not merely that the subject's assertion is blameless. Rather, the subject positively ought to make the assertion. Appealing to an excuse cannot be used to adequately explain this fact. (Nor can we explain this fact by appealing to some other, quite different, consideration.) Finally, I conclude by briefly considering whether we should replace the knowledge norm of assertion with an alternative norm (Section 7). I argue that the most plausible view is that there is no norm specifically tied to assertion.

2. Knowledge-First Epistemology

Knowledge-first epistemology was originally developed as a reaction to the post-Gettier attempts to provide a conceptual analysis of the concept of knowledge. Gettier (1963) famously argued against analyzing knowledge as justified true belief. In doing so, he provided cases in which a subject had a justified true belief that did not count as knowledge.² Gettier's paper led to a long series of papers proposing new analyses of knowledge – the no false lemmas view, the no undefeated defeater view, casual theories, sensitivity theories, and many others.³ Each of these proposed analyses were subject to further counterexamples. As time went on, proposed analyses grew increasingly baroque and ad hoc. As Williamson (2013, p. 2) writes, the program of providing an analysis of knowledge displayed the “signs of a degenerating research program”.

The repeated failure of the post-Gettier literature to provide a successful analysis of the concept of knowledge suggested that we should perhaps take the concept of knowledge to be unanalyzable. Indeed, since the force of the counterexamples did not turn on whether the

² The dialectic in Gettier's paper is more complicated than the typical presentation suggests. Gettier presents an abstract argument to the conclusion that there can be a justified true belief that does not count as knowledge. (One can competently deduce a justified true belief from a justified false belief. Such a belief won't count as knowledge.) The role of Gettier's original cases was to illustrate this general argument. The subsequent literature, however, became explicitly case-based.

³ See Clark (1963), Lehrer and Paxson (1969), Goldman (1967), and Nozick (1981), respectively. See Shope (1983) for an exhaustive presentation of the literature up to the book's publication.

proposed accounts were intended to provide conceptual analyses of the concept of knowledge or metaphysical reductions of the state of knowing,⁴ there was also some impetus in favor of taking the state of knowing to be irreducible (or, at least, to not be reducible along the lines of anything like the “traditional” justified true belief account).

Once it is claimed that knowledge is unanalyzable and irreducible (in anything like the traditional way), we have a new primitive to deploy. We can try to put this primitive to work in providing conceptual analyses and metaphysical reductions. For instance, instead of analyzing the concept of knowledge in terms of the concept of belief and the concept of justification, one might try to analyze the concept of belief and the concept of justification in terms of the concept of knowledge. Alternatively, or in addition, one might try to reduce the cognitive state of believing and the epistemic status of being justified to the state of knowing. One might try to go further and relate knowledge to other epistemic notions, such as inquiry and evidence. One might go still further and relate knowledge to non-epistemic notions, such as deliberation, assertion, and action.

Knowledge-first epistemology is the result of going down this intellectual path. Proponents of knowledge-first epistemology endorse the claim that knowledge is unanalyzable and irreducible (in anything like the traditional way).⁵ But they go much further than this. They suggest that epistemology (and perhaps the philosophy of mind) should be reoriented with knowledge at its base. Much of our mental lives – our cognitive states and their epistemic statuses – and our behavior – actions, assertions, and their assessment – should be understood in

⁴ My use of the term “state” is not meant to commit me to what Williamson (2000) means when he calls knowledge a mental state. Every claim I make about the state of knowing could be rephrased as a claim about the knowledge relation.

⁵ It is compatible with knowledge-first epistemology to endorse positive claims about the nature of knowledge. For instance, Williamson (2000) endorses a “safety from error” requirement: If one knows that p, one could not have easily been wrong in a similar case. This is not meant as a conceptual analysis or a reduction of knowledge, since according to Williamson, “similar case” should itself ultimately be understood in terms of knowledge.

terms of knowledge. How exactly this should go varies from theorist to theorist. But what doesn't vary among knowledge-first epistemologists is the endorsement of the claim that knowledge is of fundamental importance to the understanding of our mental lives – it is an “unexplained explainer”.⁶

Knowledge-first epistemology, then, shouldn't be understood as a single thesis or package of theses. Rather it is what we might call a “program”. To be a knowledge-first epistemologist requires endorsing the claims that the concept of knowledge is unanalyzable and that the state of knowing is irreducible (in anything like the traditional way). But it also requires taking a certain stance – treating knowledge as being of first importance, and trying to put the concept of knowledge and the state of knowing to use in doing significant philosophical work.⁷

There are many specific theses that are endorsed by individual knowledge-first epistemologists. They concern the concept of knowledge and its relation to other concepts; the metaphysics of the state of knowing and its relation to other cognitive states and epistemic statuses; normative connections between knowledge, assertion, and action; the purpose of cognition and communication; and much else besides. Being a knowledge-first epistemologist does not require that one endorse any specific collection of these theses. But knowledge-first epistemologists will typically endorse many such theses. In particular, knowledge-first epistemologists will typically endorse many of the following theses (or variants thereof):⁸

⁶ Williamson (2000), p. 10.

⁷ In principle, these two components of the knowledge-first program are separable. One could claim that the concept of knowledge is unanalyzable and the state of knowing is irreducible without also claiming that the concept of knowledge and the state of knowing can be used to do significant philosophical work, and vice versa.

⁸ Many of these theses can be found – explicitly or implicitly – in Williamson (2000). See McGlynn (2014) for additional citations.

Concepts

- The concept of belief (and the concepts of other cognitive states) can be analyzed in terms of the concept of knowledge. For instance, perhaps believing that p can be analyzed as treating p as if one knows it.
- The concept of justification (and the concepts of other epistemic statuses) can be analyzed in terms of the concept of knowledge. For instance, perhaps having a justified belief that p can be analyzed as either knowing that p or having a belief that p that would have been knowledge had the world cooperated.

Psychology

- Normal human subjects acquire the concept of knowledge before they acquire the concept of belief or the concept of justification.

Metaphysics

- The knowledge relation is a metaphysically privileged relation. In the terminology of Lewis (1983), it is “perfectly natural”. Alternatively, the knowledge relation is more natural than any other cognitive relation (such as the belief relation) or epistemic property (such as being justified).
- Belief and other cognitive states can be characterized in terms of the state of knowing.
- Justification and other epistemic statuses can be characterized in terms of the state of knowing.
- The characterization of belief in terms of knowledge is part of the essence of belief. It provides the real definition of belief. (And similarly for other cognitive states.)
- The characterization of justification in terms of knowledge is part of the essence of justification. It provides the real definition of justification. (And similarly for other epistemic statuses.)
- Seeing that p, remembering that p, and understanding that p each entail knowing that p. More generally, knowing is the most general propositional attitude that is factive and that constitutes a mental state.

Behavior

- In predications, explanations, and rationalizations of a subject’s behavior, one should appeal to the subject’s knowledge and not merely to the subject’s beliefs.

Connections within Epistemology

- A subject’s total evidence is all and only the propositions that the subject knows.
- Knowledge terminates inquiry. Knowing that p (or knowing that it is not the case that p) closes the question of whether p.

Normative Connections

- Knowledge norm of assertion: One ought to assert that p only if one knows that p.
- Knowledge norm of reasoning: One ought to use the proposition that p as a premise in practical or theoretical reasoning only if one knows that p.
- Knowledge norm of belief: One ought to believe that p only if one knows that p.
- Knowledge norm of action: One ought to act as if p only if one knows that p.

- The knowledge norm of assertion uniquely characterizes assertion from among the speech acts. (And similarly for the other normative connections.)
- The real definition of assertion is that it is the speech act that obeys the knowledge norm. (And similarly for the other normative connections.)

Teleology

- The central purpose of cognition is to generate knowledge.⁹
- The central purpose of communication is to transmit and pool knowledge.¹⁰
- Belief aims at knowledge.

Meaning and Reference

- What one knows (or is in a position to know) in part determines the reference of one's linguistic and mental representations.

Philosophical Methodology

- Epistemology should primarily focus on issues concerning knowledge rather than justified belief or rational credence.
- The philosophy of mind should primarily focus on issues concerning knowledge rather than belief.

What motivates these theses (or variants thereof)? Many of these theses can be motivated individually. For instance, since it will be the focus of the second half of this paper, consider the knowledge norm of assertion. According to this thesis, one ought to assert that p only if one knows that p. This thesis has been argued for in the following ways:¹¹

- (i) When someone makes an assertion, it can be questioned by asking, "How do you know that?" An assertion can be rejected by saying, "You don't know that." Requests for answers to yes/no questions can be appropriately responded to with "I don't know." And so forth. The best explanation of this body of data is provided by the knowledge norm.
- (ii) There is something inappropriate about asserting a Moorean paradoxical sentence of the form "p but I don't know that p." The best explanation of this is provided by the knowledge norm.
- (iii) If one has merely statistical grounds for believing that someone who purchased a ticket in a lottery is highly unlikely to win, it is inappropriate to tell them, "You're not going to win". The best explanation of this is provided by the knowledge norm combined with the claim that subjects cannot know lottery propositions on merely statistical grounds.¹²

⁹ There are other important purposes of cognition – e.g., planning and deliberation.

¹⁰ There are other important purposes of communication – e.g., coordination with others.

¹¹ See Williamson (2000), which develops considerations originally put forward in Unger (1975) and Slote (1979).

¹² There have been several responses to these motivations in the literature. See, for example, Hill and Schechter (2007), Lackey (2007), Kvanvig (2009), and Whiting (2013).

Beyond the individual motivations for these theses, there is a more general motivation – the theses “hang together” nicely. Taken as a whole, they provide an attractively tidy theoretical account of mental life, with knowledge at its center. A central purpose of cognition is the generation of knowledge. When all goes well, our beliefs count as knowledge. In our theoretical and practical reasoning, we reason from premises that we know. In communication, we exchange knowledge with others. We act on what we know. And so on. That’s what happens when things are going (very) well. Of course, things don’t always go so well. Sometimes we have “mere beliefs” – beliefs that don’t count as knowledge. Sometimes our mere beliefs are justified, sometimes not. We can explain what mere beliefs are and what justification is in terms of knowledge. We can explain what it is for a course of reasoning or an assertion to be improper in terms of knowledge. And so on. The simplicity and elegance of this picture provides us with reason to endorse it.

3. ...and its Limits

As attractive as this picture is, there are reasons for concern. In this section, I present three big-picture worries for knowledge-first epistemology. While they are not conclusive, they do provide reasons to be uneasy about the program.

The first worry stems from the post-Gettier literature. Proponents of knowledge-first epistemology argue that the post-Gettier literature shows that the concept of knowledge is unanalyzable and that the state of knowing is irreducible. I’m sympathetic to this view. But they also argue for the central importance of knowledge in our mental lives. My view is that the post-Gettier literature illustrates that this is a mistake.¹³

¹³ Kaplan (1985) and Kvanvig (2003) also argue that a moral of the post-Gettier literature is that knowledge is not an important epistemic status.

Consider some Gettier-style case in which a subject has a justified true belief that is not knowledge. For instance, here is a version of an example originally due to Russell (1912): Suppose that Jones wonders what time it is and briefly glances at his trusty wall clock. Suppose that the clock reads “2:05” and so Jones comes to believe that it is 2:05 pm. Suppose that it really is 2:05 pm. However, suppose that, unbeknownst to Jones, the clock stopped exactly 12 hours ago. In this case, Jones has a true belief that it is 2:05 pm. This belief is justified, since it is based on good evidence that the time is 2:05 pm. But, intuitively, Jones’s belief does not count as knowledge.

Thinking about this case, it seems clear that Jones is in a very good epistemic position. His belief about the time is true. His belief fits the evidence that’s available to him. His belief was formed on the basis of the evidence. If he acts on his belief about the time – rushing off because he has a 2:15 pm appointment, say – he’ll be acting in a way that is appropriate given the facts. He’ll be acting in a way that makes sense given his evidence. If he’s asked what time it is and he asserts that it’s 2:05 pm, he’ll be asserting a truth. He’ll be asserting something that is reasonable for him to assert given his evidence. And so on. In every sense that seems important, Jones is doing quite well.

To be sure, Jones lacks knowledge. He doesn’t know what time it is. Put baldly like that, it can sound like Jones is not in a great epistemic position, after all. But this strikes me as an illusion due to how we talk – when we say things like “S doesn’t know what time it is”, we’re typically trying to communicate that S is unsure, getting it wrong, or otherwise in some highly

problematic state with regard to his view of the time.¹⁴ Focusing on the fact that Jones's belief is both true and justified, however, it doesn't seem like Jones is missing anything important.

What this suggests is that knowledge is actually a rather marginal epistemic status. It is important to have true beliefs so that we correctly represent the world. It is important to respect the evidence and, more generally, to have justified beliefs (and rational credences). But it is difficult for me to see that, in addition, it is important to have knowledge, too.

There are features of Jones's belief that might seem problematic. It is lucky that Jones's belief turned out to be true. Jones's belief about the time is bound up with his false belief that the clock is in good working order. Jones's belief was formed via an unreliable process – glancing at a stopped clock.¹⁵ Jones's belief is not safe in the sense that Jones could easily have come to have a false belief about the time. And so on.

Nevertheless, these features of Jones's belief do not make it any more problematic than some genuine cases of knowledge. There are cases of knowledge that have similar features. For instance, someone can win a *History of the Dodo* volume in a raffle and by reading it come to believe that the first recorded mention of dodo birds was in 1598. This belief counts as knowledge despite the element of luck involved. Someone can miscount the number of chairs in a room as 11 (rather than 12), and then infer that the number of chairs in the room is less than 30. The inferred belief counts as knowledge despite the fact that it is bound up with a false belief about the exact number of chairs present. There are differences between the kind of luck involved in Jones's belief and the kind of luck involved in winning a raffle – it was lucky that Jones's belief was true whereas it was lucky that the raffle entrant came to have a belief about

¹⁴ See Weiner (2009, p. 166) for a related view. Weiner argues that the concept of knowledge is a “Swiss Army” concept in the sense that knowledge is a concept “that is not important in itself, but that provides an economical way of summing up several other concepts that are important in themselves.”

¹⁵ Jones's belief was also formed via a *reliable* process – glancing at a clock that one has no reason to think is not working.

dodos at all.¹⁶ There are also differences between the way in which a false belief is bound up with Jones's belief about the time and the way in which a false belief is bound up with the chair counter's belief about there are fewer than 30 chairs. But it is difficult to see how anything of serious epistemic significance could depend on these subtle differences.

The point is not specific to Russell's clock case. The same holds true for other Gettier-style cases. In all of these cases, when we focus on the details of the case, the fact that the subject doesn't count as knowing is of little intuitive significance. We're used to talking in terms of knowledge and treating knowledge as the most significant epistemic status. But this may be more an accident of human language and psychology than a guide to the epistemological joints of nature.

The second worry about knowledge-first epistemology can be raised by asking the following questions: What could the state of knowing *be* such that it does everything that proponents of knowledge-first epistemology suggest it does? How could there be a single kind of state that plays all (or many) of the roles listed above? Proponents of knowledge-first epistemology argue that the state of knowing is irreducible, and so cannot answer these questions by providing a reductive account of knowing. But that does not mean that the questions are illegitimate. Proponents of knowledge-first epistemology make the striking claim that there is a single kind of state that plays very many different roles in our mental life. In supporting this claim, they owe us some understanding of how this could be so. What is needed, then, is not a reductive account, but some kind of picture that explains what knowledge is such that it can be the unexplained explainer at the heart of our mental lives.

¹⁶ Pritchard (2005) distinguishes between five ways in which a belief may be epistemically lucky and argues that three of them are compatible with knowledge and two are not.

There are pictures that could explain the centrality of knowledge to our mental lives. If one of these pictures were attractive, that would buttress the case of knowledge-first epistemologists. But if there were no attractive picture, that would seem to cast doubt on the program of knowledge-first epistemology. The trouble, then, is this: Every picture that I can envision that would explain the centrality of knowledge is grossly implausible.

Consider a broadly Cartesian picture of knowledge according to which knowledge entails infallibility, indefeasibility, and rational certainty.¹⁷ Suppose, too, that we could always tell when we had such knowledge. If it were possible for us to have a wide range of such Cartesian knowledge, one could make sense of the centrality of knowledge in our mental lives.¹⁸ It would be extremely important to generate such knowledge, communicate such knowledge, act on such knowledge, and so forth.

Of course, this picture of knowledge is highly implausible.¹⁹ If knowledge were required to conform to these constraints, we would have precious little, if any, knowledge. Perhaps some of our beliefs in simple logical, mathematical, and conceptual truths are infallible, indefeasible, and rationally certain. I suspect not, given that there can be rational disagreements even over basic logical principles. But even if we could have *some* knowledge that satisfied the Cartesian constraints, none of our everyday knowledge about ourselves or the world would count.

Alternatively, consider a picture of knowledge according to which knowledge entails being in some intimate metaphysical relation with the relevant fact. Such a view could be developed by analogy to a disjunctivist view of perception according to which perception provides subjects with a distinctive kind of contact with objects in the world, at least in the “good

¹⁷ Descartes himself did not impose such strong requirements on everyday knowledge.

¹⁸ See BonJour (2011).

¹⁹ Williamson (2000) agrees. His anti-luminosity argument is targeted in part against such a Cartesian view.

case”.²⁰ On the analogous picture of knowledge, in the “good case” of believing – that is, knowing – subjects have a distinctive kind of metaphysical contact with the fact that p. Such a picture of knowledge might be able to explain the central importance of knowledge to our mental life, at least if we could explain what “contact” comes to and why having contact of this sort is itself so important. (We’d also need to explain why having a true belief or a justified true belief doesn’t entail having contact with the relevant fact.) Yet, this picture is implausible, too. Consider some of the many different ways in which one can gain knowledge. For instance, one can gain knowledge by testimony. If knowledge entails a distinctive kind of contact with the facts, the picture presumably would be that when a hearer gains knowledge by the testimony of someone who knows, the speaker’s contact with the relevant fact is somehow transmitted to the hearer along with the belief in the relevant proposition. This seems incredible. Similarly, one can gain knowledge by inductive reasoning. Again, it is difficult to believe that an inductive inference can yield some kind of intimate metaphysical relation to the truth of the conclusion, even given an intimate metaphysical relation to the truth of the premises.²¹ How could it be that being in contact (whatever that is) with assorted facts about the nearby environment can put us in contact with facts about distant locations? How could it be that being in contact with assorted facts about the present and recent past can put us in contact with facts about the distant future?²²

Given that these pictures of knowledge are implausible, it is difficult to see how else to explain the multifarious roles that knowledge plays in our mental lives according to knowledge-

²⁰ See Hinton (1967), Snowdon (1980), McDowell (1982), and Martin (2002). See Pritchard (2012) and the papers in Byrne and Logue (2008) for recent discussions of disjunctivism.

²¹ The difficulty here is perhaps most acute for an inductive inference to a conclusion about a particular case that does not go via some inductive generalization – for example, an inductive inference to the conclusion that the next marble will be red that does not go via the intermediate conclusion that all marbles are red.

²² There are metaphysical relations that are transmitted via testimony – for instance, standing in a causal relation with the fact that makes the belief true. The point is that there is no good candidate for such a relation that (a) can be transmitted via testimony and inductive inference, (b) is present in all cases of knowledge, and (c) explains the centrality and importance of knowledge. Thanks to Louis deRossett for pressing me on this issue.

first epistemology. This is not to say that knowledge-first epistemologists don't have arguments for their claims – they have provided a battery of arguments in support of their individual theses, and there is a nice theoretical unity to the combination of their views. But in the absence of an attractive picture that illuminates how it could be that there is a cognitive state that can simultaneously play all the roles that knowledge is supposed to play, there is reason for concern.

The third worry about knowledge-first epistemology stems from the diversity of our cognitive lives. There are many different kinds of cognitive states beyond outright belief.²³ For instance, there are conditional beliefs. There are degrees of confidence. There are assumptions, presuppositions, hypotheses, speculations, hunches, guesses, expectations, convictions, and so on. If part of knowledge-first epistemology is to characterize all of these cognitive states in terms of knowledge, that strikes me as a tall order. Similarly, there are positive epistemic statuses beyond justification and knowledge. For instance, there is wisdom. There are various kinds of understanding.²⁴ There are specific intellectual virtues, such as open-mindedness and creativity. If part of knowledge-first epistemology is to characterize all of these epistemic statuses in terms of knowledge, that also strikes me as a tall order.

In a similar vein, in providing an account of our cognitive lives, one should not focus entirely on beliefs and other cognitive states. There are also transitions between cognitive states. These transitions are governed by rules of inference (and belief-forming methods more generally). There are important epistemological questions about such transitions and rules. For instance, one can look at a course of reasoning that yields an unjustified belief and ask where exactly the mistake was made. Was one of the subject's initial beliefs unjustified? Or did the lack

²³ I am here distinguishing cognitive states from conative states such as desiring, wishing, planning, intending, and the like.

²⁴ See Grimm (2006), Khalifa (2011), Sliwa (2014), and Kelp (2015) for proposals on how to reduce (propositional) understanding to knowledge.

of justification stem from an epistemic mistake in one of the subsequent transitions in thought? It is difficult to see how these questions can be understood as ultimately concerning knowledge.

A knowledge-first epistemologist could, I suppose, claim that a transition between beliefs is faulty if it is not the application of a rule of inference that preserves knowledge when the world is cooperating. (This is analogous to the suggestion that a belief is justified if it is knowledge or would have been knowledge had the world cooperated.) But this proposal faces two difficulties. First, we need an account of what it is for the world to cooperate. Second, and more importantly, there is a kind of Euthyphro problem here. It is plausible that a rule of inference preserves knowledge in part *because* it doesn't involve any mistake. It is less plausible that a rule of inference doesn't involve a mistake because it preserves knowledge. This suggests that there is a more fundamental epistemic status than knowledge – namely, whether a rule of inference involves an epistemic mistake.

The diversity of our cognitive lives – involving many different kinds of cognitive states, many different epistemic statuses, and including transitions and rules of inference – suggests that the attractively tidy account provided by knowledge-first epistemologists may be a bit too tidy an account.

I should concede that these worries are not conclusive. (Indeed, I suspect that knowledge-first epistemologists won't be strongly moved by them.) There may be responses to each of them. And, even if there aren't, other considerations in favor of knowledge-first epistemology may outweigh the worries. All I would like to claim here is that these considerations provide reasons for concern about the program of knowledge-first epistemology.

To fully determine the viability of the knowledge-first approach will require philosophers to carefully evaluate the specific theses endorsed by knowledge-first epistemologists and to

determine their costs and benefits. It will also require philosophers to compare knowledge-first epistemology with alternative packages of views. In this, I agree with Williamson (2013, pp. 6–7), who writes, “In the long run, knowledge-first epistemology ... should be judged by its fruitfulness as a research program, compared to its competitors.” Indeed, I expect that any verdict on knowledge-first epistemology will only emerge over the next few decades.

In the remainder of this paper, I would like to contribute to the evaluation of knowledge-first epistemology by considering a thesis endorsed by many of its proponents (as well as by other philosophers) – the knowledge norm of assertion. According to this thesis, roughly speaking, one should assert that *p* only if one knows that *p*. In what follows, I’ll be focusing on a relatively fine-grained issue. There is a popular line of objection to the norm: In many cases, it seems appropriate for someone who has a strongly justified belief that *p*, but who doesn't know that *p*, to assert that *p*. Proponents of the knowledge norm of assertion typically explain away our judgments about such cases by arguing that the relevant assertion is improper but that the subject is not blameworthy for making the assertion because the subject possesses an excuse. In what follows, I will argue that this response to the objection doesn’t work. My discussion will not, of course, conclusively show that knowledge-first epistemology faces serious trouble. Indeed, it won’t even conclusively show that the knowledge norm of assertion is unacceptable. But it puts pressure on the knowledge norm of assertion, and derivatively on the program of knowledge-first epistemology.

4. The Knowledge Norm of Assertion

Roughly speaking, the knowledge norm of assertion is as follows: One ought not to assert that p unless one knows that p. In other words, there is something improper about asserting what one doesn't know to be true.²⁵

As stated above, there are several different motivations that have been put forward in support of this norm – for instance, it is appropriate to respond to an assertion by asking the speaker “how do you know?”, it is inappropriate to assert claims of the form “p and I don't know that p”, and it is inappropriate to assert lottery propositions believed on merely statistical grounds. To this we could add another, bigger picture, line of thought: A central purpose of communication is to transmit and pool knowledge. It is therefore important that there be a speech act tied to the transmission of knowledge. This speech act is assertion. This explains why an assertion of something that one does not know counts as improper.²⁶

The statement of the knowledge norm of assertion above is somewhat rough. There are many choice points in developing a specific thesis. Here are a few of them:

First, what is the particular normative status that appears in the knowledge norm? *Must* subjects make assertions only if they have the relevant knowledge? Or, instead, *ought* subjects to make assertions only if they have the relevant knowledge? Do they have *reason* not to make an assertion if they lack the relevant knowledge? Are they *criticizable* for asserting when they lack knowledge? Do they lack the *authority* to make an assertion when they lack knowledge? Or is

²⁵ Versions of the knowledge norm of assertion have been defended by Unger (1975), Slote (1979), Williamson (2000), DeRose (2002), Reynolds (2002), Adler (2002), Hawthorne (2004), Stanley (2005), Engel (2008), Schaffer (2008), Turri (2010), Benton (2011), and Blaauw (2012). See McGlynn (2014) for detailed discussion.

²⁶ A version of this motivation is put forward in Williamson (2000), p. 267.

some other normative status at issue?²⁷ In what follows, to use neutral terminology, I'll ask about whether an assertion is "proper" or "improper".

Second, is knowledge meant to be necessary or both necessary and sufficient for the propriety of an assertion? For simplicity, I'll assume that the knowledge norm states a necessary condition.²⁸ So far as I can tell, nothing in what follows will hang on this.

Third, does the norm directly target the propriety of the assertion or the propriety of the assertor? That is, does the norm state that there is something wrong with an assertion that is not based on knowledge or that there is something wrong with the person making the assertion? The prevailing view seems to be that the norm directly targets the propriety of the assertion rather than the assertor. (In the literature, the norm is often stated as a claim about assertors, but this seems mainly to be for presentational convenience.) I find this somewhat puzzling – the intuitions that are used to motivate the norm seem to involve judgments about the relevant assertor and not merely the assertion. But so far as I can tell, nothing in what follows will hang on this, either.

Fourth, what kind of normativity is at issue in the knowledge norm? Is it epistemic, moral, pragmatic, conventional (perhaps like the rules of etiquette or of spelling), a distinctively representational or communicative kind of normativity, or "plain vanilla" normativity (if there is such a thing)? One might think that the relevant kind of normativity must be epistemic since knowledge is, after all, an epistemic notion. But this is a mistake – one can have an epistemic requirement built into a non-epistemic norm. (Compare: an epistemic requirement can be built into a legal norm or into the rules of a game.) A better thought to have here is that the knowledge norm of assertion flows from the function or purpose of the practice of assertion. The kind of

²⁷ Williamson (2000) moves between "must", "criticizable", and "authority". It is not apparent to me how these statuses are meant to be connected.

²⁸ See Brown (2010) and Lackey (2011) for objections to the sufficiency of knowledge for proper assertion.

normativity at issue in the knowledge norm depends on what this function or purpose is. For instance, on one view, the central function of the practice of assertion is to enable the transmission and pooling of knowledge. Since this is an epistemic function, the relevant kind of normativity is epistemic. This is perhaps the most plausible view for a knowledge-first epistemologist to take. However, one could instead motivate a knowledge norm starting with the claim that the central function of assertion is to enable agents to coordinate with one another. This would support a moral or pragmatic kind of normativity. And other options are live possibilities, too.

Fifth, does the norm govern assertions by ordinary subjects or does it first-and-foremost govern assertions by idealized subjects? Here's an analogy that may help to make this contrast more salient: Probabilists typically argue that probabilistic coherence is a norm of rational credence. But this norm is usually understood as a constraint that in the first instance governs ideal reasoners. How this norm ends up bearing on ordinary reasoners turns out to be a delicate matter. What should we say about the knowledge norm of assertion? The prevailing view here seems to be that the norm applies to all assertions, whether by ordinary or idealized subjects.

Sixth, is the knowledge norm of assertion a fundamental norm? Or is it grounded in some more basic norm or norms (together with some descriptive facts)? For instance, is the knowledge norm of assertion somehow grounded in the knowledge norm of action?

Seventh, and finally, what exactly is the connection between the norm and assertion? Does the knowledge norm of assertion uniquely characterize the speech act of assertion out of all of the speech acts? Is it an analytic truth about assertion? Is it an essential feature of assertion? Does it provide the real definition of assertion? Williamson (2000) suggests that the knowledge norm is constitutive of assertion. On this view, the knowledge norm of assertion is analogous to a

rule of a game. One counts as playing a game just in case (and by virtue of the fact that) one is subject to the rules of the game. Analogously, one counts as making an assertion just in case (and by virtue of the fact that) one is subject to the knowledge norm.²⁹

Different answers to these questions generate different candidate knowledge norms of assertion. So far as I can tell, the differences between these candidate norms will largely be irrelevant to the discussion that follows. (There are exceptions: It will be important that the knowledge norm applies to ordinary subjects and not merely to idealized subjects. The distinction between the norm targeting the assertion or the assertor will also turn out to be relevant.)

What should we think about the knowledge norm of assertion, however it is developed? There is an obvious objection to the norm, one that has loomed large in the literature: In many cases, an assertion that *p* can be proper if the assertor has a strongly justified belief that *p*, even if the assertor doesn't know that *p*.³⁰

Consider a Gettier-style case, such as the case of Jones and the stopped clock. Suppose that Jones glances at the clock and is then immediately asked the time. Suppose that Jones answers that it is 2:05 pm. Intuitively, Jones's assertion is proper despite the fact that he doesn't know that it is 2:05 pm. (Moreover, Jones's assertion seems just as proper as it would have been had he genuinely known that it is 2:05 pm.) And similarly for other Gettier-style cases.

Indeed, there are many cases in which an assertion that *p* seems proper if the assertor has a strongly justified belief that *p*, even if the belief is *false*. For instance, suppose that Smith has strong evidence that the fastest way to get to Thayer Street from a certain location is to make a

²⁹ See Maitra (2011) and Johnson (2017) for objections to the analogy between the knowledge norm of assertion and the rules of a game.

³⁰ Versions of this objection appear in Douven (2006), Hill and Schechter (2007), Lackey (2007), Brown (2008), Koethe (2009), Kvanvig (2009), Gerken (2011), and McKinnon (2013). See McGlynn (2014) for discussion.

right on Benefit Street and then a left on Waterman Street. (Suppose that Smith lives in the area and knows the streets well.) When asked by someone “What’s the fastest way for me to get to Thayer from here right now?”, it is intuitively proper for Smith to respond by asserting “The fastest way for you to get to Thayer is to make a right on Benefit and then a left on Waterman.” This is so even if, unbeknownst to Smith, Benefit Street is currently closed to traffic due to an unexpected problem with a water main. The fact that Smith’s belief is false does not intuitively make his assertion improper.³¹

One can push this general line of objection still further. In many cases, an assertion that *p* can be proper if the assertor believes that *p*, even if the belief is *unjustified*.³² For example, suppose that Robinson strongly believes that there is life on Europa on the basis of wishful thinking. (In case it matters, we can suppose that Robinson does not believe that his belief that there is life on Europa is at all defective.) Suppose that Robinson is asked “Is there life somewhere other than the Earth?” Suppose that Robinson responds “Yes, there is life on Europa.” I maintain that Robinson’s assertion is proper. Since Robinson believes that there is life on Europa, it is not improper for Robinson to say so. Indeed, if Robinson has no reason to avoid answering the question or to try to mislead his interlocutor, Robinson positively ought to assert that there is life on Europa.

Of course, there is something wrong with Robinson – he has an unjustified belief. But the fact that his belief is unjustified does not entail that his assertion is improper. That’s a different matter.

One might think that there is a straightforward way to see that my claim about the propriety of Robinson’s assertion is false. In particular, one might draw an analogy between

³¹ For a similar case, see McKinnon (2013), p. 122.

³² This claim is not widely endorsed in the literature.

assertion and belief. According to this line of thought, assertion is the outer acceptance of a claim and believing is the inner acceptance of a claim.³³ Now consider a case of inference from unjustified beliefs. For instance, suppose that a thinker has unjustified beliefs in p and in $\text{if } p \text{ then } q$. Suppose that the thinker comes to believe that q by competently performing a Modus Ponens inference. The thinker's belief that q will (typically) be unjustified. By analogy, one might think, if a subject's belief that p is unjustified and hence epistemically improper, and the subject asserts that p based on this belief, the assertion will also be improper. Just as the impropriety of a belief can be inherited from the impropriety of the beliefs that it is inferred from, so too can the impropriety of an assertion be inherited from the impropriety of the belief that it is based on. Or so goes the line of thought.

This line of thought, however, is mistaken. Belief is not analogous to assertion. Believing is a state and assertion is an action. The right analogy is not between asserting that p and believing that p . It is between asserting that p and *coming to believe* that p .³⁴ And, once we see that, the case of a competent inference from unjustified premises actually tells in the other direction. If a thinker comes to believe that q by competently performing a Modus Ponens inference from unjustified premises, the output belief will be unjustified. But the forming of the belief that q – the Modus Ponens inference itself – is not improper. The inference itself does not inherit any impropriety from the impropriety of its inputs. Given that the thinker strongly believes that p and that $\text{if } p \text{ then } q$, and given that the question of whether q is salient, the thinker positively ought to draw the inference and form the belief that q . True, the thinker is in an

³³ See Adler (2002). Also see Williamson (2000), pp. 255–256, who draws the analogy between assertion and occurrent belief.

³⁴ Is there something tied to assertion to which believing is analogous? Perhaps we could say that in making an assertion, an assertor undertakes a kind of commitment. Then we could say that believing is analogous not to assertion but to assertoric commitment. I'm not sure this is correct. But even if it is, my topic concerns the norms governing assertion, not the norms governing assertoric commitment.

epistemically problematic state. But the mistake wasn't in the inference to q. The mistake was in believing that p and that if p then q to begin with.

The analogous is true for Robinson. Robinson's assertion that there is life on Europa is not improper. Robinson's belief that there is life on Europa is unjustified. But the act of asserting that there is life on Europa does not inherit any impropriety from the belief that it is based on. Given that Robinson strongly believes that there is life on Europa, and given that Robinson has no reason to avoid answering the question or to try to be misleading, Robinson positively ought to make the assertion.

The cases of Jones, Smith, and Robinson tell against the knowledge norm of assertion. If an assertion that p can be proper when the subject doesn't know that p but has a justified true belief that p, a justified false belief that p, or even an unjustified belief that p, then the knowledge norm of assertion is false. Indeed, analogous problems face other proposed knowledge norms, such as knowledge norms for reasoning, belief, and action.³⁵ So there is a serious worry here for proponents of knowledge norms.

Proponents of the knowledge norm of assertion (and other knowledge norms) are, of course, aware of this issue. There is a standard response. In the next section, I will present this response and examine some arguments for the claim that it cannot be made to work.

5. Excuses, Excuses

The intuition that an assertion can be appropriate in many cases where the asserter does not have knowledge – but instead has a justified true belief (e.g., the case of Jones) or a justified false belief (e.g., the case of Smith) – seems to be widely shared. In response to such problem

³⁵ There is no analogue of the case of Robinson that can be used against the knowledge norm of belief, but there are analogues of the cases of Jones and Smith.

cases, the standard response is to explain away our intuitions. According to this response, in a problem case, the relevant assertion is not, in fact, proper. Instead, our intuition that the assertion is appropriate tracks a different normative feature. Although the assertion violated the norm governing assertion, the assertor is blameless for this violation. This is because the assertor has an *excuse* for the violation. Our intuition about the appropriateness of the assertion tracks blamelessness, not propriety. This explains away our intuition about the case.³⁶

It is worth saying just a bit about excuses here.³⁷ Very generally, there are normative constraints on action. When the norms are complied with, the action is proper.³⁸ When the norms are violated, the action is improper. However, it can happen that the relevant norms are not complied with but the agent is not blameworthy for the violation. This is what happens when the agent has a sufficient excuse. For example, suppose that I promised to pick up my friend from the airport today. It is a norm on behavior that one must do what one has promised to do. If my car unexpectedly (and through no fault of my own) breaks down while I am driving to the airport and I'm unable to pick up my friend, I have violated the norm. But I am blameless in violating the norm – through no fault of my own, it was impossible for me to comply with it. That my car broke down excuses my violation.³⁹

³⁶ See Williamson (2000), pp. 256-7, DeRose (2002), and Hawthorne and Stanley (2008). (DeRose uses a somewhat different terminology – instead of talking about excuses, he talks about “secondary propriety”. I think this is merely a terminological difference.) See Sutton (2007), p. 80, for an analogous response to challenges to the claim that knowledge closes inquiry. See Littlejohn (forthcoming) and Williamson (forthcoming) for detailed presentations of the excuse response as targeted against the Cohen and Lehrer (1983) and Cohen (1984) new evil demon problem.

³⁷ See Austin (1956).

³⁸ More precisely, when the moral norms are complied with, the action is morally proper. When the prudential norms are complied with, the action is prudentially proper. And so forth.

³⁹ There are two other categories of blameless norm violation that should be distinguished from having an excuse. First, the violation of a norm can be *justified* if the agent had sufficiently good reason for violating it. For instance, I might have strong moral grounds not to pick up my friend from the airport, such as if there was an emergency that required me to take someone else to the hospital. (This notion of justification should not be identified with the notion(s) of justification used in epistemology.) Second, as Strawson (1962) argues, a subject can be *exempt* from the requirement to comply with a norm if the subject lacks the relevant rational capacities required to be held responsible for complying with the norm. For instance, young children are exempt from certain norms on behavior. See Littlejohn (forthcoming) for a discussion of the justification/excuse/exemption distinction in epistemology.

What could be the excuse for violating the knowledge norm of assertion in the case of Jones or Smith? There are several options. One could point to any (or several) of the following facts as excusing the norm violation: The assertor reasonably believes the asserted proposition. The assertor reasonably believes that she knows the asserted proposition. In asserting the proposition, the assertor was trying to comply with the norm that governs assertion. The assertor is generally disposed to comply with the norm governing assertion, and the assertion was an exercise of the underlying categorical basis of that disposition. It is through no fault of the assertor that the assertion violates the norm. The assertor reasonably believes that it is highly likely that she is complying with the norm. And so forth.

There are several concerns with this general line of response that appear in the literature. For instance, Gerken (2011) argues against versions of the excuse response where the assertor has an excuse because she reasonably believes that she knows the proposition asserted. Gerken points out that it requires significant cognitive sophistication to believe that one knows something. Presumably, someone in a Gettier-style case or a case of a justified false belief could lack the sophistication required to have such a belief. Nevertheless, intuitively, if such a person asserted the proposition she justifiably believed, the assertion would be appropriate.

There are two main ways an advocate of the excuse response could try to avoid this concern. First, one could identify the excuse not with the fact that the subject reasonably believes that she knows the relevant proposition, but with something less cognitively sophisticated. For instance, one could say all that is required is that the subject's evidence be sufficiently strong to make a belief that she knows the proposition reasonable, even if the subject is incapable of forming that belief. Similarly, one could identify the excuse with the fact that the subject reasonably believes the proposition or with the fact that her assertion is the exercise of the

categorical basis for her disposition to assert only what she knows. Second, one could say that a subject who lacks the cognitive sophistication to believe that she knows the relevant proposition (and is not culpable for this lack) is exempt from the requirement that she have such a belief. Instead, for such a cognitively limited subject, something less is required.⁴⁰ Either of these approaches strikes me as providing a plausible response to Gerken's challenge.⁴¹

A different concern appearing in the literature is that subjects in Gettier-style cases (and cases of justified false belief) intuitively do not owe an apology to their addressees, and so there could not be any violation of a norm.⁴² This concern can be answered by pointing out that not every norm violation requires an apology. In particular, if a norm violation is excusable, an apology is typically unnecessary.⁴³

A closely related concern is that subjects in Gettier-style cases and cases of justified false belief do not need any kind of excuse for their assertions to be appropriate.⁴⁴ I agree with this claim, but I don't see that our intuitions about the cases are clear enough that we can simply read this fact off of our intuitions. The claim is one that must be earned by careful argumentation. So I don't think that this concern is dialectically compelling, either.

There are additional worries that one might have about the excuse response.⁴⁵ Here are two more. The first concerns the target of the norm. The prevailing view in the literature seems to be that the knowledge norm of assertion concerns a normative property of the assertion, not of the assertor. But the claim that the assertor is blameless for making the assertion is a claim about the assertor, not the assertion. If our intuitions about assertions in Gettier-style cases (and cases

⁴⁰ Compare: Young children have fewer epistemic obligations to check up on their reasoning than do competent adults.

⁴¹ But see Gerken (2011) for responses to these suggestions.

⁴² See Brown (2008) and Kvanvig (2011) for versions of this concern.

⁴³ See Littlejohn (2012) and McGlynn (2014).

⁴⁴ See Douven (2006) and Lackey (2007) for versions of this concern.

⁴⁵ I'm not aware that the following concerns appear in the literature.

of justified false belief) concern the normative status of the assertor rather than the assertion, then these intuitions cannot be used to motivate the prevailing view. If instead they concern the normative status of the assertion rather than the assertor, then the excuse response cannot be used to explain them away. This suggests that there is a problem with combining the excuse response and the prevailing view. I find this worry somewhat compelling. However, it is not dialectically very strong. Proponents of the excuse response can simply respond that our intuitive judgments are not fine-grained enough to clearly distinguish between the normative status of the assertion and the normative status of the assertor.

A second additional worry concerns the assertion of a proposition that the assertor unjustifiably believes, such as in the case of Robinson's assertion that there is life on Europa. Suppose that I'm right and, when we carefully consider the case, Robinson's assertion is intuitively appropriate. An advocate of the excuse response will have to find a way to argue that Robinson has a sufficient excuse. The problem is that many of the potential excuses listed above are not ones that apply in the case of Robinson. For instance, Robinson does not reasonably believe the proposition. Robinson does not reasonably believe that he knows the proposition. And so forth. A potential response here, I suppose, is to claim that all that Robinson needs to have in order to count as blameless is a belief in the proposition asserted, or perhaps a belief that he knows the proposition asserted. The problem with this response is that it makes excusably violating the knowledge norm of assertion something that is remarkably easy to do – one need only believe that one is complying with it. That's not typically how excuses work. Merely believing that one is complying with a norm is not sufficient to count as blamelessly violating it. That said, my guess is that proponents of the knowledge norm of assertion will simply reject my

intuition about Robinson and say that his norm violation is not blameless after all. This strikes me as a bullet to bite. But perhaps it is not a very big one.

These concerns about the excuse response, then, are not dialectically very powerful. In the next section, I present a more powerful objection. The objection is a little bit involved – it has a few moving parts – but it is, I think, rather compelling.

6. No Excuses Necessary

Consider the case of an assertion in a Gettier-style case (e.g., the case of Jones) or in the case of a justified false belief (e.g., the case of Smith). For concreteness, I'll focus on the case of Jones, but nothing will hang on this. According to the excuse response, Jones's assertion is improper but Jones is blameless for asserting it, since Jones has an excuse. The trouble is that it is not merely the case that Jones's assertion is blameless. Rather, Jones *positively ought* to make the assertion.⁴⁶ When asked, he should say that it is 2:05 pm. If Jones does not say that it is 2:05 pm (and has no reason not to cooperate with his interlocutor), then he is criticizable for not asserting what he believes about the time.⁴⁷ That Jones has an excuse for making the assertion, and is therefore blameless for so doing, does not explain this fact.⁴⁸

A proponent of the excuse response might reply by saying that there are two distinct considerations involved. Jones has an excuse. That explains why he is blameless in making his assertion. In addition, there is a second consideration that explains why Jones positively ought to

⁴⁶ This already casts some doubt on the excuse response. Typically, when one positively ought to do something, no excuse is necessary in order to count as blameless in doing it, at least when there is a single kind of normativity in play. I don't want to rest much weight on this point, however, since it might be claimed that there are two kinds of normativity in play – for instance, perhaps the kind of blamelessness at issue is epistemic and the kind of ought at issue is not.

⁴⁷ At the very least, asserting that it is 2:05 pm is a good thing for Jones to do. That's all that's really needed for the argument in this section.

⁴⁸ See Schechter (2013) for a version of this point in a different context.

make the assertion. For instance, perhaps there is a norm of politeness that requires Jones to respond to his interlocutor (if he has no reason not to cooperate).

There are difficulties facing this specific proposal.⁴⁹ But there are also more general problems with the suggestion that there are two separate considerations in play. First, whenever possible, it is better to avoid multiplying considerations, especially when the phenomena being explained are, on the face of it, closely related. Second, if there is a separate consideration that can explain why Jones positively ought to make the assertion, there is the danger that this consideration could be also used to explain the intuitive data alleged to support the knowledge norm. So there is the danger that there would be no work left for the knowledge norm to do.

More importantly, recall where the knowledge norm of assertion is supposed to come from. On the most plausible versions of knowledge-first epistemology, the practice of assertion has a function or purpose. The knowledge norm of assertion flows from this function or purpose. Assuming that this is right, the reason that Jones positively ought to assert that it is 2:05 pm would also seem to flow from the function or purpose of assertion. For instance, suppose that the central function of assertion is to enable the transmission and pooling of knowledge. Assuming this view, it is plausible to think that Jones ought to respond to his interlocutor because (very roughly) doing so is the kind of thing that tends to contribute to the transmission and pooling of knowledge. Alternatively, suppose that the central function of assertion is to help agents to

⁴⁹ Here are two difficulties: First, a norm of politeness may be able to explain why Jones ought to respond in some way to his interlocutor – after all, it would be rude for him to remain silent. But such a norm cannot explain why Jones ought to respond by asserting what he believes the time to be. There are many other things that Jones could instead assert that would be equally polite – for instance, “I’m very sorry, I cannot help”. Second, a norm of politeness cannot handle all of the relevant cases. Suppose that instead of being asked the time, Jones notices that someone nearby is wondering what time it is – frantically looking around for a clock, stopping people to ask if they have a watch, and so forth. Even if Jones is not asked what the time is, so long as Jones has no reason to avoid engaging with this person, there is still some sense in which Jones positively ought to say that it is 2:05 pm. At the very least, that would be a good thing for him to do. A norm of politeness cannot explain this fact. This suggests that it would be better to appeal to a norm of helpfulness instead of a norm of politeness. In a way, I think this is right. See below.

coordinate more effectively. Assuming this view, it is plausible to think that Jones ought to respond to his interlocutor because (very roughly) doing so is the kind of thing that tends to contribute to effective coordination. And so forth. What this suggests is that there are not two separate considerations in play. The reason why Jones is blameless in making the assertion and the reason why Jones positively ought to make the assertion are the very same reason or, at least, are closely connected reasons.

The natural view for proponents of the excuse response to take is to agree that the fact that Jones has an excuse does not explain why Jones ought to make the assertion he does. It is not the existential claim that Jones has an excuse that explains why Jones ought to make the assertion. Rather, it is the fact that constitutes his excuse that explains why he positively ought to make the assertion. On this view, one and the same fact is significant in two different ways – it functions both as an explanation of blamelessness and as an explanation of an ought.

So far, the proponent of the excuse response has been able to respond to the objection. But things are about to get a bit stickier.

Proponents of the excuse response should not claim that the kind of excuse at issue in the problem cases for the knowledge norm of assertion is a special kind of excuse that is specific to the problem cases, or even to assertion more generally. Any account of the kind of excuse at issue should show how such excuses fit in with a general picture of norms and norm compliance. Otherwise, the resulting view will be unwieldy and ad hoc.⁵⁰

How can the excuses that feature in the problem cases be tied to a more general account of excuses for norm violations? The most sophisticated proposal of which I am aware is due to

⁵⁰ See Hill and Schechter (2007) for an articulation of the concern that the excuse response commits Williamson to an overly complex view.

Williamson.⁵¹ On Williamson’s account, when there is a norm governing a practice, there is typically also a secondary norm to be the sort of person who complies with the original or primary norm.⁵² Williamson suggests that this secondary norm further generates a tertiary norm to do in any situation what someone who complies with the secondary norm (that is, what someone who is disposed to comply with the primary norm) would do. For instance, there is a norm to keep one’s promises. This generates the secondary norm to be the sort of person who keeps one’s promises. This further generates the tertiary norm to do in any situation what someone who is disposed to keep their promises would do.

Williamson further claims that while some of the normative force of the original norm “rubs off” on the secondary and tertiary norms, the derivative norms typically do not “inherit the full normative status” of the original norm.⁵³ For instance, in breaking a promise, one is breaking one’s word, not living up to a commitment one has undertaken, etc. In not being the sort of person who keeps one’s promises – so long as one has not actually broken a promise – one has done no such thing. Finally, Williamson claims it is often a good excuse for violating a primary norm that one is complying with the tertiary norm it generates. (He also claims that if one complies with both the secondary and tertiary norms, this typically strengthens the excuse.) In particular, if an agent breaks a promise in a certain situation, the fact that the agent is doing what someone who keeps their promises would do in that situation will often excuse the norm violation.

There is much to say in favor of Williamson’s account. It is elegant. It fits with familiar examples of norm-governed behavior. (Williamson focuses on legal norms and the norms

⁵¹ See Williamson (forthcoming). Also see Littlejohn (forthcoming) for a related view.

⁵² “Complies with the norm” should be understood as a generic – to be the sort of person who complies with a norm does not require that one actually complies with the norm in any given situation.

⁵³ Williamson (forthcoming).

governing promising. One might also add the cases of spelling, grammar, etiquette, and games.) However, in the current context, there is a problem. Namely, Williamson's account does not sufficiently explain why it is that, in the problem cases for the knowledge norm of assertion, the assertor *ought* to make the assertion.

Applying Williamson's account, the knowledge norm for assertion generates a secondary norm: One should be the sort of person who asserts only what one knows.⁵⁴ It further generates a tertiary norm: In any given situation, one should assert what the sort of person who asserts only what they know would assert. Let us grant that Jones is the sort of person who asserts only what he knows. Let us also grant that in asserting that it is 2:05 pm, Jones is asserting what someone who is the sort of person who asserts only what they know would assert. So Jones is complying with the secondary and tertiary norms. This could explain, perhaps, why Jones is blameless in making the assertion. But it does not sufficiently explain why Jones positively ought to make the assertion and why Jones would be criticizable if he failed to make the assertion.

One might think that there is a simple reason why Jones is criticizable if he doesn't make the assertion: In not making the assertion, he would be violating the tertiary norm. This isn't quite right – the primary norm doesn't mandate any assertions at all, and so neither does the tertiary norm. But this problem is easily fixed. Presumably, there is a general norm of helpfulness that entails that one should respond to reasonable questions about the time. (Perhaps such a norm is tied to Grice's maxim of quantity.)⁵⁵ In combination with the knowledge norm of assertion, this explains why one should tell someone the time when asked, if one knows it. In combination with the tertiary assertion norm, the line of thought goes, this mandates that one should tell someone what one believes the time to be if doing so is what someone who is helpful

⁵⁴ See Lasonen Aarnio (2010), which makes use of a secondary norm on belief in order to try to explain away our intuitions that knowledge can be defeated.

⁵⁵ See Grice (1975).

and is the kind of person who asserts only what they know would assert. So the tertiary norm (in combination with a general norm of helpfulness) does mandate that Jones make the assertion, after all. What then is the problem?

The problem is that the kind of mandate that tertiary norms provide is too weak to account for the force of Jones's requirement to assert what he takes the time to be. As Williamson admits, secondary and tertiary norms typically do not "inherit the full normative status" of the primary norm. Not doing what someone who is the sort of person who keeps their promises would do is very different from breaking a promise. One ought to keep one's promises. One also ought to do what the sort of person who keeps their promises would do. But the former ought is stronger and more "full-blooded" than the latter ought. Breaking a promise is a much more serious violation than not doing what the sort of person who keeps their promises would do.

In the Jones case, however, not asserting that it is 2:05 pm seems every bit as bad – and bad in the very same way – as it would be if Jones knew that it was 2:05 pm. One way to see this is to think about Jones and twin Jones. Jones and twin Jones are in exactly the same situation – they have both just glanced at a clock that reads 2:05 and they have both been asked for the time – except for one difference. Jones glanced at a clock that stopped twelve hours ago and so doesn't know the time. Twin Jones glanced at a working clock and does know the time. I maintain that Jones and Twin Jones face equally strong pressure to assert that it is 2:05 pm. The oughts that apply to them are exactly on a par.

A different way to see the point is to consider a diachronic case where the subject moves from knowledge to a merely justified belief and back again.⁵⁶ It is a bit easier to make use of the

⁵⁶ See White (2014) for diachronic cases used against Williamson's externalist response to the new evil demon argument.

case of Smith rather than the case of Jones, so let's turn to that case. Suppose that on Monday, Smith counts as knowing that the fastest way to get to Thayer Street from a certain location is to make a right on Benefit Street and then a left on Waterman Street. Suppose that Smith is asked by a passerby "What's the fastest way for me to get to Thayer from here right now?" Smith ought to respond by saying that the fastest way to get to Thayer Street is to make a right on Benefit Street and then a left on Waterman Street. Suppose that on Tuesday, unbeknownst to Smith, Benefit Street is closed due to a water main break. Suppose that Smith is asked by another passerby "What's the fastest way for me to get to Thayer from here right now?" Smith again ought to respond by saying that the fastest way to get to Thayer Street is to make a right on Benefit Street and then a left on Waterman Street. Suppose that on Wednesday, everything is back to normal. Suppose that Smith is asked by a yet another passerby "What's the fastest way for me to get to Thayer from here right now?" Yet again, Smith ought to respond by saying that the fastest way to get to Thayer Street is to make a right on Benefit Street and then a left on Waterman Street. Indeed, the strength of the oughts on all three days seems just the same. Not giving the passerby directions would be just as bad on any of the days, and in just the same way. Williamson's account cannot accommodate this point.⁵⁷

The objection to the excuse response thus involves the interaction of three constraints on an account of the kind of excuse at issue in the problem cases for the knowledge norm of assertion. First, it is a constraint that the account not merely explain why the assertions are blameless, but also explain why the assertors ought to assert what they do. Second, it is a constraint that the account explain the full normative strength of this ought – and in particular,

⁵⁷ Williamson might claim that Smith doesn't have knowledge of the fastest route on Monday or Wednesday, since Smith's belief is not safe. We can modify the case to avoid this response. For instance, we can change the case so that Smith is asked for directions on days that are months apart. We can also specify that the water main break is not caused by a random fluke (and thus avoid any resemblance to a lottery case). Presumably, there is some variant of the case on which Smith loses and gains knowledge. We can substitute this variant for the case presented in the text.

why the ought is as strong in a problem case as it is in a corresponding knowledge case. Third, it is a constraint that the kind of excuse at issue not be some special kind of excuse particular to assertion. Instead, the account must fit with a general account of norm-governed behavior. The trouble is that it is difficult to see how these constraints can be satisfied simultaneously. I've illustrated this by focusing on Williamson's account. But it is difficult to see that an alternative account can do better.

It is worth noting that the dialectical situation would seem to be analogous for other proposed knowledge norms – including the knowledge norms of reasoning, belief, and action. It can be appropriate for a subject to use the proposition that *p* as a premise in reasoning, to believe that *p*, or to act as if *p*, despite not knowing that *p*. In response to problem cases for one of these norms, it might be suggested that in such cases, the subject has violated the norm, but has done so blamelessly since the subject has an appropriate excuse. But, again, this response underdescribes the situation. In many problem cases, the subject positively ought to reason with the proposition, believe the proposition, or act as if the proposition is true. And it seems difficult to provide an account of the kind of excuse at issue in these problem cases that also explains where this positive ought comes from, accommodates the full normative strength of the ought, and does so in a non-ad hoc way.⁵⁸

7. Conclusion: Assertion without Knowledge

Let's take stock. In the first part of this paper, I briefly presented the program of knowledge-first epistemology. I then put forward three big-picture worries for this program: (i) the post-Gettier literature suggests that knowledge is a marginal status; (ii) for knowledge-first epistemology to

⁵⁸ If anything, the problem for these other norms is worse than for the knowledge norm of assertion. The suggestion that there are two separate considerations in play – one that explains blamelessness and one that explains the positive ought – is even less plausible for the cases of reasoning, belief, and action.

be tenable, we need an explanation of what knowledge could be such that it is central to our mental lives; and (iii) the diversity of our mental lives – involving many different kinds of cognitive states and epistemic statuses and involving transitions and rules in addition to cognitive states – casts doubt on the centrality of knowledge. I also conceded that while these objections are pressing, they are not conclusive. Determining the viability of knowledge-first epistemology requires carefully evaluating the individual theses endorsed by proponents.

In the second part of the paper, I turned to a specific thesis that is widely endorsed by knowledge-first epistemologists – the knowledge norm of assertion. I presented a common objection to this thesis: In many cases, it is intuitively appropriate for someone who has a strongly justified belief that *p*, but who doesn't know that *p*, to assert that *p*. I then presented the standard response to this objection: In such a case, the assertion is improper but the assertor is blameless because the assertor possesses a sufficiently good excuse. I argued that extant objections to this response are not dialectically very powerful. Instead, I presented a new (and somewhat complicated) objection. In its essentials, the objection is that it is difficult to provide a non-ad hoc account of the assertor's excuse that explains why the assertor ought to make the assertion, where this ought is just as strong as it would be in a case of knowledge. This, I conclude, provides us with reason to reject the knowledge norm of assertion. (Analogous objections provide us with reason to reject the knowledge norms of reasoning, belief, and action.) Derivatively, this provides us with some reason to reject knowledge-first epistemology.

Supposing we reject the knowledge norm of assertion, how should we replace it? The discussion above provides some support for moving to a different norm. If we give up on the excuse response, the cases of Jones, Smith, and Robinson suggest that knowledge, truth, and

justification are not necessary conditions on proper assertion. That leaves belief, or perhaps belief that one knows.⁵⁹ Perhaps, then, we should accept one of these two norms.

I suspect not. There are assertions that are intuitively appropriate in which the asserter does not believe the proposition asserted. Suppose, for example, that one is tutoring a student in physics and knows that if one puts in all of the required hedges, the student will get lost in the weeds.⁶⁰ Similarly, suppose that some situation is urgent, and the most straightforward thing one could do to provide the required aid is to say something one knows to be false.⁶¹ In each of these cases, it strikes me as appropriate to assert claims that one knows to be false.⁶²

Proponents of various norms – the belief norm, the justified belief norm, the belief that one knows norm, the knowledge norm – will try to handle these cases by saying that the assertions violate the norm of assertion, but that the educational needs of the student or the urgency of the situation makes the violation blameless. This maneuver may work here. It is not subject to same objections that face the excuse response discussed above. But it would certainly be more satisfying if we didn't have to categorize these assertions as norm violations to begin with. It would be theoretically preferable if we could adopt an account on which these assertions are normatively on a par with other intuitively appropriate assertions.

The natural account to reach for here is a broadly Gricean account. In particular, one might think that in the case of the physics student or the case of the urgent situation, what makes the assertion appropriate is that it is the most helpful thing to say in context – the asserter takes the proposition to be more relevant to the listener's informational needs than any alternative

⁵⁹ Bach and Harnish (1979) argue that assertions express beliefs. Bach (2008) further argues that this generates a belief norm on assertion. See Williamson (2000) for discussion of the belief that you know norm.

⁶⁰ For a related case, see McKinnon (2013), p. 124.

⁶¹ For a related case, see Williamson (2000), p. 256.

⁶² For additional cases, see the examples of "selfless assertions" in Lackey (2007).

proposition that could be asserted.⁶³ On this approach, there are no norms specifically tied to assertion.⁶⁴ Instead, there are general pragmatic principles governing communication, such as a principle that tells us to do what is communicatively most helpful in context. This approach requires that we find some other way, not involving norms, to distinguish assertion from the other speech acts.⁶⁵ It also requires that we respond to the positive arguments that have been presented for the proposed norms of assertion. That will take a lot of work. But it strikes me as the most promising approach to take.⁶⁶

References

- Adler, Jonathan. (2002) *Belief's Own Ethics*. Cambridge, Mass.: MIT Press.
- Austin, John L. (1956) "A Plea for Excuses," *Proceedings of the Aristotelian Society* 57: 1–30.
- Bach, Kent. (2008) "Applying Pragmatics to Epistemology," *Philosophical Issues* 18: 68–88.
- Bach, Kent and Harnish, Robert. (1979) *Linguistic Communication and Speech Acts*. Cambridge, Mass.: MIT Press.
- Benton, Matthew. (2011) "Two More for the Knowledge Account of Assertion," *Analysis* 71: 684–687.
- Blaauw, Martijn. (2012) "Reinforcing the Knowledge Account of Assertion," *Analysis* 72: 105–108.
- BonJour, Laurence. (2011) "The Myth of Knowledge," *Philosophical Perspectives* 24: 57–83.
- Brown, Jessica. (2008) "The Knowledge Norm for Assertion," *Philosophical Issues* 18: 89–103.
- Brown, Jessica. (2010) "Knowledge and Assertion," *Philosophy and Phenomenological Research* 81: 549–566.
- Byrne, Alex and Logue, Heather. (eds.) (2008) *Disjunctivism: Contemporary Readings*. Cambridge, Mass.: MIT Press.
- Clark, Michael. (1963) "Knowledge and Grounds: A Comment on Mr. Gettier's Paper," *Analysis* 24: 46–48.
- Cohen, Stewart. (1984) "Justification and Truth," *Philosophical Studies* 46: 279–295.

⁶³ This is related to Grice's maxims of relevance and quantity. See Weiner (2005), Hill and Schechter (2007), and Lackey (2008) for applications of Gricean principles in responding to arguments in favor of the knowledge norm of assertion.

⁶⁴ See Sosa (2009) and Johnson (2017) for arguments that assertion does not have a constitutive norm.

⁶⁵ See MacFarlane (2011) for a survey of approaches.

⁶⁶ Thanks to Stewart Cohen and Timothy Williamson for discussion of these issues. Thanks to David Black, David Christensen, Louis deRossett, Mikkel Gerken, Jeremy Goodman, Christopher Hill, Casey Johnson, and David Plunkett for helpful comments on this paper. Thanks to the students in David Plunkett's undergraduate seminar at Dartmouth for useful questions and remarks. Thanks to the three editors of this volume for numerous suggestions that led to improvements. This paper was written while I was a visiting fellow at the University of Connecticut Humanities Institute. I'm grateful to the University of Connecticut for its hospitality.

- Cohen, Stewart and Lehrer, Keith. (1983) "Justification, Truth, and Coherence," *Synthese* 55: 191–207.
- DeRose, Keith. (2002) "Assertion, Knowledge, and Context," *Philosophical Review* 111: 167–203.
- Douven, Igor. (2006) "Assertion, Knowledge, and Rationality," *Philosophical Review* 115: 449–485.
- Engel, Pascal. (2008) "In What Sense is Knowledge the Norm of Assertion?," *Grazer Philosophische Studien* 77: 45–59.
- Gerken, Mikkel. (2011) "Warrant and Action," *Synthese* 178: 529–547.
- Gettier, Edmund. (1963) "Is Justified True Belief Knowledge?," *Analysis* 23: 121–123.
- Goldman, Alvin. (1967) "A Causal Theory of Knowing," *Journal of Philosophy* 64: 357–372.
- Grice, H. Paul. (1975) "Logic and Conversation," in Peter Cole and Jerry Morgan (eds.), *Syntax and Semantics*, Vol. 3. New York: Academic Press, pp. 41–58.
- Grimm, Stephen. (2006) "Is Understanding a Species of Knowledge?," *British Journal for the Philosophy of Science* 57: 515–535.
- Hawthorne, John. (2004) *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John and Stanley, Jason. (2008) "Knowledge and Action," *Journal of Philosophy* 105: 571–590.
- Hill, Christopher and Schechter, Joshua. (2007) "Hawthorne's Lottery Puzzle and the Nature of Belief," *Philosophical Issues* 17: 102–122.
- Hinton, John M. (1967) "Visual Experiences," *Mind* 76: 217–27.
- Johnson, Casey. (2017) "What Norm of Assertion?" *Acta Analytica*.
- Kaplan, Mark. (1985) "It's Not What You Know that Counts," *Journal of Philosophy* 82: 350–363.
- Kelp, Christoph. (2015) "Understanding Phenomena," *Synthese* 192: 3799–3816
- Khalifa, Kareem. (2011) "Understanding, Knowledge, and Scientific Antirealism," *Grazer Philosophische Studien* 83: 93–112.
- Koethe, John. (2009) "Knowledge and the Norms of Assertion," *Australasian Journal of Philosophy* 87: 625–638.
- Kvanvig, Jonathan. (2003) *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
- Kvanvig, Jonathan. (2009) "Assertion, Knowledge, and Lotteries," in Patrick Greenough and Duncan Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press, pp. 140–160.
- Kvanvig, Jonathan. (2011) "Norms of Assertion," in Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*. Oxford: Oxford University Press, 233–250.
- Lackey, Jennifer. (2007) "Norms of Assertion," *Noûs* 41: 594–626.
- Lackey, Jennifer. (2011) "Assertion and Isolated Second-Hand Knowledge," in Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*. Oxford: Oxford University Press, pp. 251–275.
- Lasonen Aarnio, Maria. (2010) "Unreasonable Knowledge," *Philosophical Perspectives* 24: 1–21.
- Lehrer, Keith and Paxson, Thomas. (1969) "Knowledge: Undefeated Justified True Belief," *Journal of Philosophy* 66: 225–237.
- Lewis, David. (1983) "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61: 343–77.

- Littlejohn, Clayton. (2012) *Justification and the Truth Connection*. Cambridge: Cambridge University Press.
- Littlejohn, Clayton. (2013) "The Russellian Retreat," *Proceedings of the Aristotelian Society* 113: 293–320.
- Littlejohn, Clayton. (forthcoming) "A Plea for Epistemic Excuses," in Fabian Dorsch and Julien Dutant (eds.), *The New Evil Demon Problem*. Oxford: Oxford University Press.
- MacFarlane, John. (2011) "What is Assertion?," in Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*. Oxford: Oxford University Press, pp. 79–96.
- Maitra, Ishani. (2011) "Assertion, Norms, and Games," in Jessica Brown and Herman Cappelen (eds.), *Assertion: New Philosophical Essays*. Oxford: Oxford University Press, pp. 277–296.
- Martin, Michael. (2002) "The Transparency of Experience," *Mind and Language* 17: 376–425.
- McDowell, John. (1982) "Criteria, Defeasibility and Knowledge," *Proceedings of the British Academy* 68: 455–479.
- McGlynn, Aidan. (2014) *Knowledge First?*. New York: Palgrave Macmillan.
- McKinnon, Rachel. (2013) "The Supportive Reasons Norm of Assertion," *American Philosophical Quarterly* 50: 121–135.
- Nozick, Robert. (1981) *Philosophical Explanations*. Cambridge, Mass.: Harvard University Press.
- Pritchard, Duncan. (2005) *Epistemic Luck*. Oxford: Oxford University Press.
- Pritchard, Duncan. (2012) *Epistemological Disjunctivism*. Oxford: Oxford University Press.
- Reynolds, Steven. (2002) "Testimony, Knowledge, and Epistemic Goals," *Philosophical Studies* 110: 139–161.
- Russell, Bertrand. (1912) *The Problems of Philosophy*. London: Williams and Norgate.
- Schaffer, Jonathan. (2008) "Knowledge in the Image of Assertion." *Philosophical Issues* 18: 1–19
- Schechter, Joshua. (2013) "Rational Self-Doubt and the Failure of Closure," *Philosophical Studies* 163: 429–452.
- Shope, Robert. (1983) *The Analysis of Knowing: A Decade of Research*. Princeton: Princeton University Press.
- Sliwa, Paulina. (2014) "Understanding and Knowing," *Proceedings of the Aristotelian Society* 115: 57–74.
- Slote, Michael. (1979) "Assertion and Belief," in Jonathan Dancy (ed.), *Papers on Language and Logic*. Keele: Keele University Press, pp. 177–90.
- Snowdon, Paul. (1980) "Perception, Vision and Causation," *Proceedings of the Aristotelian Society* 81: 175–92.
- Sosa, David. (2009) "Dubious Assertions," *Philosophical Studies* 146: 269–272.
- Stanley, Jason. (2005) *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Strawson, Peter. (1962) "Freedom and Resentment," *Proceedings of the British Academy*, 48: 1–25.
- Sutton, Jonathan. (2007) *Without Justification*. Cambridge, Mass.: MIT Press.
- Turri, John. (2010) "Epistemic Invariantism and Speech Act Contextualism," *Philosophical Review* 119: 77–95.
- Unger, Peter. (1975) *Ignorance: A Case for Skepticism*. Oxford: Clarendon Press.
- Weiner, Matthew. (2005) "Must We Know What We Say?," *Philosophical Review* 114: 227–251.

- Weiner, Matthew. (2009) "Practical Reasoning and the Concept of Knowledge" in Adrian Haddock, Alan Millar, and Duncan Pritchard (eds.), *Epistemic Value*, Oxford: Oxford University Press, pp. 163–82.
- White, Roger. (2014) "What is My Evidence that Here is a Hand?," in Dylan Dodd and Elia Zardini (eds.), *Scepticism and Perceptual Justification*. Oxford: Oxford University Press, pp. 298–322.
- Whiting, Daniel. (2013) "Stick to the Facts: On the Norms of Assertion," *Erkenntnis* 78:847–867.
- Williamson, Timothy. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, Timothy. (2013) "Knowledge First," in Matthias Steup, John Turri, and Ernest Sosa (eds.), *Contemporary Debates in Epistemology*, second edition. Malden, Mass.: Wiley-Blackwell, pp. 1–9.
- Williamson, Timothy. (forthcoming) "Justifications, Excuses, and Sceptical Scenarios," in Fabian Dorsch and Julien Dutant (eds.), *The New Evil Demon Problem*. Oxford: Oxford University Press.