

Omega Knowledge Matters

Abstract

You omega know something when you know it, and know that you know it, and know that you know that you know it, This paper first argues that omega knowledge matters, in the sense that it is required for rational assertion, action, inquiry, and belief. The paper argues that existing accounts of omega knowledge face major challenges. One account is skeptical, claiming that we have no omega knowledge of any ordinary claims about the world. Another account embraces the KK thesis, and identifies knowledge with omega knowledge. This position faces counterexamples, and struggles to make sense of inexact knowledge. The paper then develops a new account of knowledge, by proposing the principle of Reflective Luminosity: if you know that you know something, then you omega know it. I argue that Reflective Luminosity allows for omega knowledge while avoiding the problems for KK.

1 Introduction

Omega knowledge is the strongest kind of knowledge. When you omega know something, you know it. You know that you know it. You possess every iteration of knowledge regarding it. More precisely, you 1-know p when you know p . You n -know p when you know that you $(n - 1)$ -know p . You omega know p when you n -know p , for every n .

I have written a forthcoming book, *Omega Knowledge: What it is and Why it Matters*, which argues that omega knowledge plays an important role in philosophy, and then develops three theories of what omega knowledge is ([Goldstein Forthcoming](#)). This companion paper introduces my two arguments for why omega knowledge matters, along with one of my theories of what omega knowledge is.

In the first part of the paper, I'll argue that omega knowledge matters. In particular, I'll argue that omega knowledge is necessary for permissible assertion, action, inquiry, and belief.

I'll also argue that existing theories of knowledge don't explain the importance of omega knowledge. There are two existing approaches to omega knowledge. The first theory, Omega Skepticism, says that you omega know almost nothing.

- (1) **Omega Skepticism.** You fail to omega know most ordinary claims about the world.

According to omega skeptics, every extra iteration of knowledge imposes a stronger demand on your powers of discrimination. Infinite iterations of knowledge impose infinite demands, and these demands can't be satisfied. I think that omega knowledge is required for coordination, assertion, action, and successful inquiry. So I think that Omega Skepticism is too skeptical.

The second theory, KK, says that you omega know everything you know.

- (2) **KK.** If you know p , then you know that you know p .¹

I reject KK for two reasons. First, KK is vulnerable to counterexamples, like the unconfident examinee. Second, KK struggles to explain cases of inexact knowledge, where knowledge is governed by a Margin for Error principle. For example, when you are looking at a tree:

- (3) **Margin for Error.** For any height x , if you know that the tree is not $x - 1$ feet tall, then the tree is not x feet tall.

The goal of this paper is to make room for omega knowledge. To do so, I'll introduce a new principle about knowledge that explains how omega knowledge is possible.

- (4) **Reflective Luminosity.** If you know you know p , then you omega know p .

Reflective Luminosity avoids Omega Skepticism: you can omega know p by coming to know that you know p . But Reflective Luminosity is weaker than KK. It allows that you can know something without omega knowing it, as long as you don't know that you know it. This allows us to avoid standard counterexamples to KK. In addition, I argue that it can make sense of Margin for Error, by predicting that it is normally true rather than universally valid.

By introducing this alternative theory of knowledge, one of my goals is to undermine existing arguments for KK. These arguments do not require the full force of KK; instead, they simply require that Omega Skepticism is false.

2 Against Omega Skepticism

I'll begin by objecting to Omega Skepticism, by arguing that omega knowledge matters. My two arguments claim that omega knowledge is necessary for various permissible behavior. My first argument is that omega knowledge is necessary for permissible assertion. My second argument is that omega knowledge is necessary for permissible action more generally. If my arguments are correct, then Omega Skepticism would lead to a wide range of ordinary behavior being forbidden. Since this is not true, Omega Skepticism must be false.

¹See [Alston 1980](#), [Adler 1981](#), [Williamson 2000](#), [Hawthorne and Magidor 2010](#), [Carter 2018](#) and others for attacks on KK. See [Stalnaker 2006](#), [Stalnaker 2015](#), [Dokic and Égré 2009](#), [McHugh 2010](#), [Cohen and Comesaña 2013](#), [Greco 2014a](#), [Greco 2014b](#), [Das and Salow 2018](#), [Salow 2018](#), and others for defenses of KK.

First, I'll argue that assertion requires omega knowledge:

- (5) **Omega Assertion.** You are permitted to assert p only if you omega know p .

Omega Assertion is stronger than a simpler knowledge norm, which says that you are permitted to assert p only if you know p .² Omega Assertion implies that this rival knowledge norm is true. But it adds the further requirement of omega knowledge.

A knowledge norm on assertion can explain why it is strange to assert Moorean conjunctions, like:

- (6) p and I don't know p .

You can't know these claims. So you can't assert them. Since Omega Assertion implies that assertion is governed by a knowledge norm, Omega Assertion also predicts that Moorean conjunctions are strange to assert.

My argument for Omega Assertion is that it explains the infelicity of 'dubious assertions' like:

- (7) a. p but I don't know whether I know p .
b. p but I don't know whether I know that I know p .

Omega Assertion explains why (7-a) and (7-b) are strange to assert. You can't omega know (7-a); suppose you did. Then, since knowledge distributes over conjunction, you would omega know p .³ But since omega knowledge is factive, you wouldn't know that you know p , and so wouldn't omega know p . This is inconsistent.⁴

I accept Omega Assertion because it explains the infelicity of dubious assertions. But Omega Assertion is difficult to reconcile with Omega Skepticism. Omega Skepticism and Omega Assertion imply that you are not permitted to assert any ordinary claim.

Omega skeptics tend to believe that only knowledge is required for permissible assertion (see for example [Williamson 2000](#)). But Omega Skepticism allows for dubious assertions to be known. In theories of knowledge that embrace Omega Skepticism, you can know that p is true while knowing that you don't know that you know p . So you can know a dubious assertion. So a simple knowledge norm on assertion doesn't explain the infelicity of dubious assertions.

Similarly, the infelicity of dubious assertions cannot be explained immediately

²For defenses of the knowledge norm, see among others [Unger 1975](#), [Williamson 1996](#), [Williamson 2000](#), [DeRose 1996](#), [DeRose 2002](#), [Adler 2002](#), [Hawthorne 2003](#), [Stanley 2005](#), [Engel 2008](#), [Schaffer 2008](#), [Turri 2011](#), and [Turri 2015](#).

³Throughout the book, I assume that knowledge is closed under competent deduction: if you know some premises, and competently deduce a conclusion from these premises while retaining knowledge of the premises, you know the conclusion. See [Williamson 2000](#) and [Hawthorne 2003](#) for discussion.

⁴See [Sosa 2009](#) and [Cohen and Comesaña 2013](#), [Greco 2014a](#) and [Das and Salow 2018](#) for further defense of this style of argument. See [Benton 2013](#) and [Williamson 2013](#) for alternative explanations.

by a reflective knowledge norm, which says that you are permitted to assert p only if you know that you know p . Nothing in this position alone rules out the possibility of knowing that you know: p and you don't know that you know that you know p . If such a conjunction were reflectively known, the reflective knowledge norm on assertion would permit its assertion. In this way, the infelicity of dubious assertions seems to require Omega Assertion, which in turn suggests that Omega Skepticism is false.

On the other hand, defenders of KK have used dubious assertions to motivate the KK principle. Since KK implies that you omega know anything you know, KK can easily explain the infelicity of dubious assertions.⁵ But one goal of this paper is to show that it is not ultimately KK that does the crucial explanatory work in various cases; instead, it is the assumption that omega knowledge is abundant rather than scarce. As I'll show, there are other ways to deny Omega Skepticism, besides accepting KK.

My second argument against Omega Skepticism relies on a general principle about permission. The principle says that whenever your behavior is governed by a norm, it is also governed by a 'secondary' epistemic norm (See for example DeRose 2002, DeRose 2009, Williamson 2005, Benton 2013, and Goodman and Holguin 2021.) In order for your behavior to be appropriate, you must know that you satisfy the norm. If you don't know this, then performing that behavior would be too risky.

For example, you are permitted to drive home from the bar only if you are sober. Nonetheless, you are not automatically permitted to drive home if you are sober. What if you've drunk two beers, and you don't know whether you're sober? In that case, you aren't permitted to drive home, even if you actually are sober. Driving home would be too risky.

Norm Iteration generalizes this line of reasoning to any behavior A and proposition p :

- (8) **Norm Iteration.** If you are permitted to A only if p , then you are permitted to A only if you know p .

Norm Iteration says that whenever you are governed by one norm, you are also governed by another. If doing A requires that p is true, then doing A also requires that you know p is true.

Norm Iteration explains why it is strange to perform some behavior while simultaneously conceding that you don't know whether you're allowed to perform that behavior. If you know that you don't know whether you satisfy the conditions for permissibly performing an action, then you know that you are behaving impermissibly.

Norm Iteration implies that omega knowledge is necessary for permissible behavior. Return to the case of drunk driving. Suppose again that you're permitted to drive home only if you're sober. Norm Iteration implies that you're permitted to drive home only if you know you're sober. But Norm Iteration also applies to this second norm. Norm Iteration implies that you're permitted

⁵See for example Cohen and Comesaña 2013, Greco 2014a, and Das and Salow 2018.

to drive home only if you know that you know you're sober. An infinite series of applications of Norm Iteration then imply that you are permitted to drive home only if you omega know you're sober. Generalizing from this example, Norm Iteration implies that if p is necessary in order for some behavior A to be permitted, then omega knowledge of p is also necessary in order for A to be permitted.⁶

Norm Iteration explains why Omega Assertion is true. Suppose that you are permitted to assert p only if p is true. Norm Iteration then implies that you are permitted to assert p only if you omega know p .

Omega Skepticism says that you fail to omega know most ordinary claims. If Norm Iteration is true, it poses a threat to Omega Skepticism. To threaten Omega Skepticism about some particular claim p , one strategy is to find some behavior that is permitted only if p is true. For example, imagine you are looking at a tree that appears to be 100 feet tall. Now consider the ordinary proposition that the tree is at least 80 feet tall. Omega skeptics claim that you fail to omega know this claim. But now suppose you're correctly informed by a reliable informant that an innocent person will be killed if the tree is less than 80 feet tall and you press a certain button. I think it follows that you're only permitted to press the button if the tree is at least 80 feet tall. It then follows from Norm Iteration that you're only permitted to press the button if you omega know the tree is at least 80 feet tall.

Omega skeptics may reject Norm Iteration. One strategy would be to restrict the application of Norm Iteration. You are permitted to drive only if you know you are sober. You are permitted to drive only if you know that you know you are sober. But even if you don't omega know you are sober, you can be permitted to drive. For some value of n , you are permitted to drive only if you n -know that you are sober; but you are permitted to drive even if you don't $n + 1$ -know that you are sober. Perhaps it is vague where exactly Norm Iteration fails.⁷

Some omega skeptics have tried to explain dubious assertions and other phenomena by appeal to secondary norms. For example, even though p and I don't know that I know that p can satisfy the primary knowledge norm of assertion, [Benton 2013](#) suggests that the sentence is defective because it violates the secondary requirement of knowing that you satisfy the primary norm of assertion. In order to exploit this strategy, the omega skeptic must deny Norm Iteration at higher orders of knowing.

On this picture, each application of Norm Iteration generates a requirement

⁶The normative requirements governed by Norm Iteration could be interpreted in either a narrow scope or wide scope manner. In the narrow scope interpretation, the normative requirements would be material conditionals of the form [you are permitted to A] only if p . By contrast, [Williamson 2000](#) understands norms using the wide scope configuration: you should [A only if p]. I won't take a stand on how best to interpret normative requirements in this paper.

⁷See [Marušić 2013](#) p. 1997. Similarly, another reaction to Norm Iteration claims that the sense of 'permitted' in the consequent of Norm Iteration is different than the sense of 'permitted' in the antecedent. Each application of Norm Iteration generates a different kind of norm, with a different notion of 'permission'. This view could be combined with the idea that some senses of 'permission' are irrelevant to deliberation about what action to perform.

with weaker force than the previous application. It is bad to drive home drunk. It is not as bad to drive home ignorantly sober. It is less bad yet to drive home ignorant about whether you are ignorantly sober. Each application of Norm Iteration generates a requirement with weaker force. Granted, it is risky to act on known knowledge when you don't know that you know that you know. But it is less risky to do this than to act on knowledge when you don't know that you know. At some point, the idea goes, Norm Iteration fails to create a new norm.

One development of this theory says that for any n , it is bad to some degree to drive if you don't n -know that you're sober. But some degrees of badness are so small that they don't imply that driving home is impermissible. After all, there could also be something bad about taking a cab instead of driving. One problem with this proposal is that you might think that many actions I perform are not bad in any way. For example, imagine that after drinking several beers, I decide to take a cab home from the bar instead of driving my car. Plausibly, there would be nothing at all bad about doing this. It is permitted in every sense. But if we omega know very few things, then this proposal predicts that almost all of the actions we perform are bad in some sense.

A different development of the theory says that for some value of n , it is bad to drive if you don't n -know that you're sober; but it isn't at all bad to drive if you don't $n + 1$ -know that you're sober.

Both this version of the theory and the previous one make the wrong prediction about dubious assertions. On this proposal, it is weirder to assert *p and I don't know p* than it is to assert *p and I don't know that I know p*, and less weird still to assert *p and I don't know that I know that I know that p*. As the number of iterations increase, on this view, the apparent irrationality of your assertion should begin to lessen.

I don't think the data patterns this way. Rather, dubious assertions continue to sound strange at higher levels of iteration: *p and I don't know that I know that I know that p* (compare [Sosa 2009](#)).

At first glance, such conjunctions may seem difficult to assess as they increase in length. [Benton 2013](#) warns that it is important to distinguish 'clashes' from 'clunks'. Perhaps at high levels of iteration, dubious assertions are not infelicitous in the same way as at lower levels of iteration. They may instead simply be unparseable. Perhaps these conjunctions are knowable at some level of processing, but are so difficult to entertain consciously that they are strange to say.

To test this claim, consider dubious assertions as discourses rather than conjunctions. Suppose you say that Queen Elizabeth died in 1603, and I ask whether you are sure. Now imagine you reply that you do in fact know that Queen Elizabeth died in 1603. I can then ask whether you are sure of that, and you can reply that you know it. This pattern can repeat indefinitely without the clunkiness of asserting a conjunction. Yet if you ever acknowledge ignorance in this discourse, your original assertion seems threatened (as emphasized by [Greco 2014b](#)).

- (9) A: When did Queen Elizabeth die?
B: She died in 1603.

- A: How do you know you know that?
 B: I didn't say I know I know it.
 A: So you're saying you don't know you know when Queen Elizabeth died?
 B: I'm not saying that either. I'm saying she died in 1603. Maybe I know that I know she died in 1603, maybe I don't. Honestly, I've got no idea. But you didn't ask about what I know I know, did you? You just asked when she died.

This discourse sounds incoherent. In this way, any admission of higher order ignorance seems to require a retraction of your original assertion. This in turn suggests that dubious assertions are infelicitous at arbitrary levels of iteration.⁸

Another way to test the clunkiness hypothesis concerns interpersonal knowledge attributions. It is not especially clunky to say: it is raining, but Mary doesn't know whether John knows whether Mary knows that it is raining. But this sentence is similar in complexity to: it is raining, but I don't know whether I know whether I know that it is raining. If clunkiness is a matter of syntactic complexity, then if the first sentence is not clunky, neither is the second. (On the other hand, perhaps clunkiness involves semantic rather than syntactic complexity. Even then, however, I don't see why intrapersonal iterations of knowledge should be any more complicated than interpersonal iterations of knowledge.)

I've considered several ways of denying Norm Iteration. I now consider the range of behavior that Norm Iteration applies to. I think that Norm Iteration applies to a broad range of behavior, including assertion, action, inquiry, and subjective certainty. Returning to assertion, notice that Omega Assertion and the weaker knowledge requirement on assertion both follow from the conjunction of Norm Iteration with a truth norm on assertion, which says that you are permitted to assert p only if p is true. For example, the truth norm says that you are permitted to assert the sentence *it is raining* only if it is raining. One application of Norm Iteration implies that you are permitted to assert *it is raining* only if you know it is raining. An infinite further series of applications of Norm Iteration imply that you are permitted to assert *it is raining* only if you omega know it is raining. Now I'll consider the application of Norm Iteration to other domains.

[Williamson 2000](#) and others defend knowledge norms on action. One version of this norm says that you are permitted to act as if p only if you know p . Another version of this norm says that knowing p is not only necessary but also sufficient for being permitted to act as if p .⁹ Norm Iteration implies that if you

⁸On the other hand, one might accept Omega Assertion while denying that omega knowledge is relevant to other actions besides assertion. While Norm Iteration fails in general, it holds in the special case where the behavior is an assertion. When combined with Omega Skepticism, this response would still predict that you are never permitted to assert any ordinary claim.

Yet another response to the argument would appeal to shifts in context. In different contexts, a different number of iterations of knowledge is required for permissible action. If you are in a context where a dubious assertion involving n levels of knowledge is explicitly mentioned, then at least n iterations of knowledge is required for permissible action in that context.

⁹For endorsements of some version of a knowledge norm, see [Hawthorne 2003](#), [Williamson](#)

are permitted to act as if p only if you know p , then you are permitted to act as if p only if you omega know p . If knowledge is distinct from omega knowledge, then Norm Iteration is incompatible with knowing being necessary and also sufficient for permissible action.

Here, I think Norm Iteration potentially makes a good prediction. There are many potential counterexamples to the sufficiency of knowledge for action.¹⁰ In these cases, you know p even though you are not rationally permitted to rely on p in practical reasoning.

One kind of counterexample concerns high stakes.

- (10) **Jellybean.** Hugo is an expert in Roman History, and is participating in a study where the researcher asks him about it. For every correct answer, Hugo gets a jellybean. For every incorrect answer, he receives a painful shock. He can also remain silent, which will result in neither jellybeans nor shocks. (Reed 2010)

In Jellybean, Hugo knows that Caesar was born in 100 BC. But if Hugo is asked whether Caesar was born in 100 BC, it is rational for him to remain silent instead of answer the question. This suggests that he cannot rationally rely on his knowledge in his practical reasoning.

Other counterexamples involve low stakes. Consider Survey:

- (11) **Survey.** You are participating in a survey. Each question has a pair of claims, and you select exactly one true claim from each pair. If you get at least half of the questions right, you get a keychain. The first survey question contains two propositions: (i) Boethius wrote *The Consolations of Philosophy*, and (ii) either $1=1$ or Boethius wrote *The Consolations of Philosophy* (adapted from Beddor 2021 and Roeber 2018).

In Survey, you are rationally permitted to write (ii), and are not permitted to write (i). But you know (i) is true. If you were permitted to rely on (i) in your practical reasoning, then you would be permitted to write (i) instead of (ii). This shows that you can know something without being permitted to rely on it in practical reasoning. Norm Iteration explains what is going on in these cases. Knowledge isn't enough; instead, omega knowledge is required.

Norm Iteration implies that omega knowledge is required for both permissible assertion and permissible action. In this way, Norm Iteration provides a unified explanation of the requirements on assertion and action. See Brown 2010, Montminy 2013, McKenna 2013 and Gerken 2014 among others for further discussion of whether there is a unified norm governing both assertion and action.

Norm Iteration applies not only to action, but also to intellectual inquiry.

2005, Hawthorne and Stanley 2008, Weatherson 2012, Moss 2016, and Weisberg 2013. For endorsements of the sufficiency of knowledge for practical reliance, see Fantl and McGrath 2009 and Ross and Schroeder 2014.

¹⁰See for example Brown 2008a, Brown 2008b, Brown 2012, Gerken 2011, Reed 2010, Lackey 2010, Locke 2015, Roeber 2018, and Beddor 2020a.

Recent research has considered the conditions under which it is rational to stop inquiring into a question. Many claim that the aim of inquiring into a question is to come to know the answer to that question. On this proposal, knowledge stops inquiry: you are permitted to conclude your investigation into a question once you know the answer. Moreover, question-directed attitudes like *curiosity* and *wondering* share this aim: you are permitted to stop wondering about a question once you know the answer.¹¹ This thesis about inquiry connects to a traditional idea about intellectual humility: intellectually humble people are those who acknowledge the limits of their knowledge.

Again, Norm Iteration makes trouble for this idea. If knowledge is required for the permissible cessation of inquiry, then Norm Iteration predicts that knowledge is not sufficient for the permissible cessation of inquiry; instead, omega knowledge is also required. Again, this prediction may be a good one. [Beddor 2021](#) produces counterexamples to the sufficiency of knowledge for the permissible cessation of inquiry:

- (12) **Murine Research.** Mia is a scientist who forms the hypothesis m : Accuphine causes hyperactivity in mice. Mia conducts a number of experiments that support m . Eventually, she conducts enough experiments to know that m is true. But she still is not completely certain of m . One day Mia receives an email from a researcher at another university. Their email announces that they have just completed the most comprehensive study to date on whether Accuphine causes hyperactivity. As a courtesy, they have provided all their data as an attachment.

In Murine Research, Mia does not seem rational to avoid the email. But this suggests that her knowledge does not give her rational permission to stop inquiry.

[Beddor 2020b](#) and [Beddor 2021](#) argue that in the case of both action and inquiry, the problem is that knowledge is fallible. In this book, I'll understand fallibilism as saying that you can know something without being rationally permitted to be certain that it is true. The problem is that when you are not certain of p and know p , you can be required to continue inquiry into p (provided the stakes are right), and you can be required to act differently than you would if you had a complete guarantee of p .

[Beddor 2020a](#), [Beddor 2020b](#), [Beddor 2021](#), and [Goodman and Holguin 2021](#) respond to the cases above by adopting certainty norms on action and inquiry. This is related to the idea that intellectually humble people are those who, even when they may know something, still acknowledge when they are not certain of it. According to the proposal, the states of subjective and epistemic certainty play a crucial role. You are epistemically certain of something when you are permitted to be subjectively certain of it. The state of epistemic certainty reflects what is certain for you, given your evidence. Beddor suggests that epistemic certainty is a stronger state than knowledge. But Beddor takes certainty as primitive, using

¹¹For discussion of various versions of these theses, see among others [Kvanvig 2009](#); [Kappel 2010](#); [Kelp 2011](#), [Kelp 2014](#), [Kelp 2021](#); [Rysiew 2012](#); [Friedman 2013](#), [Friedman 2017](#); [Whitcomb 2017](#); and [Sapir and van Elswyk forthcoming](#).

it to do theoretical work, including to analyze knowledge.

Here, one extension of Norm Iteration would identify epistemic certainty with omega knowledge. On this proposal, something is certain for you just in case you know that you know ... that you know it.

- (13) **Omega Infallibilism.** You are permitted to be subjectively certain of p iff you omega know that p .

The left to right direction of Omega Infallibilism follows from a truth norm on subjective certainty. That is: suppose that you are permitted to be subjectively certain of p only if p is true. Norm Iteration then implies that you are permitted to be subjectively certain of p only if you omega know that p .

Say that you know something for sure when you know it, and you are permitted to be subjectively certain of it. According to Omega Infallibilism, you omega know something exactly when you know it for sure.

The left to right direction of Omega Infallibilism says that rational certainty implies omega knowledge. One piece of evidence in favor of this claim is that it is odd to assert the following:

- (14) I am certain of p , but I don't know that I know ... that I know that p .

For example, if you concede that you don't know that you know that you know that dinosaurs used to walk the Earth, it would be strange to continue to maintain your certainty that dinosaurs used to walk the Earth. Any concession that you lack some iteration of knowledge regarding p requires you to also acknowledge the possibility that p is false.

Conversely, it is hard to imagine a case where someone maintains that they possess every iteration of knowledge regarding some claim, while acknowledging that they are not certain of it. It is fine to say that you know dinosaurs used to walk the Earth, but you aren't certain of it. But what if you also maintain that you know that you know dinosaurs used to walk the Earth, and know that you know this latter fact, and so on? In that case, it is harder to make sense of any further dimension of uncertainty.

The resulting picture embraces fallibilism about knowledge: you can know something without being rationally certain of it. But omega knowledge is different than knowledge: if you omega know something, then you should be certain that it is true.

If omega knowledge is infallible, it avoids the counterexamples to norms on action and inquiry. In each of the above cases, someone knew p without being subjectively certain of p . Since they thought there was some chance that p could fail, they had to take account of this chance in their action and inquiry. By contrast, Omega Infallibilism says that omega knowledge implies you are permitted to be subjectively certain, and so implies that you are permitted to act as if p is true and stop inquiring about whether p is true.

Omega Infallibilism and Omega Assertion imply that assertion is also governed by a certainty norm: assert only what is certain for you (see [Beddor 2020b](#)). This seems right. Following [Unger 1975](#), [Stanley 2008](#) observes that the following

is infelicitous:

(15) It's raining but it's not certain that it's raining.

This is just as strange as *it's raining but I don't know that it's raining*. Yet the knowledge norm on assertion doesn't immediately explain the infelicity of (15). If epistemic certainty is a stronger state than knowledge, then why couldn't you know it was raining while also knowing that it wasn't epistemically certain that it was raining? By contrast, if certainty is omega knowledge, then Omega Assertion immediately explains the infelicity of (15). (15) is felicitous only if (15) is omega known. But this implies you omega know it's raining. But this implies that you know it is certain that it's raining, contradicting the second conjunct. In this way, Omega Infallibilism leads to the correct predictions about the interaction of certainty and assertion.

Omega Infallibilism has downstream consequences for the theory of evidence. I think that your rational credences should match the result of conditionalizing a prior probability distribution on your evidence (see [Williamson 2000](#), [Beddor 2020b](#)). I think that the rational prior probability distribution assigns positive probability to all contingent claims. The result is that you are permitted to be rationally certain of all and only the claims implied by your evidence. Combined with Omega Infallibilism, this means that you omega know p iff p is implied by your evidence. In slogan form: evidence is omega knowledge. This thesis departs from [Williamson 2000](#), who identifies evidence with ordinary knowledge.

Norm Iteration generates a vast array of necessary conditions for permissible action. A further strengthening of Norm Iteration and Omega Infallibilism says that omega knowledge is not only necessary, but also sufficient for the satisfaction of all normative requirements. In the case of action, one could accept:

(16) **Omega Action.** when p is relevant to your decision, you are permitted to rely on p in practical reasoning if and only if you omega know p .

More generally, one could allow that if you omega know p , then you are permitted to act as if p , assert p , stop inquiring about p , and be subjectively certain of p .

Let's take stock. In recent years, 'knowledge first' epistemology has explained many disparate data points in terms of knowledge. This includes all of the puzzles about assertion, action, inquiry, evidence, and belief described above. Yet I've surveyed various reasons to think that mere knowledge is not enough to explain these data points. In the case of assertion, the problem concerned dubious assertions. In the case of action and inquiry, we looked at cases like Survey and Murine Research, where knowledge was also not sufficient. There appear to be reasons, systematized by Norm Iteration, to expect these problems to generalize across the hierarchy of iterated knowledge. To stop the regress, a natural thought is to appeal to omega knowledge. But the problem is that existing theories predict that omega knowledge is trivial or impossible. If KK is true, then omega knowledge is the same thing as knowledge. On the other hand, extant theories of knowledge which deny KK imply that omega knowledge is scarce, because every further iteration of knowledge requires a further power of discrimination.

In the later part of this paper, I'll offer an alternative theory that denies KK, and also avoids the need for infinite powers of perceptual discrimination.

I've now finished exploring my two main arguments against Omega Skepticism. The first argument was that the infelicity of dubious assertions suggests that Omega Assertion is true, which says that omega knowledge is required for permissible assertion. The second argument was that Norm Iteration is true, which implies that omega knowledge is required for permissible behavior more generally. In the course of considering Norm Iteration, I have also explored several strengthenings of Norm Iteration, including Omega Infallibilism, and the thesis that omega knowledge is not only necessary but also sufficient for permissible action.

3 Against KK

I've now presented my case against Omega Skepticism. In the existing literature, the main strategy for avoiding Omega Skepticism has been accepting KK, which identifies knowledge and omega knowledge. In this paper, I take a different route. In section 4, I propose a new principle about knowledge, which makes omega knowledge abundant without identifying omega knowledge and knowledge. One reason that I take this approach is that there are important arguments against KK, which I'll now summarize. In this way, I hope to offer a compromise between the arguments for and against KK.

One challenge for KK is that it is vulnerable to counterexamples. Knowledge requires belief. But you can plausibly know without believing you know. In that case, you don't know that you know. For example, consider the unconfident examinee or 'unwitting historian' (Radford 1966, Feldman 2005). The unconfident examinee studied English history in high school, and retained a bunch of information without remembering the class. If forced to guess, she can reliably identify the year of Queen Elizabeth's death. But she doesn't believe she knows the year, since she has no memory of studying the question. Although she knows when Queen Elizabeth ruled, she doesn't know that she knows this.

Another kind of counterexample to KK involves concept possession. If you know that you know p , then you have the concept of knowledge. Many animals lack the concept of knowledge. But they know things. A dog can know that there is food in his bowl, without knowing he knows this.

Another argument against KK and in favor of Omega Skepticism concerns the connection between knowledge and reliability. Many have thought that knowledge requires reliably true belief. Opponents of KK say that each iteration of knowledge requires an extra level of reliability (Hawthorne and Magidor 2010, p. 387). This line of thought quickly leads to Omega Skepticism: infinite iterations of knowledge require infinitely reliable belief forming processes. This demand can't be satisfied for ordinary claims about the world. Consider even the claim that you have hands. Your perceptual faculties reliably tell you that you have hands. But your perception isn't infinitely reliable. We can imagine some possible state of affairs where your perception of your hands is misleading. This

state of affairs would be very strange. But between the actual state of affairs and that strange one, we can imagine a long chain of states of affairs, each slightly stranger than the actual state of affairs. Each member of the chain is possible by the lights of the previous one. Infinitely reliable perception would demand that your perceptual faculties perform accurately in every state of affairs in the chain.

To make these ideas more precise, consider the thesis that knowledge is constrained by a ‘margin for error’. On this view, you know p only if p could not easily have been false (Williamson 1992). Margins for error characterize cases of inexact knowledge. Imagine you are looking at a tree that appears to be 100 feet tall. Your knowledge of the tree’s height is inexact. You know that the tree’s height falls in some interval around 100 feet; but you do not know that it is exactly 100 feet tall. How much you know about the tree’s height depends on how tall the tree is. If the tree is 100 feet tall, then you know a lot about the tree’s height. But if the tree is 90 feet tall, you know less. This suggests something like the following:

- (17) **Margin for Error.** For any height x , if you know that the tree is not $x - 1$ feet tall, then the tree is not x feet tall.

Margin for Error leads to the failure of KK. If KK is true and you know Margin for Error, then you don’t know anything about the tree’s height. If you know anything about the tree’s height, then there must be some n where you know the tree is not $n - 1$ feet tall. But if KK holds, then you know that you know that the tree is not $n - 1$ feet tall. But now suppose you know Margin for Error. Then you can deduce that the tree is not n feet tall. After all, here are two things you would know: first, that you know the tree is not $n - 1$ feet tall; second, that if you know the tree is not $n - 1$ feet tall, then the tree is not n feet tall. You can thereby know by deduction that the tree is not n feet tall. By KK, you know that you know that the tree is not n feet tall. Extending this reasoning, you know for any x that the tree is not x feet tall.

The validity of Margin for Error also leads to Omega Skepticism. If Margin for Error is valid, then you should be able to omega know it. But if you omega know Margin for Error, then you don’t omega know anything about the tree’s height. If you omega know anything about the tree’s height, then there must be some n where you omega know the tree is not $n - 1$ feet tall. But now suppose you omega know Margin for Error. Then you omega know that you know Margin for Error. So two things you omega know imply that the tree is not n feet tall. In particular, you omega know that you know the tree is not $n - 1$ feet tall, and you omega know that if you know the tree is not $n - 1$ feet tall, then the tree is not n feet tall. Since you omega know these two things, you also omega know the tree is not n feet tall. Iterating this reasoning, you omega know for any x that the tree is not x feet tall. But then something you omega know would be false.

I’ve considered two arguments against KK. The first argument is that KK is vulnerable to counterexamples, like the unconfident examinee. The second

argument is that KK makes bad predictions about inexact knowledge, because it is incompatible with Margin for Error.

Zooming out, we're left with a problem. First, dubious assertions and Norm Iteration both suggest that omega knowledge matters. But the main existing theory of omega knowledge is KK, and KK faces problems. In the rest of this paper, I offer a new theory of omega knowledge that allows for abundant omega knowledge while avoiding the problems facing KK.

4 Reflective Luminosity

You have reflective knowledge iff you know that you know. A mental state is luminous iff whenever you are in the state, you know you are in the state. This section of the paper explores the principle that reflective knowledge is luminous:

- (18) **Reflective Luminosity.** If you know that you know that p , then you omega know that p .

Reflective Luminosity says that reflective knowledge is the same thing as omega knowledge. If you know that you know that p , then you know that you know that you know that p . You also know that you know that you know that you know that p . If you know that you know that p , then you possess every iteration of knowledge that p .

Reflective Luminosity allows that ordinary knowledge is not luminous, so that you can know p without knowing that you know p . Reflective Luminosity says that there is a special kind of knowledge that is luminous: knowledge about your own knowledge. In this way, Reflective Luminosity distinguishes ordinary knowledge from reflective knowledge.

Reflective Luminosity says that there are at most two levels of knowledge. Either you know p without knowing that you know p , or you omega know p . There is knowledge, and then there is reflective knowledge. There is no other kind of knowledge.

Reflective Luminosity is weaker than the KK principle, which says that all knowledge is omega knowledge. Instead, Reflective Luminosity says that reflective knowledge is omega knowledge.

[Williamson 2000](#) claims that knowledge is a mental state. If knowledge is a mental state, then there is a difference in subject matter between your knowledge that it is raining, and your knowledge that you know it is raining. The first state is knowledge about something that is not a mental state; the second state is knowledge about your mental states. Reflective Luminosity follows from the idea that knowledge of your own mental states is special. Ordinary knowledge is not luminous; but knowledge of your own mental states is luminous.

Omega Skepticism says that every iteration of knowledge is a further cognitive achievement. Reflective Luminosity disagrees. It says that knowing p is one cognitive achievement, and reflectively knowing p is a further cognitive achievement. But knowing that you reflectively know p requires no further cognitive achievement than reflectively knowing p .

In section 2, I objected to Omega Skepticism, arguing that omega knowledge plays several important roles. I'll now argue that Reflective Luminosity vindicates those arguments.

Consider assertion. Recall that dubious assertions are sentences of the form *p* and *I don't know that I know p*. Extant accounts that deny KK predict that these claims can be known. Such accounts also say that knowledge is the norm of assertion. Such theories therefore predict that dubious assertions satisfy the primary norm on assertion.

Reflective Luminosity can explain the infelicity of dubious assertions without validating KK. Suppose again that Omega Assertion is true, so that an assertion of *p* is permissible only if you omega know *p*. Dubious assertions are infelicitous because they cannot be omega known, and so (given Reflective Luminosity) cannot be reflectively known. For example, consider *p* and *I don't know that I know that I know p*. Reflective Luminosity implies that this cannot be reflectively known. For suppose you reflectively knew it. Then you'd know that you know *p*. So by Reflective Luminosity you'd know that you know that you know *p*. But by factivity, you wouldn't know that you know that you know *p*.

In Section 2, I suggested that omega knowledge is sufficient for permissible assertion. In that case, Reflective Luminosity says that you may assert *p* once you know that you know *p*. In this way, Reflective Luminosity ensures that if reflective knowledge is abundant, then permissible assertion is too.

In Section 2, I suggested that assertion is merely one of a wide class of behaviors governed by omega knowledge. I introduced the principle of Norm Iteration, which says that if you are permitted to *A* only if *p*, then you are permitted to *A* only if you know that *p*. I showed that Norm Iteration implies that if you are permitted to *A* only if *p*, then you are permitted to *A* only if you omega know *p*. In this way, Norm Iteration implies that some kind of omega knowledge is required for any permissible behavior.

If Reflective Luminosity is true, then Norm Iteration only generates norms with two applications. If you are permitted to *A* only if *p*, then you are also permitted to *A* only if you know *p*, and only if you know that you know *p*. But once you satisfy this last condition, you automatically count as omega knowing *p*, and so are guaranteed to satisfy all of the normative requirements imposed by Norm Iteration.

In Section 2, I considered how Norm Iteration would apply to action, inquiry, and certainty. In the case of action, I suggested that you are permitted to act as if *p* only if you omega know *p*. In the case of inquiry, I suggested that you are permitted to stop inquiring about *p* only if you omega know *p*. In the case of certainty, I suggested that you are permitted to be subjectively certain of *p* (if and) only if you omega know *p*. If Reflective Luminosity is correct, then reflective knowledge guarantees omega knowledge. The result would be that action, inquiry, and certainty all require reflective knowledge, but may not require anything more than that. On this view, once you reflectively know *p*, you are perfectly safe to act as if *p*, to stop inquiring about *p*, and to be certain of *p*. Reflective knowledge shields you from any chance of error.

On the other hand, when you reflectively know *p*, you know *p*. For this

reason, my account is more stringent than an ordinary knowledge account. I claim that when you know p without knowing you know p , you should not rely on your knowledge. On this picture, reflective knowledge is infallible, while ordinary knowledge is not.

In section 3, I argued against KK. I argued that KK was subject to counterexamples, and makes bad predictions about inexact knowledge. Reflective Luminosity avoids these problems.

Consider again the unconfident examinee, who knows the year of Queen Elizabeth's death, without knowing that she knows. Reflective Luminosity is compatible with the failure of KK. So Reflective Luminosity is compatible with cases like the unconfident examinee.

Likewise, Reflective Luminosity is compatible with an animal knowing there is food in his bowl, without knowing that he knows. Reflective Luminosity says however that if an animal can successfully deploy the concept of knowledge to know that he knows something, he thereby possesses every iteration of knowledge. In this way, Reflective Luminosity allows exactly the failures of KK that are pre-theoretically compelling.

Now I'll argue that Reflective Luminosity offers a satisfying treatment of Margin for Error.

Again imagine you are looking at a tree that appears to be 100 feet tall. For some height n , you know that the tree is at least n feet tall. This knowledge and KK are incompatible with knowledge of Margin for Error. Again, suppose that you know the tree is not 89 feet tall. By KK, you know that you know this. If you know Margin for Error, it follows that you know that the tree is not 90 feet tall. So you know that you know the tree is not 90 feet tall. This reasoning leads to the result that you can rule out every possible height. In this way, knowledge of Margin for Error counts against KK.

Reflective Luminosity blocks this argument. Reflective Luminosity allows you to know the tree is not 89 feet tall, without knowing that you know this.

While Reflective Luminosity avoids the argument above, Reflective Luminosity faces the threat of revenge. If Reflective Luminosity is true, then anyone who knows they know Margin for Error cannot know that they know anything about the tree's height. For suppose you know that you know anything about the tree's height. Then there must be some n where you know that you know the tree is not $n - 1$ feet tall. But if Reflective Luminosity holds, then you know that you know that: you know that the tree is not $n - 1$ feet tall. But now suppose you know that you know Margin for Error is true. Then your known knowledge implies that the tree is not n feet tall. Since known knowledge is closed under deduction, it follows that you know you know the tree is not n feet tall. Extending this reasoning, you can know for any x that you know the tree is not x feet tall.

KK forbids knowing that Margin for Error is true, while Reflective Luminosity merely forbids knowing that one knows Margin for Error is true. This disanalogy is important. Reflective Luminosity cannot validate Margin for Error unrestrictedly, at the risk of leading to knowledge of knowledge of Margin for Error. But Reflective Luminosity can nonetheless explain the appeal of Margin

for Error.

If Margin for Error were valid, it would be possible to know that you know it. So Reflective Luminosity rules out the validity of Margin for Error. But Reflective Luminosity is compatible with Margin for Error having other positive epistemic statuses that explain its appeal, short of validity.

In the accompanying book, *Omega Knowledge: What it is and Why it Matters*, I develop an alternative conception of the appeal of philosophical principles like Margin for Error. This alternative conception involves the concepts of normality, justification, and what you know in the good case where your belief forming processes are functioning reliably.

Here is the idea. Epistemology studies general principles governing knowledge. These principles have often been interpreted as necessary universal generalizations. But another way of thinking of these principles is as rules that are normally true: default modes of inference involving knowledge (Reiter 1980). We can vindicate the appeal of Margin for Error if we can predict that Margin for Error is normally true. And something can be normally true without anyone being able to know that they know it.

In the accompanying book, I develop a model of knowledge that validates Reflective Luminosity, and which also validates the thesis that Margin for Error is normally true. This can explain philosophical judgments about knowledge and error, if we interpret these judgments as part of a philosophical practice of investigating default rules of inference rather than necessary universal generalizations.

The success of this strategy depends on what normality is. In the book, I understand normality by connecting it to justification and knowledge. We can distinguish different situations based on how favorable they are for knowing. In good cases, you know quite a lot; in bad cases, you know little. For example, in cases of perceptual knowledge, how much you know might depend on the extent to which reality matches how things appear. When they match, you are in the good case; when they diverge, you are in a bad case.

We can use this idea to enrich the concept of normality. First, something is normally true when you are justified in believing it regardless of whether you're in the good case or the bad case. In this sense, you have a default entitlement to make certain inferences on the basis of your evidence, regardless of whether you are forming beliefs in unfavorable conditions. Second, something is normally true when you would know it if you were in the good case.

In the accompanying book, I develop a model of knowledge that precisely articulates the ideas above. The model predicts that in the good case, you know Margin for Error; but in bad cases, Margin for Error sometimes fails. No matter what case you are in, you are justified in believing Margin for Error.

I claim that you can know Margin for Error is true in the good case. While this is consistent with Reflective Luminosity, it is inconsistent with KK. (KK would imply that you omega know Margin for Error in the good case, which would imply that in the good case you know nothing about the tree's height.) In this way, Reflective Luminosity but not KK is consistent with Margin for Error being normally true.

5 Conclusion

In this paper, I've argued that omega knowledge matters. In addition, I argued against KK, the leading theory of knowledge that allows for omega knowledge. I argued that KK was subject to counterexamples, and failed to explain the appeal of Margin for Error.

In its place, I proposed an alternative to KK: if you know that you know, then you omega know. I argued that this alternative allows for abundant omega knowledge, but also avoids counterexamples to KK, and can explain the appeal of Margin for Error.

In the accompanying book, I develop two more theories of omega knowledge. The first theory is defined by the principle of Fragility, which says that if you know p , then it is epistemically possible for you that you omega know p (so that you don't know that you don't omega know p). The second theory replaces Margin for Error with a weaker principle, Variable Margins. Variable Margins allows that your knowledge of a tree's height is not governed by a fixed margin. Instead, it says that for every height x , there is some margin m greater than 0, where if you know that the tree is not $x - m$ feet tall, then the tree is not x feet tall. I then critically compare these three theories of omega knowledge. I show that each theory has its significant own costs and benefits. I tentatively suggest that one of Reflective Luminosity or Fragility should be accepted. But I don't settle on which principle is better.

References

- J. Adler. *Belief's Own Ethics*. MIT Press, 2002.
- Jonathan E. Adler. Skepticism and universalizability. *Journal of Philosophy*, 78(3):143–156, 1981.
- William P. Alston. Level-confusions in epistemology. *Midwest Studies in Philosophy*, 5(1):135–150, 1980. doi: 10.1111/j.1475-4975.1980.tb00401.x.
- Bob Beddor. Certainty in action. *Philosophical Quarterly*, 70(281):711–737, 2020a. doi: 10.1093/pq/pqaa006.
- Bob Beddor. New work for certainty. *Philosophers' Imprint*, 20(8), 2020b.
- Bob Beddor. Knowledge and the aim of inquiry. September 2021.
- Matthew Benton. Dubious objections from iterated conjunctions. *Philosophical Studies*, 162:355–358, 2013.
- Jessica Brown. Subject-sensitive invariantism and the knowledge norm for practical reasoning. *Noûs*, 42(2):167–189, 2008a. doi: 10.1111/j.1468-0068.2008.00677.x.
- Jessica Brown. Knowledge and practical reason. *Philosophy Compass*, 3(6): 1135–1152, 2008b. doi: 10.1111/j.1747-9991.2008.00176.x.

- Jessica Brown. Knowledge and assertion. *Philosophy and Phenomenological Research*, 81(3):549–566, 2010. doi: 10.1111/j.1933-1592.2010.00369.x.
- Jessica Brown. Practical reasoning, decision theory and anti-intellectualism. *Episteme*, 9(1):1–20, 2012. doi: 10.1017/epi.2011.4.
- Sam Carter. Higher order ignorance inside the margins. *Philosophical Studies*, pages 1–18, 2018.
- Stewart Cohen and Juan Comesaña. Williamson on gettier cases and epistemic logic. *Inquiry*, 56:15–29, 2013.
- Nilanjan Das and Bernhard Salow. Transparency and the kk principle. *Noûs*, 52(1):3–23, 2018.
- Keith DeRose. Knowledge, assertion and lotteries. *Australasian Journal of Philosophy*, 74(4):568–580, 1996. doi: 10.1080/00048409612347531.
- Keith DeRose. Assertion, knowledge, and context. *Philosophical Review*, 111(2): 167–203, 2002. doi: 10.1215/00318108-111-2-167.
- Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press, 2009.
- Jérôme Dokic and Paul Égré. Margin for error and the transparency of knowledge. *Synthese*, 166(1):1–20, 2009. doi: 10.1007/s11229-007-9245-y.
- Pascal Engel. In what sense is knowledge the norm of assertion? *Grazer Philosophische Studien*, 77(1):45–59, 2008.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, 2009.
- Richard Feldman. Respecting the evidence. *Philosophical Perspectives*, 19, 2005.
- Jane Friedman. Question-directed attitudes. *Philosophical Perspectives*, 27(1): 145–174, 2013. doi: 10.1111/phpe.12026.
- Jane Friedman. Why suspend judging? *Noûs*, 51(2):302–326, 2017. doi: 10.1111/nous.12137.
- Mikkel Gerken. Warrant and action. *Synthese*, 178(3):529–547, 2011. doi: 10.1007/s11229-009-9655-0.
- Mikkel Gerken. Same, same but different: The epistemic norms of assertion, action and practical reasoning. *Philosophical Studies*, 168(3):725–744, 2014. doi: 10.1007/s11098-013-0156-0.
- Simon Goldstein. *Omega Knowledge: What it is and Why it Matters*. Oxford University Press, Forthcoming.
- Jeremy Goodman and Ben Holguin. Thinking and being sure. August 2021.

- Daniel Greco. Could *kk* be *ok*? *Journal of Philosophy*, 111(4):169–197, 2014a. doi: 10.5840/jphil2014111411.
- Daniel Greco. Iteration and fragmentation. *Philosophy and Phenomenological Research*, 88(1):656–673, 2014b.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, 2003.
- John Hawthorne and Ofra Magidor. Assertion and epistemic opacity. *Mind*, 119(476):1087–1105, 2010.
- John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008. doi: 10.5840/jphil20081051022.
- Klemens Kappel. On saying that someone knows: Themes from craig. In Adrian Haddock, Alan Millar, and Duncan Pritchard, editors, *Social Epistemology*. Oxford University Press, 2010.
- Christoph Kelp. What’s the point of “knowledge” anyway? *Episteme*, 8(1):53–66, 2011. doi: 10.3366/epi.2011.0006.
- Christoph Kelp. Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51(3):227–32, 2014.
- Christoph Kelp. *Inquiry, Knowledge, and Understanding*. Oxford University Press, 2021.
- Jonathan Kvanvig. The value of understanding. In Pritchard, Haddock, and Millar, editors, *Epistemic Value*, pages 95–112. Oxford: Oxford University Press, 2009.
- Jennifer Lackey. Acting on knowledge. *Philosophical Perspectives*, 24(1):361–382, 2010. doi: 10.1111/j.1520-8583.2010.00196.x.
- Dustin Locke. Practical certainty. *Philosophy and Phenomenological Research*, 90(1):72–95, 2015. doi: 10.1111/phpr.12036.
- Berislav Marušić. The self-knowledge gambit. *Synthese*, 190(12):1977–1999, 2013. doi: 10.1007/s11229-011-9949-x.
- Conor McHugh. Self-knowledge and the *kk* principle. *Synthese*, 173(3):231–257, 2010. doi: 10.1007/s11229-008-9404-9.
- Robin McKenna. Why assertion and practical reasoning are possibly not governed by the same epistemic norm. *Logos and Episteme*, 4(4):457–464, 2013. doi: 10.5840/logos-episteme2013446.
- Martin Montminy. Assertion and practical reasoning. *Pacific Philosophical Quarterly*, 94:57–68, 2013. doi: <https://doi.org/10.1111/j.1468-0114.2012.01444.x>.

- Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, 2016.
- Colin Radford. Knowledge-by examples. *Analysis*, 27:1–11, 1966.
- Baron Reed. A defense of stable invariantism. *Noûs*, 44(2):224–244, 2010. doi: 10.1111/j.1468-0068.2010.00738.x.
- Ray Reiter. A logic for default reasoning. *Artificial Intelligence*, 13(1-2):81–137, 1980. doi: 10.1016/0004-3702(80)90014-4.
- Blake Roeber. The pragmatic encroachment debate. *Noûs*, 52(1):171–195, 2018. doi: 10.1111/nous.12156.
- Jacob Ross and Mark Schroeder. Belief, credence, and pragmatic encroachment1. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014. doi: 10.1111/j.1933-1592.2011.00552.x.
- Patrick Rysiew. Epistemic scorekeeping. In Jessica Brown and Mikkel Gerken, editors, *Knowledge Ascriptions*. Oxford University Press, 2012.
- Bernhard Salow. The externalist’s guide to fishing for compliments. *Mind*, 127 (507):691–728, 2018. doi: 10.1093/mind/fzw029.
- Yasha Sapir and Peter van Elswyk. Hedging and the ignorance norm on inquiry. *Synthese*, pages 1–23, forthcoming. doi: 10.1007/s11229-021-03048-0.
- Jonathan Schaffer. Knowledge in the image of assertion. *Philosophical Issues*, 18(1):1–19, 2008. doi: 10.1111/j.1533-6077.2008.00134.x.
- David Sosa. Dubious assertions. *Philosophical Studies*, 146(2):269–272, 2009.
- Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128: 169–199, 2006.
- Robert Stalnaker. Luminosity and the kk thesis. In Sanford Goldberg, editor, *Externalism, Self-Knowledge, and Skepticism: New Essays*, pages 19–40. Cambridge University Press, 2015.
- Jason Stanley. *Knowledge and Practical Interests*. Oxford University Press, 2005.
- Jason Stanley. Knowledge and certainty. *Philosophical Issues*, 18(1):35–57, 2008. doi: 10.1111/j.1533-6077.2008.00136.x.
- John Turri. Knowledge in the image of assertion. *Australasian Journal of Philosophy*, 89(1):37–45, 2011.
- John Turri. Knowledge and the norm of assertion: A simple test. *Synthese*, 192 (2):385–392, 2015. doi: 10.1007/s11229-014-0573-4.
- Peter Unger. *Ignorance: A Case for Scepticism*. Oxford University Press, 1975.

- Brian Weatherson. Knowledge, bets, and interests. In Jessica Brown and Mikkel Gerken, editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press, 2012.
- Jonathan Weisberg. Knowledge in action. *Philosophers' Imprint*, 13, 2013.
- Dennis Whitcomb. One kind of asking. *Philosophical Quarterly*, 67(266), 2017. doi: 10.1093/pq/pqw027.
- Timothy Williamson. Inexact knowledge. *Mind*, 101(402):217–242, 1992. doi: 10.1093/mind/101.402.217.
- Timothy Williamson. Knowing and asserting. *Philosophical Review*, 105(4):489, 1996. doi: 10.2307/2998423.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.
- Timothy Williamson. Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *Philosophical Quarterly*, 55(219):213–235, 2005. doi: 10.1111/j.0031-8094.2005.00396.x.
- Timothy Williamson. Response to cohen, comesaña, goodman, nagel, and weatherson on gettier cases in epistemic logic. *Inquiry*, 56(1):77–96, 2013.