Iterated Knowledge

1 Omega Knowledge Matters

1.1 Overview

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 omega know p when you p-know p when you know that you p-know p you omega know p when you p-know p for every p.

There are two existing approaches to omega knowledge. One view, KK, says that you omega know everything you know.

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

The other view, Omega Skepticism, says that you omega know almost nothing.

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

Much of the debate between KK and Omega Skepticism concerns the thesis that 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.

Margin for Error quickly leads to the failure of KK, and the truth of Omega Skepticism.

This book rejects both KK and Omega Skepticism. Against Omega Skepticism, in Chapter 1 I argue that omega knowledge is necessary for permissible assertion and action (§1.1). To do so, I explore and defend two theses about knowledge:

- (4) **Omega Assertion.** You are permitted to assert something only if you omega know it.
- (5) **Norm Iteration.** If you are permitted to do A only if p, then you are permitted to do A only if you know p.

Against KK, I argue that omega knowledge isn't the same thing as knowledge (§1.2). To achieve this balance, I consider several weakenings of Margin for Error.

I develop three theories that make room for omega knowledge without accepting KK. Each theory has its own costs and benefits. Chapter 2 introduces the principle of Reflective Luminosity:

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.

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

This 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.

In Chapter 3, I introduce an alternative account of omega knowledge, Fragility. Say that p is true 'for all you know' iff you don't know that p is false iff it is 'epistemically possible', or 'consistent with your knowledge', that p. Then:

(7) **Fragility.** If you know p, then for all you know, you omega know p.

I'll show that this principle also avoids Omega Skepticism, and is also weaker than KK.

In Chapter 4, I consider a third approach to omega knowledge. This approach avoids Omega Skepticism by replacing Margin for Error with a weaker principle. The weaker principle denies that your knowledge of a tree's height is governed by a fixed margin. Rather, the weaker principle says:

(8) **Variable Margins.** 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.

In Chapter 4, I show how this weakening of Margin for Error can avoid Omega Skepticism, and explain why omega knowledge is important.

By introducing these three theories 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.

Throughout the book, I explore theories of justification as well as knowledge. I consider the following principles about justification:

- (9) **Possible Permission.** You are justified in doing something if and only if for all you know, you are permitted to do it.
- (10) **Possible Omega Knowledge.** You are justified in believing something if and only if for all you know, you omega know it.
- (11) **Possible Knowledge.** You are justified in believing something if and only if for all you know, you know it.

I'll show that in the presence of principles like these, the theories of knowledge I develop have surprising consequences for the theory of justification.

The principles above are the main focus of the book. In Chapter 5, I develop several mathematically precise models of knowledge and justification that draw out the consequences of these various principles in detail. To do so, I model knowledge in terms of the normality of your belief forming processes. In addition, I look in detail at the particular case of perceptual knowledge, where what you know depends on how things appear, and how closely reality matches appearance.

All of Reflective Luminosity, Fragility, and Variable Margins can avoid Omega Skepticism without implying KK. But each principle has its own costs and benefits. In Chapter 6, I compare the various approaches. Although all three principles are compatible with one another, I argue that the best theory of knowledge should accept only one of them. All three principles have the same central benefit: offering an explanation of why Omega Skepticism fails. But each principle has its own costs; so accepting more than one of the principles would lead to an accumulation of cost without a corresponding accumulation of benefit. I outline the costs and benefits of each principle, and leave it to the reader to decide which if any to accept.

1.2 Against Omega Skepticism

I'll begin by arguing against Omega Skepticism. 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.

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

(12) **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:

(13) 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.

Lotteries provide another reason to accept a knowledge norm. It is strange to assert on statistical grounds that a fair lottery ticket will lose. Williamson 2000 and others argue that you don't know on statistical grounds that the ticket will lose. If you don't have lottery knowledge, then the knowledge norm of assertion explains why you can't assert that your ticket will lose. Omega Assertion makes the same prediction, since whenever you fail to know your ticket will lose, you also fail to know that you know your ticket will lose.

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.

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

- (14) 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 (14-a) and (14-b) are strange to assert. You can't omega know (14-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) Stand to attention!—and I don't know whether I have authority to order you to stand to attention.

I agree, but I think this needs to be explained. I think this imperative is weird to assert because imperatives are governed by a knowledge norm: command only if you (omega) know you have authority. But imperatives have their own form of dubious assertion. It is strange to command someone to stand to attention while acknowledging that you don't know whether you know that you have the authority to do so. Abundant theories of omega knowledge can explain why these claims are bad. (See Dorst 2019 for a similar response.)

Mandelkern 2021 and Mandelkern and Dorst forthcoming argue that neither Moorean sentences nor dubious assertions provide good evidence for a knowledge norm on assertion. Building on Silk 2015, Mandelkern 2021 observes that it is strange to say:

(ii) Open the door, but I don't know whether you will.

Generalizing, it is strange for me to order you to do something while asserting that I don't know whether you will (or that I don't know whether I know you will). But Mandelkern 2021 argues that there is not a knowledge norm on orders. For example, if you are kidnapping me, I can order you to let me go, even when I obviously don't know that you will do so.

Mandelkern 2021 and Mandelkern and Dorst forthcoming explain Moorean assertions by embracing 'posturing' rules. In asserting p, you should act as if you are absolutely certain of p. In ordering someone to A, you should act as if you are absolutely certain of p. These norms don't require you to actually be certain; they simply require you to pretend to be certain when you perform speech acts.

I reject this proposal, in part because I think that orders and assertions place different normative requirements on the speaker. First, return to the kidnapping case. While it is appropriate to order the kidnapper to release you, it is stranger to assert to the kidnapper that he will release you. The posturing account doesn't explain this. Second, it is often natural to reply to an assertion by asking how the speaker knew what they said. For example, if I say that the living room door is closed, you can ask how I know that. Moreover, once I concede that I don't know that the living room door is closed, then I cannot go on to assert that the living room door is closed. By contrast, it is bizarre to reply to an order by asking how the speaker knew that the order would be followed. For example, if I tell you to close the living room door, you cannot ask me how I knew that you would close the door. (On the other hand, it may be fine for you to ask what makes me think you will comply.)

If ordering and asserting were governed by the same posturing norm, then they would not

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 2013b for alternative explanations. Williamson 2013b argues that dubious assertions are similar to the following imperative:

Consider more complex dubious assertions, involving mixed attitudes of belief and knowledge:

- (15) a. p but I doubt that I know that p.
 - b. p but I believe that I don't know that p.
 - c. p but I am not justified in believing that I know that p.
 - d. p but I am justified in believing that I don't know that p.

These are all infelicitous. Omega Assertion rules out omega knowing any of these conjunctions. If you omega know the whole conjunction, then you omega know p. But this precludes the truth of the second conjunct. For example, imagine that you omega know that you have hands, but that you doubt that you know that you have hands. This is inconsistent: if you omega know you have hands, then you know you know you have hands, and so you don't doubt that you have hands.⁵

differ in these ways. I think that assertions and orders are governed by different norms. In particular, I think the norm on ordering is weaker than the norm on assertion. For example, one way to explain the infelicity of Moorean orders is to embrace the norm that you can give an order only if for all you know, you omega know that it will be followed. When you say a Moorean order, you assert that you don't know that the order will be followed. But by Omega Assertion, this requires you to know that you don't omega know p. This would violate the norm on ordering.

This weaker norm on ordering explains why it is strange to respond to an order by asking the speaker how he knows that the order will be followed. When the speaker orders the addressee, the speaker does not represent himself as knowing the order will be followed. Instead, he merely represents himself as possibly omega knowing that the order will be followed.

Consider again the kidnapping case. I order you to let me go, even though I obviously don't know that you will release me. In this case, I plausibly violate the proposed norm on ordering: plausibly, I know that I don't know that I will be released. But this does not imply that ordering isn't governed by the norm. Rather, I think that in this case, the speaker is pretending to satisfy the norm on ordering. (See Hawthorne et al. 2016 for a related proposal, and Mandelkern and Dorst forthcoming p. 21 for critical discussion.) In general, when a speech act is governed by a norm, speakers may sometimes violate the norm while pretending to satisfy it.

⁵Omega Assertion similarly explains the infelicity of 'ignorance conditionals' of the form If I don't know p, then p (from Dorst 2019):

- (i) a. If I don't know Padua is in Italy, Padua is in Italy.
 - b. Even if I don't know Padua is in Italy, Padua is in Italy.
 - Whether or not I know Padua is in Italy, Padua is in Italy.

Many have argued that indicative conditionals presuppose that their antecedent is possible. (See for example Stalnaker 1975 and Willer 2017. I assume that the relevant presupposition is also an entailment.) On one version of this theory, an utterance of an indicative conditional is true only if the antecedent is epistemically possible for the speaker. It follows that you omega know the ignorance conditional if I don't know that p, then p only if you omega know that for all you know, you don't know p.

It follows that ignorance conditionals cannot be omega known. Here's why. Suppose you omega know that if you don't know p, then p. Trivially, you omega know (the material conditional) that if you know p, then p. These two conditionals imply p, and you omega know both conditionals. So by deduction you can omega know p. This contradicts the claim that you omega know that for all you know, you don't know p. Summarizing: ignorance conditionals can't be omega known, because omega knowing the ignorance conditional would imply omega knowing that the antecedent is impossible. In this way, the infelicity of ignorance conditionals

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 as I'll discuss in greater detail in Chapter 5, 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 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.⁶ One of the contributions of this book is to develop weakenings of KK that preserve much of its appeal while avoiding some of its costs. Throughout the book, I argue 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.

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:

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

is analogous to the infelicity of the dubious assertion p but I don't know that I know p.

⁶See for example Cohen and Comesaña 2013, Greco 2014a, and Das and Salow 2018.

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 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.

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

⁷The 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.

Suppose that you don't know p. On the narrow scope interpretation, it follows that you aren't permitted to assert p. But on the wide scope interpretation, this does not follow. On the other hand, the wide scope norm immediately predicts that it is impermissible to both assert p and fail to know p. This does not follow immediately from the narrow scope norm.

I won't take a stand on how best to interpret normative requirements in this book. But there is at least one kind of case where the wide scope requirement makes more natural predictions. Imagine that you are permitted to swim at the beach only if the lifeguard is on duty. You haven't checked whether there is a lifeguard today, but decide not to go. In response to a beach-loving friend, it would be bizarre to explain: I am not permitted to swim today, because I don't know that the lifeguard is on duty. But on the narrow scope interpretation, Norm Iteration has this consequence. By contrast, on the wide scope interpretation, Norm Iteration merely implies that you are not permitted to swim while remaining ignorant of the lifeguard's status. Thanks to Dmitri Gallow and Dan Greco for discussion.

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 could respond in a few ways. One response is internalist about permission. On this proposal, the only conditions on permissible behavior appeal to facts about your own mental states. You have omega knowledge about your own mental state; but you don't have omega knowledge about the external world. For example, the internalist may say that you are permitted to drive home only if you justifiably believe that you aren't sober. Actual sobriety has nothing to do with it. Similarly, you are permitted to press the button only if you justifiably believe that the tree is at least 80 feet tall. Norm Iteration then implies that you are permitted to drive home only if you omega know that you justifiably believe that you aren't sober. But this does not require omega knowledge of ordinary claims about the world. Instead, it only requires omega knowledge of facts about what you are justified in believing. (In Chapter 5, I'll review a model of Omega Skepticism that licenses omega knowledge about justification while rejecting omega knowledge about the external world.)

Similarly, the internalist may deny that permissible assertion requires truth; instead, they may say that you are permitted to assert p only if you are justified in believing p. Norm Iteration would then imply that you are permitted to assert p only if you omega know that you are justified in believing p. Later in this chapter, I'll consider theories of justification in terms of knowledge. According to one such theory, you are justified in believing something only if it is possible (in other words, consistent with what you know) that you know it. This sort of theory allows a justified belief condition on assertion to explain the data points that motivate a knowledge condition. Moorean conjunctions and lottery propositions would be infelicitous because the speaker knows that the conjunction is not known.

Another response to Norm Iteration replaces knowledge with another state. For example, a variant of Norm Iteration says that if you are permitted to A only if p, then you are permitted to A only if you justifiedly believe p. Iteration of this principle would then imply that omega justification rather than omega knowledge is required for permissible action. At the end of this chapter, I sketch a rival theory of justification, according to which you are justified in performing a behavior iff for all you know, you are permitted to A.

Another response restricts 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.⁸

⁸See Marušic 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 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 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 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 unparsable. 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

some senses of 'permission' are irrelevant to deliberation about what action to perform.

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).

- (17) 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.⁹

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

⁹On 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.

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. Norm Iteration implies that if you 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. ¹¹ 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.

(18) **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:

(19) **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).

¹⁰For endorsements of some version of a knowledge norm, see Hawthorne 2003, Williamson 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.

 $^{^{11}\}mathrm{See}$ for example Brown 2008a, Brown 2008b, Brown 2012, Gerken 2011, Reed 2010, Lackey 2010, Locke 2015, Roeber 2018, and Beddor 2020a.

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. 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. ¹² 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:

(20) **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

¹²For 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.

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 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.

(21) **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:

(22) 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.¹³

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.

I will tentatively take Omega Infallibilism on board in much of the book, and explore the theory that results. The resulting picture embraces fallibilism about

¹³Thanks to an anonymous referee for help here.

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:

(23) 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 (23). 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 (23). (23) is felicitous only if (23) 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:

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. At the end of this chapter, I propose a general theory of justification which

coheres with the view that omega knowledge is sufficient for permissibility. According to that theory, you are justified in A iff for all you know, you are permitted to A. I use this theory to explore a thesis about justified belief: that you are justified in believing p iff for all you know, you omega know p. This argument relies on the assumption that omega knowledge is necessary and sufficient for permissible belief.

Omega knowledge is luminous: when you omega know, you know you omega know. ¹⁴ If omega knowledge is sufficient for permissibility, it follows that permission is luminous. If you are permitted to perform a behavior, whether it is action, assertion, inquiry, or subjective certainty, then you know you are permitted to perform that behavior.

On the other hand, I deny negative introspection. I think you can fail to omega know p without knowing that you fail to omega know p. For this reason, I think that impermissibility is not luminous. There are cases where you are not permitted to behave in some way, even though for all you know you are permitted to behave in that way.¹⁵

 15 Norm Iteration has some other powerful consequences. Say that a behavior A is 'negative' just in case for some p, you are permitted to A iff you fail to omega know p. (Thanks to Kyle Blumberg for discussion.) For example, consider the activities of inquiring about p, and of failing to assert p. You might think that these are negative activities: you are permitted to inquire about p iff you don't omega know p, and you are permitted to refrain from asserting p iff you don't omega know p.

Negative behaviors would lead to a strange result when combined with Norm Iteration. Suppose that you are permitted to inquire about p iff you don't omega know p. Norm Iteration then implies that you are permitted to inquiry about p only if you omega know that you don't omega know p. But this implies that you are never permitted to fail to omega know p without omega knowing you fail to omega know p. Any failure of negative introspection would violate a norm

I deny that these behaviors are negative. In particular, I think that you can omega know p and still be permitted to inquire about p; and I think that you can omega know p and still be permitted to fail to assert p. In such cases, you are permitted to inquiry about p, and are also permitted to stop inquiring about p. You are permitted to assert p, and you are also permitted to not assert p.

Say a behavior is 'Moorean' when its permissibility depends on a Moorean truth obtaining. (Thanks to John Hawthorne and Cameron Domenico Kirk-Giannini for discussion.) As a toy example, suppose that a research proposal about p should be funded only if p is true and you don't know whether p. Norm Iteration would then imply that the research proposal should be funded only if you know that: p is true and you don't know whether p. But this condition can't obtain, and so Norm Iteration would predict that it was impermissible to fund the research proposal. In this way, Norm Iteration is inconsistent with permissible Moorean behaviors.

In response, the defender of Norm Iteration might weaken the principle, to say that if you are permitted to A only if p and it is possible to know that p, then you are permitted to A only if you know that p. But this weakening of Norm Iteration is too weak. Another example of a Moorean behavior is asserting p and I don't know p. Principles like Norm Iteration explain why Moorean and dubious assertions are infelicitous. But the weakening of Norm Iteration risks permitting these behaviors.

Earlier, I suggested a constructive method for generating behaviors that are permitted only if p, for many choices of p. I imagined that you were correctly informed by a reliable informant that an innocent person would be killed if p is false and you press a certain button. I suggested that in that case, you are only permitted to press the button if p. But what if you are looking

¹⁴Here, I assume that knowledge is closed under infinite agglomeration. That is, if you know every member of a set of premises, and that set of premises implies a conclusion, then you know the conclusion.

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. This book is the first extended discussion of this dilemma. To solve it, I'll develop new theories of omega knowledge that deny KK, and also avoid 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.

Before continuing, I'll briefly note a few more arguments in the literature in favor of KK. In both cases, I think that the argument does not require KK, but instead requires something weaker.

at a tree, and a reliable informant correct informs you that an innocent person will be killed if you press a button and the tree is less than 80 feet tall and you don't know it? The defender of Norm Iteration is pressured to either deny that Moorean behaviors can be generated in this way, or must say that pressing the button is impermissible.

Finally, say that behavior A is exclusive about p iff you are permitted to A only if p, and you are permitted to not A only if not p. For example, imagine that you see someone acting suspiciously on the train. Perhaps you are permitted to say something to the conductor only if it is 10% likely that they are committing a crime, and you are permitted to say nothing to the conductor only if it is less than 10% likely that they are committing a crime. In that case, saying something would be an exclusive behavior about whether it is 10% likely that the suspect is committing a crime. Norm Iteration implies that if A is exclusive about p, then either you omega know whether p, or you are in a dilemma. For suppose you don't omega know whether p. Since you don't omega know that p, you aren't permitted to A; but since you also don't omega know not p, you aren't permitted not to A either. In response to this problem, I suggest that there are no exclusive behaviors. In the above example, I might grant that you are permitted to say something iff you omega know that it is 10% likely that the suspect is committing a crime. But I deny that you are permitted to say nothing iff you omega know that it is less than 10% likely the suspect is committing a crime. Rather, if you fail to omega know that it is 10% likely that the suspect is committing a crime, then you are permitted to say nothing.

The first argument is from Bonnay and Égré 2009 and Greco 2014a. They claim that KK explains why it is difficult to make sense of somebody having three iterations of knowledge without the fourth:

Without bringing in heavy-duty philosophical theory, there is no natural way to interpret ... "I grant that Jane knows that she knows that she knows what time the movie starts, but does she know that she knows that she knows what time the movie starts?" (Greco 2014a p. 196; see also Bonnay and Égré 2009 p. 200).

If KK is valid, then the question is incoherent. If KK is invalid, by contrast, then the question seems open. In particular, many who deny KK also embrace Omega Skepticism. According to such theories (presented in detail in Chapter 5), you could possess any number of iterations of knowledge without knowing you are in this position. Every time you gain an extra layer of knowledge, you have performed an even more difficult feat.

The opponent of KK may respond that the question about Jane is difficult to interpret because it is complicated. Here, it is worth contrasting intrapersonal and interpersonal iteration. It is easy to interpret analogous questions involving multiple knowers: "I grant that Billie knows that Carrie knows that Danny knows what time the movie starts, but does Alex know that Billie knows that Carrie knows that Danny knows what time the movie starts?" Yet this claim is at least as complicated as the analogous question about only Jane. The opponent of KK should explain why the two questions differ regarding ease of interpretation. ¹⁶ I think that the full force of KK is not required to explain this data point. In Chapter 2, I will introduce the principle of Reflective Luminosity, which says that whenever you know you know p, it follows that you omega know p. I'll show that Reflective Luminosity can explain this data without the full force of KK.

A second argument for KK, from Greco 2015, concerns rational coordination between multiple agents. The argument is that some cases of rational coordination require that a group of agents has common knowledge. But common knowledge requires omega knowledge; and so omega knowledge cannot be as scarce as omega skeptics say.

A group has common knowledge that p iff everyone in the group knows that p, everyone knows that everyone knows that p, everyone know that everyone knows that everyone knows that p, and so on. Common knowledge is used throughout the social sciences. ¹⁷

¹⁶One response may appeal to lack of belief. It is easy to imagine failures of interpersonal iteration because it is easy to imagine that Billie doesn't have a belief about whether Carrie knows that Danny knows. In fact, however, I think that questions of interpersonal iteration are easy to understand even in settings where it is clear that every agent has the relevant belief. Imagine a version of the case above where each agent has received evidence in support of their belief through a letter from a friend. So Carrie gets a letter saying that Danny knows what time the movie starts; and Billie gets a letter saying that Carrie knows that Danny knows what time the movie starts. It is easy to make sense of interpersonal knowledge failures, by imagining cases in which various letters are sent from unreliable sources. Yet in the intrapersonal case, it is harder to make sense of analogous questions.

 $^{^{17}}$ Lewis 1969 uses common knowledge to understand conventions. Linguists use common

One reason to believe in common knowledge concerns coordination.¹⁸ Some have argued that without common knowledge, groups could not coordinate in certain ways. For example, consider:

(25) Coordinated Attack. "Two divisions of an army are camped on separate hilltops overlooking a valley. In the valley awaits the enemy. If both divisions attack the enemy simultaneously they will win the battle, while if only one division attacks it will suffer a catastrophic defeat. Each of the generals [North and South] commanding these hilltop divisions wants to avoid a catastrophic defeat: neither of them will attack unless he believes that the general commanding the other division will attack with him. During the night a thick fog descends over the hilltops; the only way the generals can communicate is by sending a messenger through the enemy camp." (Lederman 2018b)

In this case, the generals do not have common knowledge that they will attack. The problem is that each can send a messenger to the other, but there is a chance that the messenger will be lost in the fog. At most, the generals can send a finite number of messages, which will produce at most a finite number of iterations of mutual knowledge.

Many have argued that since the generals lack common knowledge, they are not rationally permitted to attack.¹⁹ In particular, North can rationally attack only if he knows South will attack; but North also knows that South will rationally attack only if he knows North will attack. Iterating this reasoning, North and South can rationally attack only if they commonly know that they will both attack.²⁰ Summarizing, the argument is that some kinds of coordination require common knowledge, and common knowledge requires omega knowledge, and so omega knowledge is required for some kinds of rational coordination.

Omega Skepticism implies that common knowledge is scarce. After all, common knowledge implies omega knowledge: if a group has common knowledge that p, then each member has omega knowledge that p. So skepticism about omega knowledge implies skepticism about common knowledge.

The omega skeptic must deny that any important kind of rational coordination requires common knowledge. One strategy would explain coordination using a different common attitude, such as common belief or common certainty. Another strategy denies that coordination requires any kind of common attitude. For example, Lederman 2018b argues that in these cases, agents can coordinate because they do not have common knowledge that they are rational.²¹

knowledge to understand speech acts (Stalnaker 1973). Game theorists use common knowledge to understand decision making (Aumann 1976). Computer scientists use common knowledge to understand distributed systems (Fagin et al. 2003).

¹⁸See for example Heal 1978 and Clark and Marshall 1981.

¹⁹See for example Fagin et al. 2003.

 $^{^{20}\}mathrm{See}$ also Rubinstein 1989's example of the electronic mail game, introduces decision theoretic structure.

 $^{^{21}\}mathrm{See}$ Lederman 2018a for an argument that common knowledge is impossible, regardless of whether Omega Skepticism is true.

1.3 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 book, I take a different route. I develop new theories of knowledge that make omega knowledge abundant without identifying omega knowledge and knowledge. One reason that I take this approach is that there are important arguments against KK. 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.

Greco 2014b defends KK from this type of counterexample. He argues that the unconfident examinee suffers from fragmentation rather than higher order ignorance. On this proposal, the unconfident examinee has two different beliefs: relative to one fragment, she knows and omega knows when Queen Elizabeth died; relative to the other fragment, she does not know when Queen Elizabeth died. On this proposal, questions that are explicitly about higher order knowledge tend to activate the ignorant fragment. When the unconfident examinee is attending to the question of whether she remembers studying English history, the ignorant fragment may be active. But when the examinee takes her guess, the knowing fragment can activate instead.

One challenge for this account is to explain why asking about the unconfident examinee's higher order knowledge always makes salient the fragment relative to which she is ignorant. According to KK, higher order knowledge is the same state as first order knowledge. On this proposal, the question of whether the examinee knows is identical to the question of whether the examinee knows that she knows. In the context of the story, and holding fixed KK, it is clear that she would know if and only if she knows that she remembers. So asking whether she knows she remembers would be contextually equivalent to asking whether she knows. For these reasons, it is unclear why discussing higher order knowledge would make salient a different fragment than asking about first order knowledge.²²

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.²³

²²Thanks to Kyle Blumberg for help here.

 $^{^{23}}$ Norm Iteration implies that omega knowledge is required for permissible behavior. In this way, Norm Iteration implies that animals that lack the concept of knowledge are not permitted

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:

(26) **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. 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,

to engage in any behavior. One way to avoid this consequence is to replace 'knowledge' in Norm Iteration with 'being in a position to know' (a distinction I'll return to soon), and then allow that conceptually impoverished animals can be in a position to know facts about what they are in a position to know. Another option is to restrict Norm Iteration to agents that possess the concept of knowledge, and claim that agents that do not possess the concept of knowledge are not responsible for their actions in the sense relevant to wondering whether they are permissible.

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 n that the tree is not n feet tall. But then something you omega know would be false

In Chapter 4, I explore a weakening of Margin for Error, which I call Variable Margins. According to this weaker principle, the margin for error required for knowing varies based on the height of the tree. This variance produces a weaker principle than Margin for Error. According to Variable Margins, 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'll show in Chapter 4 that this weakening avoids the argument above, making room for a theory that respects inexact knowledge while rejecting Omega Skepticism.

One reason to accept Margin for Error is that it follows quickly from the principle that knowledge requires safety from error:

- (27) Safety. If you know p, then you could not easily have believed p falsely.²⁴ Margin for Error follows from Safety given a further assumption:
- (28) **Possibility.** For any height x, if the tree is x feet tall and you believe it is not x 1 feet tall, then it could easily have been x 1 feet tall while you believed it was not x 1 feet tall.

This premise itself follows from two theses. First, whenever the tree is a certain height, it could easily have been a slightly different height. Second, what you believe about the tree's height is independent of what height the tree is.

 $^{^{24}}$ For defenses of Safety, see among others Sosa 1999, Williamson 2000, Pritchard 2005, Manley 2007, Williamson 2009 and Goldstein and Waxman 2020. Safety principles are standardly relativized to methods: you know p using a method only if you couldn't easily have falsely believed p using that method. Throughout this book, I mostly suppress method-relativity. In addition, this statement of Safety is proposition relative. Whether your belief that p varies across nearby possibilities could be independent of whether your belief that q varies across nearby possibilities. So knowledge of p would require considering a different set of worlds than knowledge of q. I suppress this complexity throughout. To do so, I assume that what you believe is luminous to you, so that if you believe p, then you know you believe p. In addition, I assume that belief is closed. Given these assumptions, Safety is equivalent to the condition that you know p only if you couldn't easily have had the same overall set of beliefs while p were false. Then I can consider which set of possible worlds are those where you could easily have had the same set of beliefs.

Safety and Possibility imply Margin for Error. Take an arbitrary height, say 100 feet. By possibility, if the tree is 100 feet tall and you believe it is not 99 feet tall, then it easily could have been 99 feet while you believed it was not 99 feet tall. By Safety, it follows that you don't know that it is not 99 feet tall.

Throughout this book, I accept Safety. But I reject the validity of Margin for Error. In order to do this, I distinguish two interpretations of Safety and Possibility: similarity, and normality.²⁵ I'll show that the normality interpretation can embrace Safety while rejecting Possibility and Margin for Error. (In accepting Safety without Margin for Error, I follow Greco 2014a among others.)

The similarity interpretation says that in order to know p, you cannot falsely believe p in any possibility that is similar to the actual world. The similarity of two worlds is connected to counterfactuals: to know p, you could not have falsely believed p if the actual world had been slightly different. If knowledge involves similarity, KK fails. Similarity is intransitive. v may be similar to w and u may be similar to v, even though u is not similar to v. For example, when the tree is 100 feet tall, you know that it is not less than 90 feet tall, because 90 feet is a similar height to 100 feet while 89 is not. When the tree is 90 feet tall, you know that it is not less than 80 feet tall, because 80 feet is a similar height to 90 feet while 79 is not. But when the tree is 100 feet tall you do not know that you know that the tree is not less 90 feet tall, because this claim is not true in all cases that are similar to all cases that is similar to the actual case.

The similarity interpretation accepts Omega Skepticism. Omega knowledge is scarce, because it is disrupted by intransitivity in the chain of epistemic accessibility. Almost any claim you know is false somewhere down the chain of epistemic accessibility, because this chain can extend indefinitely far across the space of possible worlds.

Later in the book, I'll develop theories of omega knowledge that reject Omega Skepticism and KK. These theories do not sit well with the similarity interpretation of Safety.

For example, the first theory I'll consider is the principle of Reflective Luminosity: if you know that you know something, then you omega know it. If knowledge requires that your belief is true across counterfactually similar possibilities, this principle fails in the same way as KK. Again, similarity produces indefinitely long chains of intransitivity. A 100 foot tree is similar to a 90 foot tree but dissimilar to an 80 foot tree. The 90 foot tree is similar to the 80 foot tree and dissimilar to a 70 foot tree, and so on. So when the tree is 100 feet tall you know that you know it at least 80 feet tall, but you do not omega know that it is at least 80 feet tall.

Later, I'll explore the principle of Fragility: if you know p, then it is consistent with your knowledge that you omega know p. This too sits poorly with the similarity interpretation of Safety. When the tree is 100 feet tall, you know it is at least 90 feet tall. But you don't know that you know this. And you can know that you don't know that you know this, because in every counterfactually

²⁵For defense of the similarity interpretation, see Bacon 2014. For defense of the normality interpretation, see Greco 2016, Goodman and Salow 2018, Beddor and Pavese 2019, Goldstein and Hawthorne forthcominga.

similar possibility, you fail to know that you know that the tree is at least 90 feet tall.

I am not an omega skeptic. So I reject the connection between knowledge and similarity. I have independent reason to reject this link. Those who understand Safety in terms of counterfactual similarity must accept the following principle:

(29) Counterfactual Closure. If it could easily have been that p, and if p had been the case, q would have been the case, then it could easily have been that q.

But Counterfactual Closure is false. At noon, a fair coin is flipped and Betty bets a hundred dollars that it will land heads. Betty won't learn the result until 1 o'clock. So at noon, Betty doesn't know whether she will be sad when she sees the result. This means that Betty could easily have been sad when she saw the result (p). As a matter of fact, the coin landed heads. So if Betty were sad when she saw the result, she'd be sad about winning the bet. But at noon, Betty knows that she won't be sad about winning the bet. So Betty could not easily have been sad about winning the bet (q).

Here is another example. There is a room connected to a cage by a passage. A tiger prowls back and forth between the room and the cage. There is a trap door above the room. If Alex opens the trap door, the passage seals, locking the tiger in its current location. Alex plans to peek into the room and then spend the night just in case the tiger isn't there. In fact, the trap door is slightly ajar and so the passage is closed. The tiger is trapped in the room. Right before Alex peeks, there is a paradigmatically easy possibility that the tiger is in its cage. Alex knows he'll spend the night just in case the tiger is in its cage. So there's an easy possibility that he'll spend the night in the room. But since there is actually a tiger in the room, Alex would have been eaten by the tiger if he spent the night. But Alex knows he won't be eaten by a tiger. So there is no easy possibility where he's eaten by a tiger.

If knowledge does not involve similarity, what is it? An alternative interpretation connects knowledge to the normality of an agent's belief forming process. 27 To know p, your belief forming process could not have produced a false belief in any situation that is as normal as your actual predicament.

Normality explains how knowledge is and is not related to probability. There's a trillion-sided die. You can't know whether it will land 1017 before looking. But if it comes up 1017, you can know it did by looking. Yet the risk of hallucinating is greater than one in a trillion. Before looking, any roll of the die is equally

²⁶Some may resist the relevant counterfactual judgments. One way to access the relevant reading focuses on causal independence. The coin landing heads is causally independent of Betty being sad when she saw the result. But it is natural to hold fixed causally independent events when counterfactually imagining that Betty had been sad when she saw the result. On the other hand, it is possible to access 'backtracking' readings of the counterfactual, where you counterfactually modify the outcome of the coin toss to accommodate Betty's reaction. For further discussion of these matters, see Goldstein and Hawthorne forthcominga.

²⁷See for example Dretske 1981, Greco 2014a, Stalnaker 2015, Goodman and Salow 2018, Carter 2018, Beddor and Pavese 2019, and Carter and Goldstein 2021.

normal; but after looking, 1017 is more normal than other outcomes, because hallucination is abnormal.²⁸ There are various conceptions of normality in the literature. On one conception, you know p only if your state carries the information that p. A state carries the information that p only if normally p is true when you are in that state.²⁹ Here, the normality of a state has something to do with whether it is functioning optimally. Other conceptions of normality, such as that in Carter 2018, appeal directly to ordinary judgments about the normality of various situations. This book focuses on the structure of normality, staying neutral on what exactly normality is.

The main claims in this book are neutral on the exact nature of normality. But for the sake of concreteness, it's worth saying a bit more. My preferred interpretation of normality involves evidence and accuracy (see Carter and Goldstein 2022). You form beliefs about the world on the basis of evidence. Sometimes your evidence misleads you about the world; in other cases, your evidence gives you a very good grip on what world the world is like. I think how much you know depends on how accurate your evidence is. When conditions are most normal, your evidence is maximally accurate. In some abnormal situations, your evidence is quite inaccurate. For example, in skeptical hypotheses, your evidence is radically misleading about the world. Imagine that you are a brain in a vat, receiving electrical signals from a team of scientists, which create sensory experiences suggesting that you are living an ordinary life. In this case, your evidence is misleading. Your evidence suggests that you have hands, and that you're walking around in the real world. But this isn't true.

How do we measure normality, conceived in terms of the accuracy of evidence? My preferred take is to rely on a measure of 'evidential probability', which says how likely the evidence makes various claims. To find a measure of evidential probability, you need two things. First, you start with a prior probability distribution, which measures the 'intrinsic plausibility' of various hypotheses (Williamson 2000). Then, you conditionalize this prior on a body of evidence. The resulting measure says how strongly this body of evidence supports various hypotheses. The accuracy of your evidence at a possible world is then proportionate to the probability assigned to the possible world by the evidential probabilities (see Goldstein and Hawthorne forthcomingb, Goodman and Salow forthcoming for further discussion). None of the main theses in the book depend on this particular conception of normality, in terms of evidential accuracy, which is itself understood in terms of evidential probability. But this is one way of interpreting the theses about the structure of normality and knowledge which follow.

Ultimately, though, I don't have a strong opinion about exactly what normality is. Instead, I prefer to develop a theory of normality by considering which principles about knowledge are normally true. In Chapters 2 and 3, I'll explore two key theses about normality. The first thesis is that Margin for Error is normally true; the second thesis is that KK is normally true. I'll show how

²⁸See Nelkin 2000, Smith 2010, and Carter 2018 for similar examples.

 $^{^{29}\}mathrm{See}$ for example Stampe 1977; Dretske 1981; Millikan 1984; Stalnaker 1999; and Greco 2014a.

varying theories of omega knowledge can be supplemented with these principles about normality, in ways that have consequences for the theory of knowledge and justified belief.

In Chapter 5, I develop several theories of knowledge in detail, which make specific claims about the structure of normality. Until then, the main aspect of normality I'll rely on is that you can compare different situations for how normal they are. In the good case, your belief forming processes are functioning most normally, and you know the most. In the bad case, your belief forming processes are functioning least normally, and you know the least.

Beddor and Pavese 2019 offer a simple approach that validates KK. The strongest thing you know at any world w is that you inhabit a world at least as normal as w. On this proposal, v could easily have been the case iff v is at least as normal as w. This validates KK. At any world v compatible with what you know at w, conditions are at least as normal as at w. So you know at least as much at v as you do at w.

Other normality models reject KK, and embrace Omega Skepticism. For example, Goodman and Salow 2018 describe the conditions under which w is significantly less normal than v. One of their theories claims roughly that at any world w, the strongest thing you know is that conditions are not significantly less normal than v. On this proposal, v could easily have been the case iff v is not significantly less normal than w.

This leads to failures of KK. There are chains of worlds w, v, and u, where v is less but not significantly less normal than w, and u is less but not significantly less normal than u, and yet u is significantly less normal than w. At w, you know but do not know that you know that you are not in u.

In chapter 5, I develop models knowledge in terms of normality that validate Reflective Luminosity and Fragility. I illustrate the key concepts with the case of perceptual knowledge, where someone gains information about a quantity like tree height. The tree appears to be some height, and the tree is some height. What the agent knows depends on what is normal. What is normal depends on the distance between the real and apparent height of the tree.

1.4 Justification

Throughout the book, I consider not only knowledge but also justification. I am especially interested in a general notion of justification for any behavior A, defined in terms of knowledge:

(30) **Possible Permission.** You are justified in A iff for all you know, you are permitted to A.

The idea of Possible Permission is that you are justified in engaging in a behavior iff for all you know, you satisfy all of the requirements for that behavior. One way to motivate Possible Permission is to think of justification as a kind of excuse. There is a special kind of excuse you can have for your behavior just in case you are in a state that is epistemically indiscernible from doing what you

are actually allowed to do.³⁰

Variations on Possible Permission replace epistemic possibility with stronger epistemic positions. For example, one alternative view says that you are justified in A iff it is sufficiently likely on the evidence that you are permitted to A. For example, Littlejohn and Dutant 2020 propose that you are reasonable in believing p iff there is a high probability that you know p, conditional on what you know. I focus on Possible Permission instead, in part because this makes it easier to explain why justification is entailed by knowing.

Now consider how Possible Permission applies to various behaviors. First, consider assertion. I've argued that you are permitted to assert p iff you omega know p. It follows that you are justified in asserting p iff for all you know, you omega know p.³¹ For example, this theory says that you can be justified in asserting that you have hands even when you are a brain in a vat. In that case, you are not actually permitted to assert that you have hands, since you don't have hands, and you don't omega know you have hands. But in this scenario, it is epistemically possible that you omega know you have hands, and so you are justified in asserting it. In this way, Possible Permission systematizes the sense in which victims of skeptical scenarios have a good excuse for behaving in the same way as ordinary people.

Now consider action and inquiry. You are justified in acting as if p iff for all you know, you omega know p; and you are justified in stopping inquiry about p when for all you know, you omega know p. On this picture, all of these activities are permitted iff you omega know, and are justified iff for all you know they are permitted.

Imagine that you know that one of your 30 students cheated, but you don't know which. You are permitted to fail a student only if they cheated. Does Possible Permission imply that you are justified in failing every student, since for each student, it is epistemically possible that they cheated? Not if Norm Iteration is true. Norm Iteration implies that you are permitted to fail a student only if you omega know they cheated. If you know you don't omega know that a student cheated, then you are not justified in failing them. In this way, Norm Iteration combines with Possible Permission to raise the standards required for justification. ³²

Now consider justified belief. There are many interesting concepts of belief; in this book, I focus on 'strong' belief, which I understand as subjective certainty. I tentatively take on board Omega Infallibilism: you are permitted to be subjectively certain of p iff you omega know p. By Possible Permission, it follows that you are justified in being subjectively certain of p iff for all you know, you

³⁰Others, like Williamson forthcoming and Greco forthcoming distinguish excuses from justifications. But Williamson forthcoming understands excuses in a similar way: you are excused from the requirements when you behave in a way that would satisfy the requirement if conditions were normal. In Chapter 5, I develop models where this condition is equivalent to Possible Permission. See also Kelp and Simion 2017, which argues that if you are blamelessly ignorant of violating the knowledge norm, then you are blameless in acting.

³¹Throughout the book, in these kinds of arguments I assume that you omega know all of the relevant facts about knowledge.

³²Thanks to Bob Beddor for discussion.

omega know p. If belief is subjective certainty, this implies:

(31) **Possible Omega Knowledge.** You are justified in believing p iff for all you know, you omega know p.

Possible Omega Knowledge reduces facts about justification to facts about knowledge.

Possible Omega Knowledge implies that dubious assertions are never justifiably believed. You cannot justifiedly believe p and I don't know that I know p, because you know that you don't omega know this claim. Omega knowing this claim would imply omega knowing p, and knowing that you don't know that you know p. But this is absurd.

Some have argued that belief is not subjective certainty, because subjective certainty is scarce. But I deny Omega Skepticism, and think that omega knowledge is abundant. If you are permitted to be subjectively certain of anything you omega know, then subjective certainty is scarce.

Chapter 3 explores Fragility, the principle that if you know p, then for all you know, you omega know p. In the presence of Omega Infallibilism and Possible Permission, this implies that when you know p, you are justified in being subjectively certain of p. In this way, Fragility can make do with justified subjective certainty as its primary notion of justified belief.

On the other hand, Chapter 2 explores Reflective Luminosity, according to which whenever you know that you know p, you omega know p. Reflective Luminosity allows that there are situations where you know p while knowing you don't omega know p. In this way, it allows you to know p when you are not justified in being subjectively certain of p. This view requires another notion of justified belief, if justified belief is required for knowing.

Here, one option is to analyze justification in terms of possible knowledge instead of possible omega knowledge:

(32) **Possible Knowledge.** You are justified in believing p iff for all you know, you know that p.³³

According to this principle, you are justified when you are in a state that is epistemically indistinguishable from knowing. In Chapter 5, I develop models of knowledge in terms of normality. In that setting, Possible Knowledge corresponds to the principle that you are justified in believing whatever you would know when conditions are most normal. 34

In Chapter 3, I show that Fragility implies that possible knowledge is identical with possible omega knowledge. On this view, Possible Knowledge is equivalent to Possible Omega Knowledge. By contrast, in Chapter 2 I show that Reflective Luminosity is compatible with possible knowledge being different than possible omega knowledge.

³³See Lenzen 1978, Williamson 2000 (p. 46), Stalnaker 2006, Williamson 2013a, Rosenkranz 2018, and Carter and Goldstein 2021. For similar views, see Bird 2007 and Ichikawa 2014.

³⁴See Williamson forthcoming for further discussion of the connection between normality, justification, and excuses.

If possible knowledge is distinct from possible omega knowledge, then Possible Knowledge may be difficult to reconcile with Norm Iteration and Possible Permission. Norm Iteration implies that some kind of omega knowledge is necessary for permissible belief. But then Possible Permission will imply that you are only justified in believing p if for all you know, you have some kind of omega knowledge.

Some epistemologists think there is an important difference between 'primary' and 'secondary' norms. One hypotheses is that the primary norm on assertion, action, inquiry, and subjective certainty is to know, while omega knowing is only a secondary norm. (This doesn't affect my arguments against Omega Skepticism.) Then one might say that you are 'primarily' permitted to A iff you satisfy the primary norm on A. In this case, one might modify Possible Permission, to say that you are justified in A iff for all you know, you are primarily permitted to A. In this case, one could accept Possible Knowledge instead of Possible Omega Knowledge. I am somewhat skeptical of this strategy, because I don't see what other theoretical role the primary / secondary distinction would play other than to define this notion of justification. But ultimately I don't have a strong opinion about which of Possible Omega Knowledge or Possible Knowledge is a better theory of justified belief: I will explore both principles in detail throughout the book.

Both Possible Omega Knowledge and Possible Knowledge predict that belief is 'strong' in a few senses. First, both principles predict that you are not justified in believing Moorean claims of the form p and I don't know p. They also potentially predict that you are not justified in believing that a fair lottery ticket will lose, if you know that you don't know (or omega know) that a fair lotterv ticket will lose.³⁵ In this way, such principles distinguish justified belief from rational high credence.³⁶ Possible Omega Knowledge predicts that belief is strong in another sense: if you are justified in believing p, then you are justified in believing you know p. After all, omega knowing p implies omega knowing that you know p. So if it is possible that you omega know p, then it is also possible that you omega know that you know p.

Both principles also offer a solution to the new evil demon problem (Lehrer and Cohen 1983). The new evil demon problem is that you are justified in believing the same things in the good case, where the world is as it appears, as you are in the bad case where the world departs radically from appearance, for example when you are deceived by an evil demon. But externalist theories of justification like Goldman 1979 and others do not make this prediction.

Possible Omega Knowledge and Possible Knowledge can avoid the new evil demon problem (Williamson 2013a). What you know depends on whether you are deceived by a demon. But what you possibly know may not depend on your environment in the same way. In Chapter 5, I'll consider existing work that vindicates this idea in the case of perceptual knowledge, by interpreting normality in terms of the difference between the real and apparently perceived

 $^{^{35}\}mathrm{See}$ Williamson 2000 p. 255 and Marušic 2013 among others for defense of the claim that you are not justified in believing that a lottery ticket will lose.

36 See for example Foley 1993 and Christensen 2005.

values of a quantity.

Both Possible Omega Knowledge and Possible Knowledge are theories of propositional, not doxastic justification. This means that they don't describe the conditions under which you are justified in believing what you actually believe, on the basis you actually have. Instead, they describe the conditions under which you would be justified to believe p if you did believe p. Suppose you almost believe p, but just barely lack the requisite confidence. Still, you are unsure of what it takes to believe. For all you know, you believe and know p. But as a matter of fact, you don't. Since you don't believe p, you are not doxastically justified in believing p. This case is an immediate counterexample to Possible Omega Knowledge or Possible Knowledge when interpreted in terms of doxastic justification; but not in terms of propositional justification.

Many epistemologists distinguish knowledge from being in a position to know. You are in a position to know p when you satisfy the evidential, modal, and reliabilist conditions of knowing, but have not yet formed a belief. In my models of knowledge, following the literature cited above, I focus on people who know exactly what they are in a position to know. On the other hand, in Chapters 2 and 3 I will consider whether agents who violate this assumption make trouble for Reflective Luminosity or Fragility.

Possible Knowledge and Possible Omega Knowledge are best interpreted as governing being in a position to know, rather than knowledge itself. It is easy to know that you don't believe p, when you don't. In such a case, you know that you don't know that p. But for all that, you may be propositionally justified in believing p, since the only barrier to your knowing may be that you have not yet formed a settled opinion. Similarly, imagine that you have hardly any evidence at all in favor of p; but what if you haven't considered the question of whether you know p? Then for all you know, you do know p; but you aren't justified in believing p. In this case, you are in a position to know that you aren't in a position to know p; the only problem is that you haven't carefully considered the question yet. 373839

³⁷A further challenge for these principles concern people who lack the concept of knowledge. Such people do not know that they do not know anything. But it doesn't follow that they are justified in believing anything at all. Again, the problem may be avoidable by distinguishing what they know from what they are in a position to know (imagining that they gained possession of the concept of knowledge.)

On the other hand, perhaps even without having the concept of knowledge, what you know could guarantee that you don't know p. If something is blue, this guarantees it is not green. So if you know something is blue, this guarantees that you don't know it is green. So if you know something is blue, you are not justified in believing it is green, since what you know guarantees that you don't know it is green. In this way, another version of these principles would define justification in terms of what your knowledge guarantees about what you know. Thanks to John Hawthorne for discussion.

 $^{^{38}\}mathrm{See}$ Rosenkranz 2018 for detailed defense of Possible Knowledge, and critical discussion of being in a position to know.

 $^{^{39}}$ Another question is whether to interpret Omega Assertion and related norms as governing being in a position to know, rather than knowledge itself. See Willard-Kyle 2020 for arguments that assertion involves what you are in a position to know, rather than what you know. In this case, one could accept a variant of Norm Iteration, which says that if you are permitted to A only if p, then you are permitted to A only if you are in a position to know p.

My main arguments in the book involve knowledge rather than justification. For this reason, I am not committed to either Possible Omega Knowledge or Possible Knowledge. But I explore these principles in the book because each of Reflective Luminosity and Fragility have powerful consequences for the theory of justification in the presence of either Possible Omega Knowledge or Possible Knowledge. I'll show that facts about justified belief may help to decide which of Reflective Luminosity, Fragility, and/or Variable Margins to accept.

I have written this book so that extensive use of mathematical symbols is confined to Chapter 5. My hope is that readers who are not interested in detailed questions about epistemic logic can read the other chapters, and understand the main claims and contributions of the book.

2 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 chapter explores the principle that reflective knowledge is luminous:

(33) **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 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. But Reflective Luminosity shares one structural property in common with KK. KK says that some inferences are knowledge preserving without being truth preserving. p does not imply that you know p. But KK say that knowing p implies knowing that you know p. Even though p is weaker than knowing p, KK says that knowing p is the same as knowing that you know p. I deny KK. So I say that knowing p does not imply knowing that you know p. But I agree that some inferences are knowledge preserving without being truth preserving. In particular, I say that knowing you know p implies knowing you know that you know p. Even though knowing p is weaker than reflectively knowing p, knowing you know p is the same as knowing you reflectively know p.

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.

Recent work in cognitive science suggests that human cognition involves two different systems. 40 The fast system makes a snap judgment, and the slow system follows up with reflection. One route to Reflective Luminosity says that the slow and fast system differ in the kinds of knowledge they provide. When your fast system reliably judges that p, you know p without knowing that you know. But when your slow system reliably judges that p, you know that your judgment is knowledge. Now suppose that judgments about what you know all issue from the slow system. The result is Reflective Luminosity. This provides one interpretation of the unconfident examinee. When asked the historical question, their fast system delivers a verdict; but their slow system does not. The result is a hesitant guess that is known, but not omega known.

Whether Reflective Luminosity is true depends on the nature of knowledge. In Chapter 5, I explore the prospects for Reflective Luminosity within a theory of knowledge as belief that is safe from error. I interpret the relevant kind of safety as belief that is normally true, and I show how a particular interpretation of normality leads to the validity of Reflective Luminosity.

2.1 Benefits of Reflective Luminosity

I'll start by considering several benefits of Reflective Luminosity.

First benefit: Reflective Luminosity explains how Omega Assertion and Norm Iteration can be true.

First, Reflective Luminosity vindicates the arguments from Chapter 1 suggesting that omega knowledge plays an important role in assertion and behavior more generally.

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.

⁴⁰See Kahneman 2011 for an overview.

In Chapter 1, 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 Chapter 1, I suggested that assertion is merely one of a wide class of behaviors governed by omega knowledge. In particular, 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 Chapter 1, 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. 41

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.

Finally, return to the discussion of common knowledge. Reflective Luminosity can explain how common knowledge is possible. KK fails. You can know without knowing you know. But reflective knowledge is omega knowledge. In this way, Reflective Luminosity can explain the possibility of common knowledge without resorting to the KK principle.

Reflective Luminosity is only compelling if knowledge of knowledge is abundant rather than scarce. Otherwise, Reflective Luminosity leads to Omega Skepticism. According to Reflective Luminosity, any known knowledge possesses every iteration of knowledge. One worry is that this kind of knowledge is so difficult to come by that it is scarce, applying only to our knowledge of basic

⁴¹Marušic 2013 argues that you should believe p only if you reflectively know p, but denies that you should believe p only if you omega know p.

mathematical facts and perceptual seemings. In Chapter 5, I assuage this worry by developing a model of knowledge where reflective knowledge is abundant.

Second benefit: Reflective Luminosity explains why three iterations of knowledge seems like the same thing as four.

Recall from Chapter 1 that one motivation for KK is that it is difficult to make sense of somebody having three iterations of knowledge without the fourth. This is exactly what Reflective Luminosity predicts. By contrast, omega skeptics allow indefinitely long chains of intransitivity between possibilities. So omega skeptics allow you to possess any number of iterations of knowledge without knowing that you are in this state.

Third benefit: Reflective Luminosity explains the unconfident examinee case.

In Chapter 1, I argued that KK is vulnerable to counterexamples. 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.

Fourth benefit: Reflective Luminosity explains why it is unusual to ask people whether they possess second-order knowledge.

I have argued that omega knowledge plays an important role in assertion, in action, and in several other domains. One challenge for this thesis is that when we criticize the assertions and actions of others, it is much more natural to discuss knowledge than it is to discuss iterations of knowledge. For example, while it is natural to respond to an assertion with the question 'how do you know?', it is much stranger to instead ask 'how do you know that you know?'.

KK offers one explanation of what is going on. The more complicated question is equivalent to the simpler one. So there is no reason to ask the more complicated question (see Haziza forthcoming p. 20 for discussion).

Reflective Luminosity offers an alternative interpretation of these facts. In Chapter 1, I considered Omega Infallibilism, the thesis that you are rationally permitted to be certain of exactly what you omega know. I suggested that you know something for sure just when you know it with rational certainty. If these theses are accepted, then Reflective Luminosity implies that you know something for sure if and only if you know that you know it. This can explain why it is strange to ask about iterations of knowing. It is not strange at all to ask whether someone know something for sure. Reflective Luminosity predicts that there are two natural questions you can ask someone when they rely on p: whether they know p, and whether they know p for sure. Asking about iterated knowledge is dispreferred because it is equivalent to asking about knowing for sure. By contrast, any defender of KK who accepts Omega Infallibilism will be forced to equate these two questions, which appear to be distinct. (In Chapters 3 and 4, I develop two alternative theories of omega knowledge. On these theories, each iteration of knowledge is distinct from omega knowledge. These theories have a harder time explaining why ordinary people so rarely ask whether someone 2-knows, or 3-knows, and so on.)

Fifth benefit: Reflective Luminosity predicts that Margin for Error is normally true.

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:

(34) **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.

Suppose for example 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 that 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 Lumi-

nosity 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.

In Chapter 1, I considered the idea that what you know depends on the normality of your situation. In the good case, your belief forming processes are operating normally; in the bad case, they are operating abnormally. In Chapter 5, I develop precise models of this idea. One of my goals there will be to show that Reflective Luminosity is consistent with the idea that you know Margin for Error is true in the good case:

(35) Known Margins. In the good case, you know Margin for Error.

The idea is that Margin for Error seems true because the principle is known in cases where conditions are maximally normal (in 'the good case'). Here is another way of thinking about it. 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 say that a principle is normally true when it is true in any situation that is almost as normal as the maximally normal one. In this setting, Reflective Luminosity then allows that knowledge is normally governed by a margin for error. 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. (In the next chapter, I'll consider an alternative thesis: that normally, KK is true.)

Known Margins and Reflective Luminosity are consistent. By contrast, Known Margins leads to absurdity when combined with KK. In the presence of KK, Known Margins implies that Margin for Error is omega known in the good case. As I discussed above, this would imply that in the good case you know nothing about the tree's height. 42

Some may balk at this explanation of Margin for Error. At first glance, Margin for Error seems to be a principle that is true iff it is a conceptual truth. But if Margin for Error is a conceptual truth, then plausibly it can be reflectively known. I am unsure of how to test whether Margin for Error is a conceptual truth. I believe that almost all instances of Margin for Error are true, and so

 $^{^{42}}$ One variant of Margin for Error replaces the material conditional with the subjunctive conditional. On this view, for any height x, if the tree had been x feet tall, then it would have been epistemically possible for you that the tree was x-1 feet tall. Reflective Luminosity is consistent with this subjunctive version of Margin for Error being known. The relevant question here is how tall the tree would have appeared if it had been x feet tall. Suppose that if the tree had been x feet tall, then it would have appeared x feet tall. Then if the tree had been x feet tall, it would have been epistemically possible that it was x-1 feet tall. In Chapter 5, I give a model of Reflective Luminosity that validates Known Margins. This model also allows knowledge of the subjunctive variant of Margin for Error, assuming that the tree would have appeared to be strictly within 1 of x feet tall, if it had been x feet tall. Thanks to Sam Carter for discussion.

Margin for Error appears at first glance to be a conceptual truth. But in the final analysis I think it is not a conceptual truth, since it implies Omega Skepticism. Here, perhaps it depends on why you find Margin for Error compelling in the first place. One reason could be that you reflect on the concept of knowing, and intuit that Margin for Error holds. Another reason could be that you reflect on various examples of knowing, and infer that Margin for Error holds as the best explanation of your judgments about cases. My explanation of Margin for Error's appeal makes better sense if your basis for Margin for Error is this kind of case based reasoning. In Chapter 5, I pursue this strategy with greater precision, delineating the conditions under which Margin for Error can be known along with Reflective Luminosity. Here's another reason to doubt that Margin for Error is a conceptual truth. The validity of Margin for Error should depend on contingent facts about our perceptual faculties. It is logically possible that our perceptual faculties could be infinitely discriminating. For example, an omniscient God might have this power. (On the other hand, inspired by Williamson 1994, one might derive Margin for Error from facts about semantic plasticity. Perhaps it is a conceptual truth that you can't know that John is tall if the sentence 'John is tall' could easily have expressed a falsehood. And maybe it is a conceptual truth that some sentences in our language could easily have expressed falsehoods. At any rate, in this book I'll try to avoid taking a stand on these difficult questions.)

Since Reflective Luminosity requires that you don't know you know Margin for Error, Reflective Luminosity implies that Margin for Error can possibly fail. But I showed in Chapter 1 that Margin for Error follows from Safety, combined with natural assumptions about what could easily have been the case.

- (36) Safety. If you know p, then you could not easily have believed p falsely.
- (37) **Possibility.** For any height x, if the tree is x feet tall and you believe it is not x-1 feet tall, then it could easily have been x-1 feet tall while you believed it is not x-1 feet tall.

In Chapter 5, I develop models of knowledge that validate Reflective Luminosity and invalidate Margin for Error. I respond to the argument above by preserving Safety, and rejecting Possibility. I do so by interpreting Safety in terms of normality rather than counterfactual similarity.

Sixth benefit: Reflective Luminosity can predict that you are justified in believing Margin for Error.

I'll now argue that Reflective Luminosity can strengthen its account of inexact knowledge, through the thesis that you are always justified in believing Margin for Error.

In Chapter 1, I considered two theories of justified belief. Possible Knowledge said that you are justified in believing p iff for all you know, you know p. Possible Omega Knowledge said that you are justified in believing p iff for all you know, you omega know p.

In the presence of Reflective Luminosity, Possible Knowledge and Possible Omega Knowledge offer different theories of justification. When you are in the good case, there are claims that you possibly know, but which you know you don't omega know. What you are justified in believing will then differ based on whether Possible Knowledge or Possible Omega Knowledge is true. In this section, I'll focus on the benefits for defenders of Reflective Luminosity if they accept each of Possible Knowledge or Possible Omega Knowledge. Then, in the next section on costs, I'll show that Reflective Luminosity faces important costs if it accepts either principle.

First, consider Possible Omega Knowledge. Given Reflective Luminosity, Possible Omega Knowledge is equivalent to the principle that you are justified in believing p iff for all you know, you reflectively know p.

Alternatively, the defender of Reflective Luminosity could instead accept Possible Knowledge. Given Possible Knowledge, Reflective Luminosity is consistent with the principle that agents in any condition are justified in believing Margin for Error:

(38) **Justified Margins.** You are justified in believing Margin for Error.

In the models I develop in Chapter 5, Possible Knowledge implies that Known Margins and Justified Margins are equivalent. You know p in the good case iff in any case, it is epistemically possible for you that you know p. In Chapter 5, I develop a model of Reflective Luminosity that validates Known Margins, and would validate Justified Margins when combined with Possible Knowledge.

By contrast, KK and Possible Knowledge lead to absurdity when combined with Justified Margins. If KK is valid and Justified Margins and Possible Knowledge are true, then you are not justified in believing anything about the tree's height. Suppose there is a height n where you are justified in believing the tree is not n feet tall (say it is 49). It follow that for all you know, you know the tree is not 49 feet tall. By KK, it follows that for all you know, you know you know the tree is not 49 feet tall. Similarly, from Justified Margins it follows that for all you know, you know that Margin for Error is true. Now suppose that justified belief agglomerates, so that whenever you are justified in believing p and justified in believing q, you are justified in believing p and q. It follows that for all you know: you know that you know the tree is not 49 feet, and you know Margin for Error. It follows that for all you know, you know that the tree is not 50 feet tall. The result is that you are justified in believing that the tree is not x feet tall, for any height x. In this way, Reflective Luminosity but not KK is compatible with Justified Margins and Possible Knowledge.

Justified Margins strengthens the error theory for Margin for Error. No matter how abnormal your situation, you are justified in believing Margin for Error. Small wonder that the principle seems true.

Possible Knowledge allows Reflective Luminosity to capture Justified Margins.

⁴³Many models of justified belief that accept Possible Knowledge validate agglomeration, including the ones I discuss in Chapter 5. On the other hand, see Carter and Goldstein 2021 for models that reject agglomeration in this setting in order to account for the preface paradox.

By contrast, Reflective Luminosity cannot validate Justified Margins when combined with Possible Omega Knowledge. (I clarify this point further in Chapter 5. The problem is that when conditions are most normal, you know that you don't omega know Margin for Error.)

Reflective Luminosity has one more potential benefit related to justified belief. It predicts that whenever you are justified in believing you know something, you are also justified in believing you know that you know. This offers an error theory for KK, by explaining why the principle seems attractive.

Consider two precisifications of the thesis that you are justified in believing KK:

- (39) a. If you are justified in believing you know p, then you are justified in believing you know that you know p.
 - b. You are justified in believing that if you know p, then you know that you know p.

Given Possible Knowledge, Reflective Luminosity accepts the first thesis, and rejects the second. Possible Knowledge and Reflective Luminosity imply that if you are justified in believing you know p, then you are justified in believing you know that you know p. In this sense, failures of KK are elusive. Whenever you seem to know p, you seem to know that you know p. Small wonder that KK seems true. (Possible Omega Knowledge immediately implies this thesis, independently of Reflective Luminosity.)

Now consider the second thesis. KK implies that you are justified in believing that if you know p, then you know that you know p. Possible Knowledge and Reflective Luminosity do not imply this. (Nor does Possible Omega Knowledge). In Chapter 5, I illustrate this by developing a model of Reflective Luminosity. By contrast, in the next chapter I'll show that Fragility can deliver this very strong error theory, predicting that you are justified in believing that the KK thesis is true.

In this way, Reflective Luminosity suggests that some confusion surrounding KK involves a scope ambiguity. Whenever you seem to know, you seem to know that you know. But you don't always seem to: know that you know if you know.

2.2 Costs of Reflective Luminosity

Reflective Luminosity also faces several challenges. In this section, I consider each challenge in detail. In the next section, I'll lay out a few ways of weakening Reflective Luminosity, and I'll suggest that these weaker principles avoid some of the costs of Reflective Luminosity.

First cost: Reflective Luminosity does not allow testifiers with different levels of reliability to produce different amounts of knowledge.

One challenge for Reflective Luminosity involves testimony. One popular idea is that reliable testimony transmits knowledge: if a reliable speaker tells you that p, and the speaker knows p, and you trust the speaker, then you come to know that p (see Leonard 2021 for an overview). This basic transmission principle could be extended to theses about iterated knowledge. If a very reliable person tells you p, and they know that they know p, then you come to know that you know p. If an even more reliable person tells you p, and they know that they know that they know that you know p.

Now imagine you could speak with two different speakers who are difficult to distinguish: Very Reliable and Less Reliable. Very Reliable's testimony is very reliable, and so would give you reflective knowledge of p; Less Reliable is less reliable, and so would give you knowledge of p without reflective knowledge of p. Now imagine that Less Reliable tells you that you know that p. Since Less Reliable's testimony generates knowledge, this should let you know that you know that p. But plausibly you fail to know that you know that you know that p, since you don't know whether you spoke with Less Reliable or Very Reliable. (Similarly, suppose that you spoke with Very Reliable, but could have easily spoken to Extremely Reliable, whose reliability ensures that when he says p, you come to know that you know that you know that p. In that case, you could know that you know that you know that p without knowing that you know that you know that p.

If such chains of speakers were possible, each governed by a corresponding transmission principle, then Reflective Luminosity would fail. In response to such cases, defenders of Reflective Luminosity must deny that there are pairs of speakers like Less Reliable and Very Reliable. One response would distinguish subject matters. Perhaps there can be a Less Reliable and Very Reliable speaker regarding ordinary claims about the world; but there cannot be such pairs of speakers regarding claims about what you know. In addition, such a response must also deny that there are pairs of Very Reliable and Extremely Reliable speakers about any topic, since this would immediately violate Reflective Luminosity.⁴⁴

Second cost: Reflective Luminosity does not generalize to revenge versions of Murine Research involving second-order knowledge.

A second challenge for Reflective Luminosity concerns higher-order analogues of Murine Research. Recall that Murine Research threatened the sufficiency of knowledge for the rational cessation of inquiry. In Murine Research, we imagined a researcher who knew the answer to a question, but who was also rational in searching for more information about the question, because they weren't certain of the answer.

Now consider:

(40) **Revenge Murine Research.** Mia is again a scientist who forms the hypothesis m: Accuphine causes hyperactivity in mice. Mia conducts a number of experiments that support m. Eventually, she conducts a large

 $^{^{\}rm 44}{\rm Thanks}$ to John Hawthorne for discussion.

number of experiments, and this allows her to know that she knows 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 Revenge Murine Research, Mia's experiments don't only give her knowledge that Accuphine causes hyperactivity. They also let her know that she knows this fact. Some readers may judge that in this case, Mia may permissibly fail to be absolutely certain that Accuphine causes hyperactivity. In this case, she is plausibly rational to open emails that provide more information about whether Accuphine causes hyperactivity.

Reflective Luminosity and Omega Infallibilism deny this possibility. They imply that in Revenge Murine Research, Mia can be rationally certain of m, and so can rationally refrain from gathering more evidence about m. Those who deny this conclusion have reason to reject Reflective Luminosity or Omega Infallibilism.

Third cost: Reflective Luminosity does not generalize to revenge versions of the unconfident examinee involving second-order knowledge.

A third challenge for Reflective Luminosity concerns revenge versions of the unconfident examinee involving higher-order knowledge.

Imagine an examinee who is answering questions on a strange kind of exam. Instead of simply asking about ordinary facts, the exam asks whether the examinee knows various facts. The examinee reaches the question of whether they know when Queen Elizabeth died. The examinee can't remember whether they learned this in a class, but decides to guess that they did. They guess correctly, in a way that is reliably caused by taking the class.

In this case, symmetry with the original unconfident examinee suggests that the higher order unconfident examinee knows that they know when Queen Elizabeth died. Reflective Luminosity then implies that they omega know when Queen Elizabeth died. But further symmetry with the original case implies that they don't know that they know when Queen Elizabeth died. In this way, Reflective Luminosity requires radically distinguishing first and second-order knowledge, so that unconfident examinee cases can never arise at the second order.⁴⁵

 $^{^{45}{\}rm The}$ unconfident examinee also raises challenges for the combination of Reflective Luminosity, Omega Infallibilism, and Possible Knowledge.

In at least some versions of the unconfident examinee, you are plausibly rational in asserting that you believe you know p even though you aren't certain of p. This suggests that you can know that: you are justified in believing you know p even though you aren't rationally certain of p. But this is incompatible with Reflective Luminosity, Possible Knowledge, and Omega Infallibilism. If you know you aren't rationally certain of p, then you know you don't omega know p. So by Reflective Luminosity, you know you don't know that you know p. But suppose Possible Knowledge is true. Then if you know you don't know that you know p, then you

Fourth cost: Reflective Luminosity doesn't allow you to remember seeing something without being certain of it.

Another challenge for Reflective Luminosity concerns iterated factive mental states.⁴⁶ Following Williamson 2000, suppose that factive mental states like remembering and seeing imply knowing. Plausibly, you can remember seeing that p, without being rationally certain that p. Moreover, Reflective Luminosity allows you to see that p without being certain of p, and allows you to remember that p without being certain of p; it would be strange if iterating these two uncertain states suddenly produced certainty.

But this is ruled out by the combination of Reflective Luminosity and Omega Infallibilism. If remembering and seeing both imply knowing, then remembering that you saw that p implies knowing that you know that p. So Reflective Luminosity implies that remembering you saw that p implies that you omega knew p, and so by Omega Infallibilism implies that you were rationally certain of p. (One response to this argument would deny that remembering and seeing imply knowing.)

Fifth cost: if Possible Omega Knowledge is true, then knowledge does not imply justification and Margin for Error is not justifiably believed; and if Possible Knowledge is true, then dubious assertions are justifiably believed.

I now consider the costs for Reflective Luminosity of adopting either theory of justified belief considered in Chapter 1.

Consider Possible Omega Knowledge. First, Reflective Luminosity and Possible Omega Knowledge together lead to failures of the implication from knowledge to justification.

The problem is that Reflective Luminosity is compatible with cases where you know p and know you don't omega know p (in Chapter 5, I develop a model that features these cases). But then Possible Omega Knowledge would imply that knowledge does not entail justification. (In Chapter 3, I introduce the Fragility principle, which says that if you know p, then for all you know, you omega know p. This principle is compatible with Reflective Luminosity, and would avoid this problem. Nonetheless, I ultimately argue that there are few reasons to accept both Reflective Luminosity and Fragility together; instead, the best theory should accept at most one of the principles.)

The second problem for combining Reflective Luminosity and Possible Omega Knowledge concerns Margin for Error. Above, I showed that Reflective Lumi-

aren't justified in believing you know p. Summarizing: if Reflective Luminosity is true, if omega knowledge is rational certainty, and if justified belief is possible knowledge, then you can't know that: you are justified in believing you know p without being rationally sure of p. (On the other hand, Possible Omega Knowledge and Omega Infallibilism are incompatible with this data point, regardless of Reflective Luminosity.)

⁴⁶Thanks to Ben Holguin.

nosity and Possible Knowledge together imply that you are justified in believing Margin for Error. But Possible Omega Knowledge loses this advantage. The problem is that in the good case, you know that you don't omega know that Margin for Error is true, and so you aren't justified in believing Margin for Error.

In the model I develop in Chapter 5, you would also omega know that you know you don't omega know Margin for Error, and so you would omega know that you aren't justified in believing Margin for Error. In that case, it is hard to see why Margin for Error would be attractive.

These challenges all suggest that the defender of Reflective Luminosity should accept Possible Knowledge instead of Possible Omega Knowledge. This preserves the inference from knowledge to justification, and predicts that Margin for Error is justifiably believed.

Unfortunately, this route faces its own cost: now you would be justified in believing dubious assertions.

(41) **Dubious Justification.** You justifiably believe that: p and you don't know that you know p.

The problem is that Reflective Luminosity allows that dubious assertions can be known; it only rules out their being omega known. So Reflective Luminosity imposes no barrier to possibly knowing dubious assertions. More carefully, as I illustrate in detail in Chapter 5, Reflective Luminosity allows the following: for all I know, I know (p and I don't know that I know p) So Reflective Luminosity would allow justified belief in dubious assertions. Again, in this book I am thinking of belief as subjective certainty. So the relevant prediction is that you could be justifiably certain of conjunctions like p and I don't know that I know that p.

Marušic 2013 and Greco 2014a claim that dubious assertions are as strange to believe as they are to assert. But another perspective denies that dubious assertions are strange to believe. Belief and assertion are different. Belief is not the 'inner analogue' of assertion. Assertion represents the speaker as omega knowing, while belief only represents the believer as knowing.

In summary, the defender of Reflective Luminosity must choose between allowing justified beliefs in both or neither of Margin for Error and dubious assertions. This depends on which of Possible Knowledge or Possible Omega Knowledge you accept.

2.3 Weakenings

Reflective Luminosity offers various benefits, but also has several costs. One strategy is to explore weakenings of Reflective Luminosity that preserve the benefits while avoiding the costs.

Reflective Luminosity is the strongest in a family of weakenings of KK. The next strongest member of this family says that if you know that you know p, then you omega know p. This weaker principle allows that you can know you know p without knowing you know you know p. But it says that there

is no difference between possessing three versus four iterations of knowledge. On this proposal, there are three kinds of knowledge: knowing without reflectively knowing, reflectively knowing without knowing you reflectively know, and omega knowing.

Generalizing, say that you n-know p when you possess n iterations of knowledge that p. Let n-Reflective Luminosity be the principle that n-knowledge implies omega knowledge:

(42) n-Reflective Luminosity. If you n-know that p, then you omega know that p.

KK is the special case where n=1. Reflective Luminosity is the special case where n=2.47

In Chapter 1, I considered the thesis that knowledge requires safe belief, and safe belief depends on counterfactual similarity. Such theses deny n-Reflective Luminosity for every value of n. No matter the value of n, I can construct a chain of n worlds where each is counterfactually similar to the next, and yet the last is not counterfactually similar to the last. The case of tree height, discussed in detail in Chapter 5, provides such a construction.

Those who reject Reflective Luminosity and accept n-Reflective Luminosity can preserve many of the benefits of Reflective Luminosity. The principle still allows for abundant omega-knowledge, provided that n-knowledge is also abundant.

If 3-Reflective Luminosity is true, the theory can still predict that three iterations of knowledge seems like the same thing as four iterations.

 $n ext{-Reflective}$ Luminosity can accommodate the unconfident examinee, since it denies the KK thesis, and so allows for knowledge without omega knowledge.

On the other hand, if Reflective Luminosity fails, then it is difficult to explain why interlocutors so rarely ask about second-order knowledge. No higher version of *n*-Reflective Luminosity will explain this, in the way that Reflective Luminosity was able to identify reflective knowledge with knowing for sure.

 $^{^{47} \}rm Bonnay$ and Égré 2009 §3 discusses n-Reflective Luminosity in detail. That paper provides a resource sensitive two-dimensional 'token' semantics for knowledge that can validates n-Reflective Luminosity. The system is stronger than any I discuss in my book, in two respects. First, it validates a principle that is strictly stronger than n-Reflective Luminosity. In the case where n is 2, this stronger principle says that: if p is false and it is epistemically possible that (p is true and also that it is epistemically possible that it is epistemically possible that q is true), then it is epistemically possible that (p is true and it is epistemically possible that q is true).Second, that system cannot validate Reflective Luminosity without also validating a version of negative introspection, which says that if it is consistent with what you know you know that p, then it is epistemically possible that you know that it is epistemically possible that p (p. 195-6). I reject this principle, because it faces similar challenges to negative introspection. For example, imagine you are deceived by an evil demon. You don't have hands. But for all you know, you are in the good case where you know you have hands. Moreover, you know that it is epistemically possible that you know you have hands, since this is always epistemically possible for you, no matter how bad your predicament. Since you don't have hands, it is consistent with what you know you know that you don't have hands. But it is not epistemically possible that you know that it is epistemically possible that you don't have hands. You can never know that you don't know you have hands, and you know this about yourself.

Higher orders of n-Reflective Luminosity offer even stronger versions of the error theory for Margin for Error. n-Reflective Luminosity is compatible with normally n-1-knowing that Margin for Error is true, while also knowing quite a bit about the tree's height. Nor does n-Reflective Luminosity lead to any difficulty regarding Justified Margins.

Now consider whether n-Reflective Luminosity can avoid some of the costs of Reflective Luminosity. First, n-Reflective Luminosity will allow for chains of testifiers to produce different amounts of knowledge at any order lower than n. For example, if 3-Reflective Luminosity holds and 2-Reflective Luminosity fails, then there could be Very Reliable and Less Reliable testifiers who generate reflective knowledge and ordinary knowledge when they testify to you about what you know. On the other hand, if such testifiers tell you about what you reflectively know, then the same problem arises. Generalizing, n-Reflective Luminosity will not allow chains of testifiers to differ in how much knowledge they produce regarding the claim that you (n-1)-know a proposition.

n-Reflective Luminosity can potentially avoid revenge versions of Murine Research and the unconfident examinee. Similarly, n-Reflective Luminosity can allow you to remember that you see something, without being certain of it. You can have up to n-1 iterations of factive mental states without achieving certainty. On the other hand, some readers may think that arbitrary chains of iterated factive mental states can be possessed without rational certainty. For example, perhaps you can remember realizing that you saw that you realized p, without being certain. Similarly, consider interpersonal iterations: perhaps even without being certain that it rained, you can remember that Alex realized that Billy saw that Claire remembered...that Danny was sad that it rained.

More generally, Omega Infallibilism offers one strategy for deciding which version of Reflective Luminosity to accept, if any. If omega knowledge is the unique state that licenses rational certainty, then the relevant question is whether there is some n such that n-knowing p licenses rational certainty. If there is, then n-Reflective Luminosity is valid. If you think that there are people who can n-know p without being permitted to act as if p, or who can continue to inquire about whether p even when they n-know p, then you should deny n-Reflective Luminosity (or else deny Omega Infallibilism).

One hypothesis is that n-Reflective Luminosity is first valid for some value of n that exceeds the working memory capacity of humans. n might be ten for example. On this picture, there are systematic counterexamples to the luminosity of knowledge, of known knowledge, of known knowledge, and so on. But once you achieve ten iterations of knowledge, you possess omega knowledge. On this proposal, omega knowledge is difficult but not impossible to possess. But there are tradeoffs. If n-Reflective Luminosity is first valid at a higher value of n, then various behavior governed by Norm Iteration becomes more demanding. In this way, deciding on the right version of Reflective Luminosity involves trading off the ability to express distinguish more kinds of knowledge against increasing rational demands.

The final cost of Reflective Luminosity was that it led to bad consequences for the theory of justified belief. If Possible Omega Knowledge is true, it leads

to the failure of knowledge to imply justification, and to the failure of justifiably believing Margin for Error. If Possible Knowledge is true, then it allows that dubious assertions can be justifiably believed.

Weakening Reflective Luminosity to n-Reflective Luminosity does not avoid these problems. The problems emerged because Reflective Luminosity is too weak rather than too strong. In the next chapter, I'll consider the Fragility principle, which has exactly the kind of logical strength needed to make the right predictions in these cases. (Fragility will also have significant costs of its own.)

Another variant of Reflective Luminosity applies to ignorance rather than knowledge:

(43) **Negative Reflective Luminosity.** If you know that you don't know p, then you know that you know that you don't know p.

Negative Reflective Luminosity says that knowledge of ignorance is luminous. Negative Reflective Luminosity is another weakening of KK. It says that a special kind of knowledge is luminous: knowledge of ignorance. When combined with Reflective Luminosity, Negative Reflective Luminosity implies that if you know that you don't know p, then you omega know you don't know p.

Negative Reflective Luminosity is also weaker than Negative Introspection, the principle that if you don't know p, you know you don't know p. Negative Reflective Luminosity avoids the most serious problem for Negative Introspection. Take any claim p that you justifiedly believe you know. Now consider a skeptical scenario in which you are deceived by an evil demon, and p is false. In this scenario, you believe p, but you do not know p. Yet you don't know that you don't know p. After all, you believe you do know p; so how could you know you don't know p. Moreover, you don't even seem to be in a position to know that you don't know p; after all, you are justified in believing you know p. Negative Reflective Luminosity avoids this issue, since its antecedent is false in skeptical scenarios. Negative Reflective Luminosity is part of a broader package according to which knowledge of our own ignorance and knowledge is part of our 'cognitive home'.

Another weakening of Reflective Luminosity restricts the principle to certain canonical methods of knowing, such as perception, induction, or testimony. Imagine that you know that Jones knows that p, but don't know whether Jones knows that Jones knows that p. Now you learn that you are Jones. In that case, you come to know that you know that p, and so Reflective Luminosity implies that you omega know that p. But in learning that you are Jones, your epistemic position with respect to p does not seem to improve. Here, one might reject the application of Reflective Luminosity, by claiming that in this case you do not know that you know that p by relying on a canonical belief forming method, such as perception. Rather, perhaps in this case perception only gives you knowledge of p, and knowledge that Jones knows that p, while your higher order knowledge is generated by a further deductive argument that relies on the premise that you are Jones, and that Jones knows that p.

Restricting Reflective Luminosity to canonical belief forming methods can

help avoid some of the challenges to Reflective Luminosity. Chains of testimony may produce different amounts of knowledge, if testimony about what you know is non-canonical. This strategy may also help avoid revenge versions of the unconfident examinee. Imagine that you are in group therapy for obsessive compulsive disorder. The goal of the therapy is to come to believe and know that you are protected from germs. You know that the evidence regarding germ protection is strong, and so you know that if you believe you are protected from germs, then you know you are protected from germs. The problem is you aren't sure whether you believe you're protected. Some people in the group haven't progressed at all far enough to believe; other people in the group are right below the threshold for believing; others are right above the threshold; and others are far above the threshold. In fact, you are far above the threshold for believing, but you don't realize this is the case. Nonetheless, when someone asks whether you believe, you wager the guess that you do believe you are protected. Like the regular unconfident examinee, your belief is safe from error, and plausibly could count as knowledge. But then by deduction you could come to know that you know that you are protected, since you know that if you believe you are protected, then you know you are protected. Nonetheless, plausibly you fail to omega know that you are protected, for structurally analogous reasons to why the unconfident examinee fails to know that she knows. As with the previous case, one could resist the application of Reflective Luminosity by claiming that your method of knowing that you know you are protected from germs is not canonical, because it is a strange concatenation of deduction and reliable guessing.⁴⁸

2.4 Conclusion

In this chapter, I defended Reflective Luminosity, the thesis that if you know you know p, then you know you know you know p. I argued that this principle offers a compromise between defenders and opponents of KK. It explains the role of omega knowledge in assertion and other kinds of behavior. But at the same time, Reflective Luminosity allows for counterexamples to KK, and makes sense of inexact knowledge.

I outlined a variety of benefits and costs of Reflective Luminosity. In my eyes, the main benefits of Reflective Luminosity are that it avoids Omega Skepticism, it seems to capture exactly the range of cases where KK fails, and it offers a plausible account of inexact knowledge, by predicting that Margin for Error is normally true, and that you are justified in believing Margin for Error. On the other hand, the main costs for Reflective Luminosity were that it faced a variety of potential 'revenge' counterexamples, where higher-order knowledge did not generate rational certainty, and it also made bad predictions about justified belief.

In the face of these benefits and costs, it is difficult to assess the overall plausibility of the principle. In the next two chapters, I turn to considering two alternative theories of omega knowledge. In each case, I'll identify a collection of

⁴⁸Thanks to John Hawthorne for discussion.

benefits and costs of each theory. Then, in the book's final chapter, I'll present an overall summary of the benefits and costs of each of the three theories discussed in the book.

3 Fragility

In this chapter, I explore 'Fragility', which says that if you know p, then for all you know, you omega know p. I begin by clarifying what Fragility says, and how it relates to KK and Reflective Luminosity. Then I survey the benefits and costs of Fragility. I argue that Fragility, like Reflective Luminosity, offers a compromise between opponents and defenders of KK. Like Reflective Luminosity, it avoids Omega Skepticism, and explains the role of omega knowledge in assertion and other behavior. Like Reflective Luminosity, it is consistent with counterexamples to KK.

Whether Fragility is true depends on the nature of knowledge. In Chapter 5, I explore the prospects for Fragility within a theory of knowledge as belief that is safe from error. I interpret the relevant kind of safety as belief that is normally true, and I show how a particular interpretation of normality leads to the validity of Fragility. In particular, the model of knowledge I develop proposes that there is an epistemic asymmetry between the good case (where your belief forming processes are functioning normally) and the bad case. In the model, Fragility says that in the good case where conditions are most normal, you omega know anything that you know. But in worse cases, you can know things that you don't omega know. The model makes two more claims. First, no matter what condition you are in, it is always possible for you that you are in the good case, where reality matches appearance. Second, in the good case, you know at least as much as in any other case. The result is Fragility. Suppose you know p. Then it is possible that you know p in the good case. So it is possible that you know that you know p (and omega know p).

3.1 Clarifying Fragility

Say that p is epistemically possible for you just in case p is true for all you know, just in case p is consistent with what you know, just in case you don't know that p is false. Then Fragility represents a kind of optimism about the possibility of iterated knowledge.

(44) **Fragility.** If you know p, then for all you know, you omega know p.

Fragility allows you to know p without omega knowing p. But Fragility says that even in this case, it is epistemically possible for you that you omega know p. Fragility is optimistic, because it says that when you know p, you can never rule out the possibility that you are in the best epistemic position of omega knowing p.

Fragility provides an antidote to Omega Skepticism. Suppose that skepticism is false, so that you know something. Then you do not know that you fail to omega know p. This requires that in some possibility consistent with your knowledge, you omega know p. So omega knowledge must not be radically scarce.

On the other hand, the strict letter of Fragility allows that you know much while omega knowing very little. In Chapter 5, I develop a model of Fragility that

explains in detail how Fragility avoids Omega Skepticism. The key property of the model is that KK holds in the good case where your belief forming processes are functioning normally. In my model, Fragility says that in the good case, you omega know anything that you know. In this way, the model avoids Omega Skepticism by saying that when conditions are most normal, you omega know many things.

The contraposition of Fragility says that:

(45) If you know that you don't omega know p, then you don't know p.

(45) says that knowledge is fragile, because (45) articulates a connection between knowledge and defeat. Say that q is a defeater for p iff q is true, you know p, and if you were to learn q, then you would lose your knowledge of p. Fragility says that propositions about higher order ignorance are defeaters. If you learn that you don't omega know p, you learn that you are in some way epistemically defective with respect to p. Fragility says that if you learn that you are epistemically defective with respect to p, this knowledge defeats your knowledge of p. Knowledge of p is fragile in the face of evidence that one is not epistemically ideal with respect to p.

Fragility is strictly weaker than KK, since it replaces knowing that you know with the epistemic possibility of knowing that you know. Fragility allows that you can know p without knowing that you omega know p. But things are different if you become aware that you are in such a predicament. If you know p and then learn you don't know omega know p, something changes. New information about your non-ideal status destroys your knowledge of p.

Fragility allows that it is possible for you to know p while failing to know that you omega know p. So the kind of defeat above is not simply logical: it is not a case where you learn p and then lose your knowledge of q because q is inconsistent with p.

Fragility distinguishes the first and third person perspective. I can know that you know p and that you don't omega know p. But you can't know that you know p and that you don't omega know p. I can know something about you that you can't know about yourself. This is analogous to Moorean claims. I can know that it is raining and you don't know its raining. But you can't know that it is raining and you don't know it's raining. (Reflective Luminosity also creates an asymmetry between yourself and others. Reflective Luminosity allows you to know that I know p, without omega knowing that I know p. But Reflective Luminosity says that if you know that you know p, then you omega know that you know p.)

Fragility is neither stronger nor weaker than Reflective Luminosity; the two conditions are independent. In Chapters 5 and 6, I explore in detail the differing constraints that Fragility and Reflective Luminosity impose on knowledge (along with another principle, Variable Margins, introduced in Chapter 4). I don't take a stand on which principle is best. But I believe that they should not all be accepted.

3.2 Benefits of Fragility

In this section, I'll consider the benefits of Fragility. In the next section, I consider its costs.

First benefit: Fragility explains how Omega Assertion and Norm Iteration can be true.

Fragility can explain the role of omega knowledge in assertion and other behavior.

Consider assertion. Recall that dubious assertions are sentences of the form p and I don't know that I know p. Sentences of the form p and I don't n-know p are infelicitous for any value of n. This suggests that omega knowledge plays an important role in assertion.

In Chapter 1, I embraced Omega Assertion, the principle that you should assert p only if you omega know p. Surprisingly, Fragility can explain the infelicity of omega knowledge without Omega Assertion. Fragility only requires that you should not assert p when you fail to know p.

The key is that Fragility is equivalent to the thesis that dubious assertions are unknowable:

(46) **Ignorance of the Dubious.** You don't know that: p and you don't omega know that p.

Suppose you know the conjunction: p and you don't omega know that p. Then you know p and you know that you don't omega know that p, contradicting Fragility. Conversely, suppose that Fragility fails. Then you know p while knowing that you don't omega know that p. But then you can conjoin this knowledge, to learn the conjunction [p] and you don't omega know p, contradicting Ignorance of the Dubious. Fragility says that dubious assertions are unknowable. So Fragility explains why dubious assertions of the form p and p and p and p can't be asserted.

In this way, Fragility predicts that a whole family of dubious assertions are unknowable. Consider the following chain of dubious assertions:

- (47) a. p but I don't know that I know that p.
 - b. p but I don't know that I know that I know that p.
 - c. p but I don't know that I know that I know that I know that p.
 - d. ...

Fragility implies that p and I don't n-know that p is unknowable, for every value of n. So it explains the infelicity of each order of dubious assertion.

Other dubious assertions involve mixed attitudes of belief and knowledge:

- (48) a. p but I doubt that I know that p.
 - b. p but I believe that I don't know that p.
 - c. p but I have no justification for believing that I know that p.

d. p but I have (sufficient) justification for believing that I don't know that p.

Fragility implies that each of the conjunctions above is unknowable. In each case, the argument is roughly the same: the iterated state in the second conjunct of the dubious assertion is logically as strong as the state of not knowing that one knows. For this reason, assuming that knowledge is closed under simple deduction, knowing any of these conjunctions implies knowing the dubious assertion p and I don't know that I know p.

Start with (48-a). Knowledge is incompatible with doubt. So if you doubt that you know that p, then you don't know that you know p. So if you know that p and that you doubt that you know p, then you know that you doubt that you know p. But since this last bit of knowledge implies that you don't know that you know p, it follows that you know that you don't know that you know p. This contradicts Fragility, since you also know p. In short, this complex assertion is logically stronger than p and p and p and p and p are the weaker dubious assertion is unknowable, so is the stronger.

The same argument applies to each of the other dubious assertions above. For (48-b): if you believe that you don't know that p, then you don't know that you knows p. For (48-c): if you aren't justified in believing that you know p, then you don't know that you know p. For (48-d): if you're justified in believing you don't know p, then you don't know that you know p. (Reflective Luminosity also explains the infelicity of each of these sentences.)

Strictly speaking, Fragility can explain dubious assertions without Omega Assertion. But there is a good reason for the Fragility defender to accept Omega Assertion. Suppose that you accept a knowledge norm without accepting Omega Assertion (denying Norm Iteration). Greco 2014b worries that without KK, dubious assertions would be indefensible: 'in such cases, while [you] will be able to permissibly assert that p, if [your] permission to assert that p is challenged, [you] will not be able to permissibly defend [yourself].' More precisely, say that p is an indefensible assertion if you are permitted to assert p but forbidden to assert that you know p. Now take a case where you know p without knowing you know p. Since you know p, the knowledge norm allows you to assert p. Since there is no higher norm on assertion such as Omega Assertion, this assertion would be permissible. But now if someone asks you whether you know p, you cannot answer yes. The view permits indefensible assertion.

By contrast, suppose Omega Assertion holds. Then there are no indefensible assertions. If you are permitted to assert p, then you omega know p, and so are permitted to assert that you know p. Because Omega Assertion is true, the case above is not a dilemma. If you know without omega knowing, you simply ought

⁴⁹The forced march Sorites is another potential example of indefensible assertion. As you consider a series of decreasingly hairy men, you are permitted to call some intermediate member bald. But you are not always permitted to answer 'yes' when asked if the subsequent member of the series is bald; otherwise, you would end up permissibly claiming that someone with a full head of hair is bald. In this way, the Sorites paradox seems to create situations in which you are permitted to assert something, even though you are not permitted to always acknowledge further pragmatic commitments of what you asserted.

not assert.

Again, in the presence of Omega Skepticism, Omega Assertion would lead to skepticism about permissible assertion. Fragility denies Omega Skepticism, saying that it is always possible that you omega know, if you do know. In this way, Fragility combines with Omega Assertion to explain the impossibility of indefensible assertion.

Summarizing, Fragility can explain the infelicity of dubious assertions without appeal to Omega Assertion. But Fragility needs Omega Assertion in order to avoid indefensible assertions, cases where you are permitted to assert p and forbidden to assert you know p. (Norm Iteration provides a more systematic explanation of why indefensible assertion is impossible.)

In Chapter 1, I considered the relevance of omega knowledge to a wide class of behavior beyond assertion. In particular, I explored 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 p (and only if you omega know p). I also considered the application of Norm Iteration to behaviors like action, inquiry, and certainty. For example, I considered the thesis that you should rely on p as a reason only if you omega know p; that you should stop inquiring about p once you omega know p; and that you should be certain of p if and only if you are certain of p.

These views are untenable if Omega Skepticism is true. In that case, you could never rely in action on any ordinary claim, and you could never stop inquiring into anything. You could never be certain of anything.

Fragility blocks Omega Skepticism, and thereby makes room for permissible action, inquiry, and certainty in this setting. When you know, it is possible that you omega know. So when you know p, it is epistemically possible that you would be rational in relying on p in action. (Again, in Chapter 5 I develop a model of Fragility according to which in the good case you omega know everything that you know. According to this model, being in the good case gives you a special kind of permission to act and end your inquiry.)

Finally, consider common knowledge. My model of Fragility, developed in Chapter 5, predicts that when you are in the good case and you know p, you omega know p. This means that in the good case, a group can commonly know that p.

On the other hand, I allow that when you are not in the good case, you can know something without knowing that you know it. This means that common knowledge is difficult to obtain when you are not in the good case. On this proposal, then, common knowledge is a special state that we attain when our environment is favorable enough. Nonetheless, Fragility can predict that even when a group is not in the good case, every member of the group can be justified in believing that the group has common knowledge, and so in this way can still be justified in coordinating.

Second benefit: Fragility avoids revenge versions of Murine Research and the unconfident examinee.

Fragility provides an interesting response to counterexamples to KK. One threat to KK concerns simple agents who know without having the concept of knowing. This is no immediate threat to Fragility. Suppose dogs fail to know that they know because they lack the concept of knowledge. By the same token, they fail to know that they don't know that they know. They trivially satisfy Fragility.

In the last chapter, we saw that Reflective Luminosity struggles with revenge versions of Murine Research and the unconfident examinee. With these revenge cases, we imagine an agent who has some amount of higher-order knowledge, without permissible certainty or omega knowledge. These kinds of cases are compatible with Fragility, since Fragility allows you to reflectively know without omega knowing.

Third benefit: Fragility avoids other costs of Reflective Luminosity, involving chains of testifiers with different reliability, and iterated factive mental states.

Fragility also avoids other challenges for Reflective Luminosity. Recall that Reflective Luminosity made controversial predictions about testimony, denying the possibility of pairs of testifiers whose different reliability produces systematically different iterations of knowledge. Fragility has no problem here, since it allows that you can possess any number of iterations of knowledge without possessing a further iteration of knowledge.

Similarly, Reflective Luminosity implies that whenever you remember seeing that p, you are permitted to be certain of p. The problem is that remembering and seeing each seem to imply knowing. Fragility has no such consequence, since it allows that you can know that you know something without being permitted to be certain. (On the other hand, Fragility and Omega Infallibilism do imply that you would be justified in being certain in any such case, since for all you know you would omega know.)

Fourth benefit: Fragility is compatible with systematic anti-luminosity.

Another difference between Fragility and Reflective Luminosity concerns the status of luminosity arguments. Williamson 2000 has argued that no interesting mental state is luminous, in the sense that you know you are in the state whenever you are in it. Fragility is compatible with this conclusion. But Reflective Luminosity denies it, because it says that reflective knowledge is luminous.

Fifth benefit: Fragility explains why KK is attractive, by predicting that it is normally true.

In the previous chapter, I showed that Reflective Luminosity explains the attraction of Margin for Error, by predicting that it is normally true. Fragility offers a corresponding benefit. My model of Fragility that I'll develop in Chapter 5 says that the KK principle is normally true. Normally, if you know p, then you omega know p. This offers a kind of error theory for KK. KK seems true, because it is normally true. In particular, say that something is normally true when it is true in any situation that is almost as normal as the most normal one. In my model of Fragility, KK will hold in all such situations.

On the other hand, no theory should say that both KK and Margin for Error are both normally true. At least, this result should be avoided if the set of normally true claims is closed under logical deduction. After all, normally you know something about a tree's height, when you're looking at it and it appears 100 feet tall and is so. Say that you normally know that it is at least 90 feet tall. If KK is normally true, it follows that you normally know that you know it is at least 90 feet tall. But if Margin for Error is normally true, it would follow that you normally know that it is at least 91 feet tall. In this now familiar way, the normal truth of both KK and Margin for Error would lead to paradox. For this reason, the theorist about normality must choose between KK and Margin for Error.

Sixth benefit: Fragility fits well with Possible Omega Knowledge and Possible Knowledge together.

In Chapter 1, I explored two principles about justification. Possible Omega Knowledge says that you are justified in believing p iff for all you know, you omega know p. Possible Knowledge says that you are justified in believing p iff for all you know, you know p. In Chapter 5, I develop a model of Fragility. Again, the model will say that in the good case, you know anything that you omega know. In this model of Fragility, Possible Knowledge and Possible Omega Knowledge are both valid: you have possible knowledge iff you have possible omega knowledge iff you have knowledge when conditions are most normal.

One worry for Fragility is that it doesn't really vindicate the importance of omega knowledge. Granted, it implies that it is possible for you that you omega know many things. But it doesn't actually imply that you do omega know many things. But this means that Fragility doesn't actually imply that you satisfy the requirements of assertion and other kinds of behavior.

When supplemented with Possible Omega Knowledge, Fragility provides an answer to this worry. In Chapter 1, I considered the idea that justification is a kind of excuse. Possible Permission said that you are justified in A when you are in a state that is epistemically indistinguishable from being permitted to A. In Chapter 1, I explored omega knowledge norms. In the presence of

those norms and Possible Permission, Fragility then says that whenever you know p, you are justified in asserting p, acting as if p, and ending inquiry about p. In addition, Fragility implies that if you know p, then you are justified in believing p, in the sense that you are justified in being certain that p. Even more, Fragility implies that if you know p, then you are justified in believing that you omega know p. Applied to common knowledge, Fragility (especially as modeled in Chapter 5) therefore allows that when a group knows p, the group is justified in believing that the group has common knowledge of p, and so is justified in acting in whatever ways are permitted by common knowledge. In this way, Fragility offers a vindication of the importance of omega knowledge, even in cases when you don't actually have much omega knowledge. By contrast, omega skeptics think they know that they don't omega know ordinary claims. So omega skeptics are not entitled to analogous excuses.

The theory of justification also has interesting consequences for dubious assertions. Omega Assertion, Fragility, and Possible Permission imply that any apparent case of indefensible assertion is actually a justified but impermissible assertion. Suppose that you know p without knowing that you know p. By Omega Assertion, you are not permitted to assert p. But by Fragility, it is possible that you omega know p. If omega knowledge is sufficient for permissible assertion, it follows that it is possible that you are permitted to assert p. By Possible Permission, it follows that you are justified in asserting p (and in asserting that you know p).

Fragility also has interesting consequences for the logic of justification. Given Possible Omega Knowledge, Fragility is equivalent to the thesis that knowledge implies justification. Fragility says that if you know p, then for all you know, you omega know p. Possible Omega Knowledge says that you are justified in believing p iff this condition holds. So Fragility essentially just says that knowledge requires the special kind of justification articulated by Possible Omega Knowledge.

In the previous chapter, I showed that Reflective Luminosity faces some challenges in giving a theory of justified belief. The problem was that if it accepted Possible Omega Knowledge, it would end up denying that knowledge implied justification; but if it accepted Possible Knowledge instead, it would predict that dubious assertions are justifiably believed. Fragility resolves this dilemma. In its presence, Possible Omega Knowledge and Possible Knowledge are equivalent, knowledge implies justification, and dubious assertions cannot be justifiably believed.

3.3 Costs of Fragility

First cost: Fragility distinguishes higher iterations of knowing.

In the last chapter, I argued that one benefit of Reflective Luminosity was that it predicted that three iterations of knowledge was the same as four. Fragility does not have this consequence; it allows you to possess any number of iterations

of knowledge without having omega knowledge.

Second cost: Fragility doesn't explain why it is strange to ask people whether they know that they know.

Second, Fragility does not explain why it is more common to ask people what they know, compared to asking people what they know that they know. As I showed in the last chapter, Reflective Luminosity can explain this, because questions about second-order knowledge are equivalent to perfectly commonplace questions about what you know for sure. By contrast, Fragility distinguishes what you know for sure from what you reflectively know, and so should predict that it would be normal to ask about reflective knowledge.

Third cost: Fragility makes strange predictions about the unconfident examinee.

Third, Fragility faces trouble from people who know something without believing that they know it. One such case might be the unconfident examinee, who can reliably guess when Queen Elizabeth ruled, but doesn't remember ever learning this information. Although she knows when Queen Elizabeth ruled, she doesn't know that she knows this. Fragility implies that there is something unstable about the examinee's predicament. If she is apprised of her higher order ignorance, she either loses her ordinary knowledge or gains reflective knowledge. The sense in which her position is unstable is that there is a truth which would destroy her knowledge if she learned it.

To test this prediction, imagine that the unconfident examinee was confronted with an omniscient angel, who tells her that she don't know that she knows when Queen Elizabeth died. The examinee then asks the angel if that means that she don't know when Queen Elizabeth died. The angel replies that she isn't saying this; rather, she's merely communicating the information that she doesn't have higher order knowledge about Queen Elizabeth. Fragility predicts that this testimony would destroy the examinee's first-order knowledge.

Fragility makes even more surprising predictions. Fragility allows that the unconfident examinee believes that she doesn't know that she knows. Fragility insists, however, that the unconfident examinee can't know that she doesn't believe she knows, if she does in fact know. After all, knowledge implies belief. So knowing you don't believe you know implies knowing you don't know you know.

This is at first glance a strange prediction, which may favor Reflective Luminosity over Fragility. In many ordinary cases, the presence and absence of belief is luminous. If you don't believe, you know you don't believe. But surely in some ordinary case where KK fails, you know p without believing you know p. Any such case is a potential counterexample to Fragility.

There are a few ways of defending Fragility. First, I could modify the target of Fragility: it doesn't govern knowledge, but rather governs being in a position

to know. You are in a position to know p when the only thing preventing you from knowing p is the absence of belief. In this way, being in a position to know stands to knowledge as propositional justification stands to justified belief. Then the principle says that if the reliability of your belief forming processes, the safety of your environment, the strength of your evidence and whatever else put you in a position to know p, then they do not put you in a position to know that you are not in a position to omega know that p. On this view, you can know p while knowing that you don't actually omega know p, on the grounds that you don't believe you know p. But even though you know you don't actually omega know p, it is possible for you that you are in a position to know that you are in a position to omega know p. That is: it is possible for you that even though you don't actually believe you know p, if you did believe you know p you would omega know p.

Another response insists that Fragility governs knowledge, and says that the absence of belief is hard to know about. Granted, you can know p without believing you know p. But when you do so, you don't know that you don't believe you know p.

Another response says that when you know p, you believe you know p. One way to accept this theory is to embrace the thesis that belief is strong, so that whenever you believe p, you believe you know p. (Later, I consider the analogous thesis about justified belief.) Why would believing imply believing you know? Greco 2014b offers one model. Suppose KK is valid, and define belief as a state subjectively indistinguishable from knowing. To believe is to be in a state that 'feels the same' from the inside as knowing. Now suppose you believe p. Then you are in a state that feels the same from the inside as knowing. But if KK is valid, then you are also in a state that feels the same from the inside as knowing you know. So you believe you know. This argument assumes KK, and so is unavailable to me. But even without KK, you could accept the principle as a thesis about subjective indistinguishable. The idea would be that whenever your state feels the same as knowing, it also feels the same as knowing you know. This could be true even if knowing is not the same as knowing you know. It could also be true even if knowing does not always feel the same as knowing you know. (In Chapter 5, I develop a model of epistemic indistinguishability that has exactly the relevant feature: whenever it is epistemically possible that you know, it is epistemically possible that you know that you know. On the other hand, I think that subjective and epistemic indistinguishability are different.)

As we've seen, Fragility explains the infelicity of dubious assertions, where you try to assert while acknowledging that you don't believe that you know. So some of the motivation for Fragility extends directly to the thesis that if you know p, you don't know that you don't believe that you know p. Any theory of knowledge faces a challenge in reconciling the apparent possibility of knowing without believing that you know with the facts that it is easy to know what you do or don't believe, and that it is bizarre to assert p while conceding that you

⁵⁰See for example Goldman 1979.

⁵¹For more discussion of the notion of being in a position to know, see Marušic 2013, Rosenkranz 2018, Willard-Kyle 2020, and Yli-Vakkuri and Hawthorne forthcoming.

don't believe you know.

Fourth cost: Fragility does not predict that Margin for Error is normally true, or justifiably believed.

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:

(49) **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.

Suppose for example that you know the tree is not 100 - 1 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 100 feet tall. So you know that you know the tree is not 100 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.

Although Fragility is weaker than KK, it has a similar consequence regarding Margin for Error. Fragility leads to absurdity when combined with knowledge of Margin for Error. Suppose that you know the tree is not 100-1 feet tall, and suppose you know Margin for Error. If knowledge is closed under competent deduction, then you know the conjunction of these claims. But then by Fragility it is possible that you omega know their conjunction. By closure, you could then deduce that the tree is not 100 feet tall. It follows that it is possible that you omega know that the tree is not 100 feet tall. Now suppose that possible omega knowledge (or justified belief) agglomerates (Chapter 5 develops models in which this agglomeration principle holds). Since you know Margin for Error, it is possible that you omega know Margin for Error. It follows from agglomeration that it is possible that you omega know that: you know the tree is not 100 feet tall, and Margin for Error holds. It follows that it is possible you omega know the tree is not 100 + 1 feet tall. Iterating this reasoning, it follows that for every height of the tree, it is possible that you omega know the tree is that height. But that is absurd. After all, omega knowledge is factive: if you omega know something, it is true. So if it possible that you omega know p, then it is possible that p. So it would follow that your knowledge is consistent with the tree being any height. To avoid this result, defenders of Fragility must deny that Margin for Error is known.

Fragility differs from Reflective Luminosity regarding the status of Margin for Error. In Chapter 2, I argued that Reflective Luminosity can preserve the appeal of Margin for Error, because it can predict that Margin for Error is normally true, that Margin for Error is known when conditions are most normal, and that you can justifiably believe Margin for Error. (I defend these claims in detail in Chapter 5.) Fragility cannot validate analogous principles. The problem is that if you ever knew Margin for Error, Fragility would imply that for all you know you would omega know Margin for Error. But the reasoning above would then imply that for all you know, you know nothing about the tree's height. This is

absurd. In this way, Fragility requires a more extreme departure from Margin for Error.

In Chapter 5, I defend Fragility from the charge of flagrant Margin for Error violation. I argue that it violates Margin for Error less flagrantly than the KK principle. My defense will center around cases of cliff-edge knowledge, where the tree is x feet tall and you know that it is at least x feet tall. Fragility requires some cases of cliff-edge knowledge. But it also allows that for many heights of the tree you do not have cliff-edge knowledge. By contrast, I'll show that existing models of KK imply that you have cliff-edge knowledge of the tree's height whenever the tree's real height departs sufficiently from its apparent height. In this way, Fragility mitigates some of the more absurd features of KK regarding Margin for Error. More generally, my model of Fragility will turn out to validate Margin for Error in bad cases where conditions are not maximally normal, rather than in the best case where conditions are most normal. In this way Reflective Luminosity and Fragility both deny Margin for Error in some cases and accept it in other cases; but the principles disagree about the range of cases in which the principle holds.

As discussed in Chapter 2, my model of Reflective Luminosity in Chapter 5 will predict that normally, Margin for Error is true, because Margin for Error is known in the good case. By contrast, the defender of Fragility must deny that normally Margin for Error is true. Instead, the defender of Fragility says that normally KK is true.

Since Fragility requires that you don't know Margin for Error, Fragility implies that Margin for Error can possibly fail. But I showed in Chapter 1 that Margin for Error follows from Safety, combined with natural assumptions about what could easily have been the case.

- (50) **Safety.** If you know p, then you could not easily have believed p falsely.
- (51) **Possibility.** For any height x, if the tree is x feet tall and you believe it is not x-1 feet tall, then it could easily have been x-1 feet tall while you believed it was not x-1 feet tall.

The model of Fragility I develop in Chapter 5 responds to the argument for Margin for Error by preserving Safety and rejecting Possibility. I do so by interpreting Safety in terms of normality rather than counterfactual similarity.

Fifth cost: Fragility faces a challenge from Memory Experiment, a case where you seem to know without being justifiably certain.

Recall that Possible Omega Knowledge followed from the thesis that belief is subjective certainty, subjective certainty is permitted iff you omega know (Omega Infallibilism), and you are justified in believing whatever you are for all you know permitted to believe (an instance of Possible Permission). Given Omega Infallibilism, Fragility says that if you know p, then for all you know you are permitted to be subjectively certain of p. In this way, Fragility leads to a

weak kind of infallibilism about knowledge. Granted, you can know p without being permitted to be subjectively certain of p. But whenever you know p, it is epistemically possible for you that you are permitted to be subjectively certain of p. Similarly, Omega Infallibilism, Possible Permission, and Fragility imply that if you know p, then you are justified in being subjectively certain of p. That is: knowledge implies justified certainty.

there are potential counterexamples to the thesis that knowledge implies justified certainty. Consider the following case from Holguín forthcoming.

(52) Memory Experiment. Joan and Megan are participating in a trial of a drug whose primary effect is to swamp its subjects with an extraordinary number of fake "memories" of the events of the past 24 hours. One of the subjects will get the drug, while the other will get a placebo. Who gets which is determined by a coin-flip whose result is known only to the experimenters. During the experiment Joan and Megan are both (separately) asked 'Do you remember what you ate for dinner yesterday?' Joan appears to remember that she ate fish; Megan appears to remember that she ate spaghetti. As a matter of fact it was Joan who got the placebo and Megan who got the drug. (Holguín forthcoming, p. 5.)

In this case, some judge that Joan does remember eating spaghetti. If remembering entails knowing, then this implies that she knows she ate spaghetti. But Joan plausibly knows that she should have a roughly .5 credence that she ate spaghetti yesterday; after all, she knows that there is a 50% chance she ingested the drug. But recall that Omega Infallibilism says that if Joan omega knows she ate spaghetti last night, then Joan is permitted to be certain that she ate spaghetti. So if Joan knows that she is not permitted to be certain that she ate spaghetti, then she knows that she doesn't omega know that she ate spaghetti. While Fragility struggles to explain this case, Reflective Luminosity has no special problem, since it does not connect first-order knowledge to possible omega knowledge or justified certainty. On the other hand, Reflective Luminosity is threatened by revenge versions of Memory Experiment, in which Joan appears to remember seeing that she ate fish, and so potentially knows that she knows she ate fish.

One response to this case denies that Joan does remember eating spaghetti; instead, she only seems to remember. A second response denies that remembering entails knowing. A third response introduces a second notion of rational credence. In accepting Omega Infallibilism, I have suggested that your rational credences should match the prior probability of hypotheses, conditionalized on what you omega know. But in cases where you don't know what you omega know, another notion of rational credence is possible. In such cases, you could instead match your credences to the expectation of evidential probability, weighted by the probability that you omega know various things. The prior probability of taking the placebo was .5. So in this second sense of rational credence, Joan's credence that she ate the spaghetti should be weighted by the 50% chance she took the placebo and the 50% chance that she took the drug. This will produce a credence

of roughly .5 that she ate spaghetti. (This leaves unsettled the conditions under which Joan omega knows that she ate spaghetti. This response is under some pressure to deny that she automatically omega knows that she ate spaghetti if she took the placebo. If she did omega know, then she would omega know that she omega knows, and so the probability that she omega knows would be 1. Rather, one way of developing this proposal is that when she took the placebo, it is consistent with her knowledge but not entailed by her knowledge that she omega knows.)

Sixth cost: Fragility implies that whenever you know something, you justifiably believe that you know it.

Fragility has another interesting consequence that some will view as a cost. Given Possible Knowledge or Possible Omega Knowledge, Fragility implies:

(53) **JK.** If you know p, then you are justified in believing you know p.

Given Possible Knowledge, JK says that if you know p, then it is possible that you know that you know p. Given Possible Omega Knowledge, JK says that if you know p, then it is possible that you omega know that you know p. Since omega knowledge implies every iteration of knowledge, this is equivalent to the condition that if you know p, then it is possible that you omega know p. (Given Possible Omega Knowledge, JK is equivalent to the thesis that knowledge implies justification.)

Berker 2008 endorses JK. Berker calls a state 'lustrous' when you are justified in believing that you are in the state whenever you are in fact in the state. He suggests that knowledge is lustrous, even if it is not luminous. Berker 2008 observes that the Margin for Error arguments against KK do not immediately extend to JK, because justified belief does not satisfy the Margin for Error principle.

KK implies that if you know p, then you are justified in believing you know p. After all, knowledge implies justification, and so knowing you know implies justifiedly believing you know. In this way, whenever a state is luminous it is also lustrous.

Possible Knowledge and Reflective Luminosity do not imply that knowledge is lustrous. Reflective Luminosity allows that you can know p even when you know that you do not know that you know p (as I illustrate in Chapter 5).⁵²

 $^{^{52}\}mathrm{On}$ the other hand, Fragility is distinct from the principle that if you know p, then you know you are justified in believing p. Given Possible Knowledge, this says that whenever you know p, you know that it is possible that you know p. This principle follows from Fragility given the assumption that justified belief is consistent, so that it is never possible that you know p and also possible that you know not p. But even theories that deny Fragility may also accept this principle. For example, in Chapter 5 I discuss models of knowledge that deny KK and Fragility, but accept that justification is luminous, so that whenever you are justified in believing p, you know you are justified in believing p. This model rejects Fragility but accepts that if you know p, then you know you are justified in believing p.

3.4 Weakenings

Before concluding, I consider a few weakenings of Fragility, and explore whether they retain the benefits of Fragility while avoiding its costs.

Fragility is the strongest of a family of principles. The weakest principle says:

(54) **Weak Fragility.** If you know p, then it is possible that you know that you know p.

Weak Fragility says that when you know, it is possible that you know that you know. But it allows that you can know p while knowing that you don't omega know p. Fragility implies Weak Fragility, but not vice versa.

Generalizing from Weak Fragility, again say that you n-know p when you possess n iterations of knowledge regarding p. Then there is a family of Fragility principles:

(55) n-Fragility. If you know p, then for all you know, you n-know p.

Weak Fragility is the same as 2-Fragility. Fragility implies n-Fragility for every value of n. When you omega know p, you n-know p for every value of n. So when it is possible that you omega know p, it is possible that you n-know p for every value of n.⁵³ Fragility encodes the idea that your knowledge of p is defeated by any information that your epistemic position with respect to p is not ideal. But here you might distinguish different degrees of epistemic ideality. Failing to know that you know p is not ideal. Failing to know that you know that you know p is not ideal in another way. Perhaps the first failure defeats knowledge in a way that the second does not. In this case, one might accept Weak Fragility but deny Fragility.

Weak Fragility avoids the challenge from Memory Experiment. Even when combined with Omega Infallibilism and Possible Omega Knowledge, it escapes the conclusion that knowledge implies justified certainty. For this reason, it allows Joan to remember eating spaghetti, without having justification for begin certain that she ate spaghetti.

In the presence of Possible Omega Knowledge, Weak Fragility also avoids Fragility's endorsement of the JK thesis that if you know p, then you are justified in believing you know p. In order to be justified in believing you know p, it has to be possible for you that you omega know that you know that p. But Weak Fragility allows you to know something without being justified in believing that you omega know that you know it.

On the other hand, Weak Fragility shares with Fragility the other costs discussed in the previous section. It distinguishes higher iterations of knowledge; it doesn't explain why it is strange to ask about second-order knowledge; and it predicts the unconfident examinee can't know that she believes she doesn't know.

 $^{^{53}}$ This is a structural difference between Fragility and Reflective Luminosity. Weak Fragility is the weakest principle in the family of Fragility principles. By contrast, there is no weakest version of Reflective Luminosity. For every n, n-Reflective Luminosity is stronger than n+1-Reflective Luminosity.

Weak fragility also faces similar challenges from Margin for Error. It does not allow that Margin for Error is normally true, or justifiably believed. The problem is that Weak Fragility without Fragility leads to less disturbing but still inappropriate results, when combined with knowledge of Margin for Error. Suppose again you know the tree is not 100-1 feet tall, and that you know Margin for Error. It follows that it is possible that you know that: you know the tree is not 100 - 1 feet tall, and if you know the tree is not 100 - 1 feet tall, then it is not 100 feet tall. It follows that it is possible that you know that the tree is not 100 feet tall. This reasoning does not iterate symmetrically. I started with the premise that you know the tree is not 100-1 feet tall; but I have derived the structurally weaker consequence that it is possible that you know the tree is not 100 feet tall. From this and Weak Fragility, it follows that it is possible that it is possible that you know that: you know the tree is not 100 feet tall, and if the tree is not 100 feet tall it is not 100 + 1 feet tall. It follows that it is possible that it is possible that you know the tree is not 100 + 1feet tall. Iterating this reasoning, the result is that for any height, there is some n where it is n-possible that you know the tree is not that height. This is a surprising result. For example, extant omega skeptics do not embrace this result. According to standard omega skeptical models (outlined in Chapter 5), when the tree seems to be 100 feet tall, you know that it seems to be that height, and omega know you are justified in believing that it is between 100 - x and 100 + xfeet tall, for some value of x. So there is no n where it is n-possible that you know it is not 100 feet tall.

Still, perhaps this result could be tolerated. In that case, an open question is whether there is a viable theory that accepts Weak Fragility, denies Fragility or the agglomeration of justified belief, and accepts knowledge of Margin for Error. Since Fragility plays the crucial role of blocking Omega Skepticism, I do not pursue this question further in this book.

On the other hand, Weak Fragility does not retain all of the benefits of Fragility. n-Fragility does not explain the infelicity of dubious assertions that involve more than n iterations of knowledge. Models of Weak Fragility without Fragility will not predict that KK is normally true or justifiably believed. In addition, Weak Fragility does not fit very well with Possible Omega Knowledge; in that case, knowledge would not imply justification. Summarizing, Weak Fragility avoids the problematic prediction of Fragility that knowledge implies justified certainty. But it does so at the cost of losing the inference from knowledge to justified belief (in the presence of Possible Omega Knowledge), and also at the cost of rejecting the claim that KK is normally true or justifiably believed.

Fragility is also the strongest of another family of principles. A weaker version of the principle says:

(56) **Reflective Fragility.** If you know that you know p, then it is possible that you omega know p.

Recall Possible Permission, which says that you are justified in a behavior when for all you know it is permitted. In the presence of omega knowledge norms and Possible Permission, Fragility says that knowledge is sufficient for justified behavior (for possible omega knowledge). By contrast, Reflective Fragility would say that reflective knowledge rather than ordinary knowledge suffices for permissible (and hence justified) behavior. In some cases, merely knowing p does not seem to suffice for justifiedly relying on p. For example, perhaps scientists already know that vaping is safer than smoking. But they also know that further research about vaping is still appropriate, and they know that they should not be certain that vaping is the safer choice. In this way, mere knowledge may not suffice for justification. Reflective Fragility could explain why this is case: while scientists know that vaping is safer than smoking, they do not know that they know this. Similarly, the unconfident examinee is plausibly someone who knows when Queen Elizabeth died, even though they are not justified in being subjectively certain about whether they died. Fragility, Omega Infallibilism, and Possible Permission jointly rule out this possibility. By contrast, Reflective Fragility makes sense of it.

Generalizing, again say that you n-know p when you possess n iterations of knowledge regarding p. Then there is another family of Fragility principles:

(57) n-Reflective Fragility. If you n-know p, then for all you know, you omega know p.

Each weakening of n-Reflective Fragility raises the barriers to justified behavior. For example, 3-Reflective Fragility says that neither knowledge nor reflective knowledge suffice for justified behavior; but 3-reflective knowledge does suffice.

3.5 Conclusion

So far, I've discussed two strategies for explaining how omega knowledge could be abundant rather than scarce. In the next chapter, I introduce a third strategy, which I call Variable Margins.

Fragility and Reflective Luminosity have a wide range of consequences for epistemology. The two principles are compatible. I'll end the book with Chapter 6, which critically compares the costs and benefits of each principle, as well as Variable Margins, and their combination. Ultimately, I'll argue that the best theory of knowledge should accept at most one of the principles, but that each principle has significant benefits and costs compared to the other.

So far, Reflective Luminosity and Fragility have differed in a few important ways. First, they offer different approaches to which principles in epistemology are normally true or justifiably believed. As I'll show more carefully in Chapter 5, Reflective Luminosity allows for theories of knowledge where Margin for Error is justifiably believed and normally true. By contrast, Fragility allows for theories of knowledge where KK is justifiably believed and normally true. Second, the principles offer different approaches to justified belief. Unlike Reflective Luminosity, Fragility allows for Possible Omega Knowledge without severing the inference from knowledge to justified belief. In addition to these structural differences, Reflective Luminosity and Fragility also made a variety of different

predictions about cases involving iterated knowledge, about iterated factive mental states, about the connection between knowledge and justified or rational certainty, and about classic cases like the unconfident examinee. I'll explore these trade-offs more carefully in Chapter 6. Before doing so, I have two more goals. First, in Chapter 4 I'll introduce a third theory of knowledge, based around the idea that the margin for error for knowledge can vary in ways that allows for omega knowledge. Then, in Chapter 5, I'll develop detailed models of knowledge that vindicate some of the key claims I've made so far about the connection between knowledge, justification, and normality.

4 Variable Margins

Margin for Error says that for any height x, if you know that the tree is not x-1 feet tall, then the tree is not x feet tall. Generalizing from this specific claim, there are a family of Margin for Error principles which say that some margin m is such that for any height x, if you know that the tree is not x-m feet tall, then the tree is not x feet tall.

This chapter explores a weakening of this kind of principle:

(58) **Variable Margins.** For any height x, there is some margin m > 0, where if you know that the tree is not x - m feet tall, then the tree is not x feet tall.

Variable Margins is weaker than Margin for Error. Margin for Error says that there is some particular margin m which constraints your knowledge of the tree's height, no matter what height the tree is. Variable Margins instead allows that your knowledge of the tree's height may be constrained by a varying margin, depending on the actual height of the tree.⁵⁴

Suppose the tree appears 100 feet tall. Margin for Error says that no matter what height x the tree actually is, it is possible for you that the tree is x-1 feet tall. Variable Margins instead allows that when the tree is 100 feet tall, it is possible for you that the tree is 99 feet tall; but when the tree is 80.1 feet tall, it is not possible for you that the tree is 80 feet tall. In this way, Variable Margins allows that when the tree is 100 feet tall, your knowledge is constrained by a margin of 1 foot; but when the tree is 80.1 feet tall, your knowledge is constrained by a margin of less than 1 foot.

Variable Margins agrees that every iteration of knowledge requires an extra layer of reliability. But the defender of Variable Margins can allow for lots of omega knowledge, by saying that the marginal demand of reliability decreases with each iteration of knowledge, and approaches zero. Knowledge requires reliably true belief across a range of tree heights. Reflective knowledge requires reliably true belief across a wider range of heights. Knowing that you know that you know requires reliability across an even larger range of heights. But each widening of the range of heights is smaller than the one before. If these widenings approach a limit, then you get non-trivial omega knowledge of the tree's height. When the tree appears 100 feet tall and is so, you can omega know that the tree is greater than 80 feet tall. You know that the tree is at least 90 feet tall, because this belief is reliably true whenever the tree is at least 90 feet tall. You reflectively know that the tree is at least 85 feet tall, because this belief is reliably true whenever the tree is at least 85 feet tall. You know that you know that you know that the tree is at least 82.5 feet tall, because this belief is reliably true across an even wider range of tree heights. The strongest thing you omega know is that the tree is greater than 80 feet tall, because as the tree's

 $^{^{54} \}rm See$ Williamson 2000 §5.3 for discussion of conceptions of safety that correspond to Variable Margins. See Bonnay and Égré 2009 p. 190 for discussion of how Variable Margins can deflate challenges to KK.

height approaches 80 feet tall, all that is required for knowledge is that your belief is true at heights greater than 80 feet tall.

Following Williamson 2013a, say that you have 'cliff-edge knowledge' of the tree's height when you know that the tree is at least or at most x feet tall, and the tree is x feet tall. Variable Margins is equivalent to the thesis that you never have cliff-edge knowledge. In this way, any view that denies Variable Margins must accept that your powers of discrimination are incredibly powerful. (Since Margin for Error is stronger than Variable Margins, it also rules out cliff-edge knowledge.)

Recall the argument, reviewed in Chapter 1, that Margin for Error precludes omega knowledge. Again, 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 that 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. So you omega know the tree is not n feet tall. Iterating this reasoning, you omega know for any n that the tree is not n feet tall. But then something you omega know would be false.

Variable Margins avoids this argument. Again suppose that there is some n where you omega know that the tree is not n-1 feet tall. For example, perhaps you omega know that the tree is not 80 feet tall, and so n=81. Suppose also that you omega know Variable Margins. It follows that you know the tree is not 81-1 feet tall, and you know Variable Margins. But these two things do not imply that the tree is not 81 feet tall. For all you know, the margin for error when the tree is 81 feet tall is less than 1 foot. So it is consistent with what you know that the tree is 81 feet tall.

Variable Margins offers a substitute for Margin for Error. Your judgments about inexact knowledge can be fully systematized using Variable Margins, without accepting anything as strong as Margin for Error. Omega Skepticism can be avoided, because infinitely iterated knowledge does not require infinitely demanding reliability. The demands of increased reliability are ever increasing, but at an ever diminishing rate, and so approach a limit.

By offering an error theory of Margin for Error, Variable Margins is in some ways more satisfying than either Reflective Luminosity or Fragility. Reflective Luminosity required giving up Margin for Error. In its place, I proposed Known Margins, which said that Margin for Error held in the good case where one's belief forming processes were most reliable. But I had to concede that Margin for Error failed entirely in some bad cases. In particular, if the strongest thing that you reflectively know is that the tree is 80 feet tall, then Reflective Luminosity leads to the result that when the tree is 80 feet tall, you know that the tree is at least 80 feet tall; your knowledge is not inexact in that case. Similarly with Fragility: I suggested that Margin for Error may hold in sufficiently bad cases, but fails to govern what you know in the good case. In particular, if in the good case you know that the tree is at least 90 feet tall, then Fragility leads to the conclusion that if the tree was 90 feet tall, you would know that the tree was

at least 90 feet tall. In this way, both principles seemed to require cliff-edge knowledge, cases in which the tree is x feet tall, and you know that the tree is at least x feet tall.

On the other hand, Variable Margins is compatible with each of KK, Reflective Luminosity, and Fragility. To the extent that Variable Margins offers a better account of inexact knowledge than naive theories of knowledge which validate those latter principles, one proposal would be to combine some of KK, Reflective Luminosity, and Fragility, with Variable Margins.

One package of principles would be Variable Margins and KK. As discussed in Chapter 1, omega knowledge of Margin for Error is incompatible with KK, since together these principles predict that you know nothing about the tree's height. By contrast, omega knowledge of Variable Margins is compatible with KK. Yet in Chapter 1 I argued that the main appeal of KK is simply that it provides an antidote to Omega Skepticism. Since Variable Margins can do this without relying on KK, there is little reason to accept KK once you have already accepted Variable Margins.

Similarly, Variable Margins is also compatible with Reflective Luminosity and Fragility. But another view denies all of KK, Reflective Luminosity, and Fragility, but still accepts Variable Margins and denies Omega Skepticism. On this view, omega knowledge is abundant. But knowledge does not suffice for omega knowledge, since KK fails. Reflective knowledge does not suffice for omega knowledge, since Reflective Luminosity fails. Knowing in the good case does not suffice for omega knowledge, because Fragility fails. Still, this view says that we omega know many things, and that Margin for Error style reasoning is captured by Variable Margins. In Chapter 5, I develop a model of knowledge with these features.

Each of Reflective Luminosity and Fragility had surprising consequences that can be avoided if one accepts Variable Margins without those principles. For example, Reflective Luminosity implies that you can't have three iterations of knowledge without also having four. While Greco 2015 welcomes this result, others may be skeptical. For example, in Chapter 2 I considered the possibility that the reliability of a speaker can affect how many iterations of knowledge you can gain from their testimony: a very reliable testifier can give you three iterations of knowledge, while an extremely reliable testifier could give you four iterations of knowledge. Variable Margins is compatible with this graded conception of reliability; Reflective Luminosity precludes it.

Likewise, if you accept Variable Margins without Fragility, you can avoid some of the more controversial predictions from Fragility, while allowing that omega knowledge is abundant. One of Fragility's predictions was that when you know something, you are automatically justified in believing that you know it. Those who reject such immediate access to mental states will not be happy with this result. They could still embrace Variable Margins, and thereby accept the abundance of omega knowledge.

4.1 Costs of Variable Margins

On the other hand, Variable Margins has some significant costs of its own. These costs revolve around the fact that Variable Margins predicts that your omega knowledge about the tree's height is always an open interval.

(59) **Open Knowledge.** For any height x, there is some n where the strongest thing you omega know about the tree's height is that it is strictly greater than x - n feet tall.

Suppose the tree appears 100 feet tall. One model of Open Knowledge and Variable Margins says that when the tree is 100 feet tall, you omega know that the tree is greater than 80 feet tall. You know that the tree is at least 90 feet tall. When the tree is 90 feet tall, you know it is greater than 85 feet tall. As the tree's height approaches 80 feet tall, you continue to know that the tree is greater than 80 feet tall. No matter how close to 80 feet tall you get, Variable Margins remains true: there is always some margin m between the tree's height and 80, so that whenever the tree is x feet tall, it is possible for you that the tree is x-m feet tall. In Chapter 5, I develop a precise model of knowledge with this structure.

Without Open Knowledge, Variable Margins would not allow for omega knowledge of the tree's height. Suppose the tree appears 100 feet tall, is 100 feet tall, and that the strongest thing you omega know is that it is 80 feet tall or greater. Then if the tree were 80 feet tall, you would know that it was not less than 80 feet tall. So Variable Margins would fail when the tree was 80 feet tall. When the tree was 80 feet tall, you would have cliff-edge knowledge about the tree.

(KK says that omega knowledge is the same thing as knowledge. In the presence of Open Knowledge, KK thus implies that for any x, there is some n where the strongest thing you know about the tree's height is that it is greater than x-n. Not only your omega knowledge but also your knowledge would be open. By contrast, those who accept Variable Margins without KK could allow that some of your strongest knowledge is closed, even though your strongest omega knowledge is always open.)

Yet Open Knowledge comes with serious costs.

First cost: Variable Margins only applies to 'analogue' cases of knowledge; not to 'digital' knowledge.

Open Knowledge is unsatisfying as a general theory of inexact knowledge. The problem is that not all knowledge is 'analogue', concerning a continuous property. Variable Margins does not provide a satisfying account of omega knowledge in 'digital' cases. ⁵⁵ So far in this book, I have used tree heights as the main example. Tree height superficially appears to be continuous: for any two heights, there is another height between them. But not all knowledge is

⁵⁵Thanks to John Hawthorne for help here.

like this. Imagine that I glimpse a calendar which marks an exam, and see a date circled. My knowledge of the date is inexact. Suppose that the exam date appears to be February 27th, and is so. Plausibly, your knowledge of the date is governed by a margin for error. Seemingly, for any day x in February, if you know the exam is not on day x-1, then the exam is not on day x. Variable Margins cannot explain the appeal of this principle without embracing Margin for Error. The problem is that 1 day is the minimum difference between dates. If your knowledge of dates does not satisfy Margin for Error with a difference of 1 day, then your knowledge of dates does not satisfy Variable Margins. The only way to satisfy Variable Margins in this case would be for any day at least 1 away from the actual date to be epistemically possible.

Your knowledge of calendar dates cannot always be open while allowing abundant omega knowledge. Suppose that when the exam date appears to be February 27th and is so, the strongest thing you omega know that it is not February 13th. It follows that you omega know that the exam is February 14th or later. It follows from this that if the exam were on February 14th, you would know that it was not February 13th. If the exam were on February 14th, your knowledge would be closed rather than open. To summarize, Variable Margins is unsatisfying as a general account of inexact knowledge, because it only allows for omega knowledge in cases where your knowledge is open. But when you are learning about a digital rather than an analogue property, your knowledge can't be open.

The calendar case is one example of digital knowledge, where you have knowledge about a quantity that is not continuous. If you think that omega knowledge is abundant in the case of tree height, you should also think that it is abundant in the case of digital knowledge. After all, a glimpse at the calendar licenses assertions about when the exam is (that it is later than February 5th, for example). Moreover, digital knowledge is still inexact. When the exam appears from a brief glimpse to be on February 27th and is so, you do not know that the exam will be on February 27th. You could too easily have been wrong to count as knowing. For this reason, I am skeptical that Variable Margins provides a fully satisfying account of Margin for Error style reasoning. Variable Margins only provides an error theory in continuous cases, but does not extend to digital cases.

As another example, imagine that you are looking at a tree, and come to discover that the tree is made up of wooden blocks that are each 1 foot tall. This discovery shifts your knowledge from analogue to digital: you now know it's length is an integer. It would be bizarre if this shift led to a significant change in what you omega knew. In addition, our judgments related to margin for error style reasoning aren't affected by whether you make this discovery.

By contrast, both Reflective Luminosity and Fragility apply naturally to digital cases. Since Reflective Luminosity and Fragility are both defined in terms of knowledge, with no special reference to the subject matter of what is known, they also apply in the case of the calendar. For example, the defender of Reflective Luminosity can adopt the same error theory about Margin for Error in the case of the calendar as in the case of the tree. In the good case where the

exam appears to be on February 27th and is so, your knowledge is governed by a margin for error. This means that you are always justified in believing that your knowledge of the exam date is governed by a margin for error. Similarly, the defender of Fragility can continue to claim that your knowledge in the good case violates margin for error constraints, but that knowledge outside the good case is governed by a margin for error.

Second cost: Variable Margins comes arbitrarily close to admitting cliffedge knowledge.

Second, there is a question about whether Variable Margins and Open Knowledge capture the idea that one's knowledge is governed by a margin for error. These principles allow that one can rule out that the tree is 80 feet tall even as the tree becomes arbitrarily close to being 80 feet tall. Even when the tree is 80.000000001 feet tall, one could know that it is greater than 80 feet tall. Such powers of discrimination seem incredible, and call out for explanation. As Williamson 2000 puts it: "To be safe on the top of a cliff, a young child must be at least three feet from the edge; it is not enough to be some positive distance or other, no matter how small, from the edge" (p. 124). Moreover, the defender of omega knowledge will be forced to embrace such knowledge. Suppose that when the tree appears 100 feet tall and is so, the strongest thing you omega know is that it is greater than 80 feet tall. It follows that even when the tree is 80.000000001 feet tall, you continue to know that it is greater than 80 feet tall. After all, the tree being 80.000000001 feet tall is compatible with what you omega know when the tree is 100 feet tall. So if in that case you didn't know it was greater than 80 feet tall, then when the tree was 100 feet tall you wouldn't omega know after all that the tree is greater than 80 feet tall.

Each of Fragility and Reflective Luminosity deny Margin for Error. In the models I develop in the next chapter, each principle allows that there are some tree heights x where you know that the tree is at least x feet tall even when the tree is x feet tall. This would mean for example that when the tree is 80 feet tall, you could know that the tree is at least 80 feet tall. This is also an incredible claim, but doesn't seem significantly more incredible than that you can know the tree is greater than 80 feet tall when it is 80.0000000001 feet tall. Nonetheless, each of Reflective Luminosity and Fragility offer an explanation of the plausibility of Margin for Error. In the case of Reflective Luminosity, the explanation is that Margin for Error holds in the good case, and so you are always justified in believing it is true. In the case of Fragility, the explanation is that Margin for Error holds in bad cases, and so seems true because it reflects an important feature of our epistemic limitations.

Third cost: Variable Margins predict that the strongest proposition about a tree's height that you can assert or rely on is always an open interval.

I'll now turn to the predictions of Variable Margins for the theory of assertion

and other behavior. Again, Variable Margins offers a strategy for denying Omega Skepticism. On the resulting theory, you can omega know many things about the world. In this way, Variable Margins can make room for Omega Assertion, which says that you can permissibly assert p only if you omega know p. By making room for Omega Assertion, Variable Margins can explain the infelicity of dubious assertions, sentences of the form p and I don't know that I know ... that I know that p.

In order to account for omega knowledge, the defender of Variable Margins must embrace Open Knowledge: the strongest thing you omega know about the tree's height is an open interval. But I think that Open Knowledge leads to a strange prediction about assertion. In Chapter 1, I suggested that omega knowledge is necessary and sufficient for permissible assertion. The sufficiency of omega knowledge for permissible assertion is important for my theory of justification, since I think that you are justified in performing a behavior when for all you know it is permissible.

In addition to Omega Assertion, I also think that the practice of assertion is governed by a defeasible preference for asserting the strongest thing that you permissibly can. If you omega know p^+ and you omega know p^- , and p^+ is stronger than p^- , and both p^+ and p^- are relevant to the question at hand, then it is better to assert p^+ than p^- .

When combined with Open Knowledge, the preference for stronger assertions produces a strange result. When you are asked how tall the tree is, the best thing for you to assert about the tree's height will always be an open interval. Suppose again that the tree appears 100 feet tall, and is so. Imagine that your knowledge of the tree's height is open, so that the strongest thing you omega know is that the tree is greater than 80 feet tall. In that case, it will be best for you to assert that the tree is greater than 80 feet tall. This is not the only thing that you omega know. For example, you also know that the tree is at least 81 feet tall. But Open Knowledge implies that the very strongest thing that you omega know is always that the tree is strictly greater than some height.

This is a strange result. In many conversations about trees, it is perfectly natural to say that a tree is at least 80 feet tall, rather than that the tree is greater than 80 feet tall. This suggests that in many conversations, the strongest thing you omega know about the tree's height is that it belongs in a closed interval. 56

On the other hand, the defender of Variable Margins may reply that assertion is not governed by such strict standards. First of all, 'at least' and 'greater than' may not be scalar competitors, subject to pragmatic reasoning about

⁵⁶In January of 2022, engineers for D.C. Metro inspected the wheels of metro car trains. The Metro plan required flagging any car whose wheels were "more than 1/32 of an inch" further apart than normal. Unfortunately, the project was delayed once the engineers measured several wheels as being exactly 1/32 of an inch further apart than normal: the regulators could not agree about whether the cars should be flagged (https://www.washingtonpost.com/transportation/2022/01/08/metro-ntsb-railcar-investigation/). It would be cold comfort to the parties involved to inform the engineers that the strongest thing they omega know could not possibly be that the wheels were an exact distance apart, rather than being strictly greater than some distance apart.

which expression to use. Second of all, assertion tends to sloppy in various ways that make trouble for any simple theory like the one above. For example, most assertions about tree height will involve round numbers; but we shouldn't infer from this that the strongest thing you omega know must be a round number. (On the other hand, the preference for round numbers is plausibly explained as loose talk. But there doesn't seem to be a difference between 'at least' or 'greater' in terms of how loose or strict each expression is.) In this way, the exact predictions of Variable Margins for the theory of assertion are difficult to extract.

Analogous challenges await the application of Variable Margins to other forms of behavior. For example, consider Omega Infallibilism, the thesis that you are permitted to be certain of p iff you omega know p. In Chapter 1, I suggested that this thesis pairs naturally with the thesis that you should conditionalize on the strongest proposition that you omega know. In the presence of Open Knowledge, this leads to another surprising result. Suppose that your prior probability measure is regular, assigning positive probability to every possibility. Open Knowledge and Omega Infallibilism then lead to the result that your posterior credence measure of tree heights will only assign positive probability to an open interval of heights. For example, when the tree appears 100 feet tall and is so, you would assign positive probability to every height strictly greater than 80 feet tall, and yet you would assign no probability to the tree being 80 feet tall. Perhaps such probability measures can be rationally permissible. But it would be surprising if every single rationally permissible probability measure had this property. Similarly, another question is whether there is an important class of 'open' actions, that are permissible if and only if p, where p is an open interval describing the tree's height.

Fourth cost: Variable Margins does not predict that knowledge implies justification.

Another challenge for Variable Margins concerns the connection between knowledge and justification. In the previous chapter, I showed that Fragility could be combined with Possible Omega Knowledge to obtain the result that you are justified in believing everything you know. Without Fragility, Variable Margins and Possible Omega Knowledge would not have this implication. You could know something while knowing that you don't omega know it, and so without justifiably believing it. The defender of Variable Margins could instead opt for Possible Knowledge. But in that case, they will be forced to accept that you can justifiably believe dubious assertions. Here, the challenge for Variable Margins is similar to that for Reflective Luminosity.

4.2 Conclusion

I've now laid out the costs and benefits of Variable Margins as a theory of omega knowledge. The signature benefit of Variable Margins is that it provides

a systematic account of inexact knowledge, by offering a principled alternative to Margin for Error. Unlike Reflective Luminosity or Fragility, Variable Margins does not have to flat-footedly deny that Margin for Error holds across cases. Instead, it says that knowledge is systematically governed by a margin for error that varies as a function of the facts. In addition, I suggested that Variable Margins avoids some of the philosophical extravagance of Reflective Luminosity and Fragility. Unlike Reflective Luminosity, it does not posit any sharp cutoff in the demands that increasing iterations of knowledge impose on the reliability of one's belief forming faculties. Unlike Fragility, it does not imply that first-order knowledge immediately grants you a justification to act as if you omega know.

I also suggested that Variable Margins comes with a few major costs not shared by Reflective Luminosity or Fragility. First, the principle does not immediately extend to 'digital knowledge', like what you know when you catch a glimpse of a calendar. Relatedly, the principle may struggle to capture the spirit of margin for error requirements, since it allows (say) that when a tree is 80.000000001 feet tall, you can know that it is greater than 80 feet tall. Such a belief seems too unreliable to count as knowledge. Finally, the principle requires that the strongest thing you omega know is invariably an open interval. But this leads to the surprising prediction that there would be some kind of systematic preference to assert open rather than closed claims.

So far in this book, I've informally explored three potential theories of omega knowledge. Each theory explains how omega knowledge is possible; but each theory comes with rival costs and benefits. In the next chapter, I explore each theory in mathematical detail. I develop formally precise models of knowledge that make clear predictions about how what you know in a particular case. These models illustrate in detail how Reflective Luminosity, Fragility, and Variable Margins differ regarding their underlying conceptions of knowledge.

5 Modeling Knowledge

In this chapter, I develop theories of knowledge that validate Reflective Luminosity, Fragility, and Variable Margins. In doing so, I vindicate my earlier claims that these principles can defeat Omega Skepticism. I begin by introducing standard tools from epistemic logic, in order to illustrates the logical connections between the various principles discussed so far. After that, I develop three theories of knowledge: one that validates Reflective Luminosity, one that validates Fragility, and one that validates Variable Margins. To do so, I focus on the case of perceptual knowledge, drawing on existing work from Williamson 2013a that defines what you perceptually know in terms of the distance between reality and appearance. I use these models to explore in detail how much you omega know on the basis of perception. In doing so, I vindicate my earlier claims that Reflective Luminosity, Fragility, and Variable Margins avoid Omega Skepticism. In addition, I explore the extent to which each theory can capture the ideas behind Margin for Error.

This chapter uses lots of mathematical formalism. For readers who aren't interested in this, I recommend skipping to the next chapter, which summarizes the main arguments of the book.

5.1 Epistemic Logic

I interpret the knowledge of a single agent as a modal necessity operator K, and I interpret epistemic possibility as the dual of K, called M. p is epistemically possible for you iff you don't know not p iff for all you know, p.

Epistemic logic defines knowledge in terms of an accessibility relation E between possible worlds. You know p (Kp) at world w just in case p is true at every world v E-accessible from w. p is true for all you know (Mp) just in case there is some E-accessible world at which p is true. Ew is the set of worlds that are E-accessible from w.

You know you know p (KKp) at w just in case p is true at every world accessible from a world accessible from w. In other words, KKp is true at w just in case p is true throughout E^2w , where wE^2u iff for some v, wEv and vEu. Iterated knowledge universally quantifies over the accessibility relation E^2 derived from epistemic accessibility.

Say that v is ancestrally accessible from w iff there is a path of accessibility from w to v that traverses some number of worlds. You omega know p iff p is true at every world v ancestrally accessible from w.

Properties of knowledge correspond to properties of epistemic accessibility. Consider the principle that knowledge is factive, so that anything known is true. Factivity corresponds to the reflexivity of epistemic accessibility: every world \boldsymbol{w} is accessible from itself.

KK corresponds to the transitivity of accessibility. KK says that Kp implies KKp. This means that whenever p is true throughout Ew, A is also true throughout E^2w . This is equivalent to the requirement that $E^2w \subseteq Ew$, so that u is accessible from w whenever u is accessible from v and v is accessible from w.

Omega Skepticism says something about the ancestral of accessibility. It says that for any ordinary claim p, there is some world in the ancestral of accessibility where p is false. The idea is that even when you know p, there is some chain of worlds (potentially very long) connected by epistemic accessibility which ends in a world where p is false.

Reflective Luminosity corresponds to a weakening of transitivity. It says that whenever you can get from w to z by three steps (so that wEv, vEu, and uEz), you can get from w to z by two steps (so that wEv', and v'Ez). 57

Again say that v is ancestrally accessible from w iff there is a path of accessibility from w to v that traverses some number of worlds. You omega know p iff p is true at every world v ancestrally accessible from w. Reflective Luminosity says that whenever v is ancestrally accessible from w, it can be reached from w in only two steps. If you know you know p, then you omega know p.

n-Reflective Luminosity says that if you n-know p, then you (n+1)-know p. n-Reflective Luminosity corresponds to a more general weakening of transitivity. It says that whenever you can get from w to z through n worlds v_1 through v_n (so that wEv_1, v_1Ev_2, \ldots , and v_nEz), you can get from w to z through n-1 world v_1' through v_{n-1}' (so that $wEv_1', v_1'Ev_2', \ldots$, and $v_n'Ez$). Any journey of n steps can be reduced to a journey of n-1 steps.

Again say that v is ancestrally accessible from w iff there is a path of accessibility from w to v that traverses some number of worlds. You omega know p iff p is true at every world v ancestrally accessible from w. n-Reflective Luminosity says that whenever v is ancestrally accessible from w, it can be reached from w in only n steps.

Fragility corresponds to a different kind of weakening of transitivity. Recall that E^* , the ancestral of E, relates w to v just in case v can be reached from w by some series of worlds. Fragility says that every world w can see some world v where any world ancestrally accessible from v is accessible from w. Say that an

 $^{^{57}}$ Suppose KKKp fails at w. Then p is false at some z with wEv and vEz. But there is some v' where wEv' and v'Ez. So KKp fails at w also. Conversely, suppose that wEv, vEu and uEz, but there is no v' where wEv' and v'Ez. Now let p be true at exactly those u' where wEv'' and v''Eu', for some v''. KKp is true at w. But KKKp is false at w. After all, p is false at z, and yet wEv, vEu, and uEz; so Kp fails at u, KKp fails at v, and thus KKKp fails at w.

 $^{^{58}}$ Reflective Luminosity is weaker than another principle, which says that you know that: if you know p, then you know that you know p. Given epistemic closure, this principle implies Reflective Luminosity, but is not implied by it. This principle's says that whenever there is a path from w to z via v and u (so that $wEv,\,vEu,$ and uEz), there is also a path from w to z via v alone (so that wEv and vEz). This principle is stronger than Reflective Luminosity because it requires the skip from 3 to 2 steps to happen in a particular way. If there are 3 steps $w\to v\to u\to z$, there must not only be some 2 step pattern $w\to v'\to z$; in addition, the pattern must be $w\to v\to z$.

⁵⁹Suppose $K^{n+1}p$ fails at w. Then p is false at some z with wEv_1 and ... and v_nEz . But there are some v_1' through v_{n-1}' where wEv_1' and ... and $v_{n-1}'Ez$. So K^np fails at w also. Conversely, suppose that $wEv_1 \ldots v_nEz$, but there is no $v_1' \ldots v_{n-1}'$ where $wEv_1' \ldots v_{n-1}'Ez$. Now let p be true at exactly those u' where $wEv_1'' \ldots v_{n-1}''Eu'$, for some $v_1'' \ldots v_{n-1}'' K^np$ is true at w. But $KKK_{n+1}p$ is false at w. After all, p is false at z, and yet $wEv_1 \ldots v_nz$; so $K^{n+1}p$ fails at w.

accessibility relation E is jump transitive if and only if $\forall w \exists v \in Ew : E^*v \subseteq Ew$. Fragility is valid if and only if epistemic accessibility is jump transitive.⁶⁰

Fragility corresponds to a coherent constraint on epistemic accessibility. This constraint is compatible with the reflexivity of accessibility, so that Fragility is compatible with the factivity of knowledge. Fragility also has some small consequences for any modal operator that satisfies it: for example, it implies that accessibility is serial, so that every world sees some other world. This in turn corresponds to the requirement that when you know p, p is possible for you.

The characterization of Fragility in terms of jump transitivity also clarifies the relationship between KK and Fragility. First, KK implies Fragility. For suppose E is transitive and reflexive. Then every world w trivially sees a world v (in particular, itself) where $E^*v \subseteq Ew$. By contrast, Fragility does not imply KK. Epistemic accessibility can be jump transitive without being transitive. Jump transitivity concerns the relationship between Ew and E^*v for some v or other. Transitivity concerns the relationship between Ew and E^*w .⁶¹

n-Fragility says that if you know p, then for all you know, you n-know p. Let the n-ancestral of E (E^n) be an accessibility relation that relates w and v just in case v can be reached from w through n applications of R. n-Fragility then corresponds to a generalization of jump transitivity. It says that every world w can see some world v where any world accessible from v by the v-ancestral of v is accessible from v by v-ancestral only if v-ancestral only in v-ancestral only if

5.2 Appearance/Reality Models

So far, I have characterized the major principles in the book in terms of epistemic logic. But I have not provided a substantive interpretation of the epistemic accessibility relation. The rest of this chapter offers such an interpretation. I suggest that knowledge is about normality, not counterfactual similarity.

Before getting into the details, I want to clarify the methodology in this chapter. I'm going to develop three models of knowledge in terms of normality. In each of these models, normality has a different structure. But in each case, I will not offer a general argument that normality has this structure. Instead, I'll show that if normality did have this structure, then we would arrive at an

 $^{^{60}\}mathrm{Let}\ M^*$ denote the dual of omega knowledge. Suppose E is jump transitive and $w\models KM^*A$ for every n. Then $\forall v\in Ew: E^*v\cap A\neq\emptyset$. By jump transitivity, $\exists v^*\in Ew: E^*v^*\subseteq Ew.$ So $Ew\cap A\neq\emptyset$. So $w\models MA$. Conversely, suppose that E is not jump transitive. Then $\forall v\in Ew:\exists z\in E^*v:z\not\in Ew.$ Let $A=\{w\mid \exists v\in Ew:z\in E^*v\ \&\ z\not\in Ew\}.\ w\models KM^*A,$ since $\forall v\in Ew:v\models M^*A.$ But $w\not\models MA,$ since $\neg\exists v\in Ew:v\in A.$

⁶¹For other work on weakenings of positive and negative introspection in epistemic logic, see San 2019.

⁶²Suppose E is jump transitiveⁿ and $w \models KM^nA$. Then $\forall v \in Ew : E^nv \cap A \neq \emptyset$. By jump transitivityⁿ, $\exists v^* \in Ew : E^nv^* \subseteq Ew$. So $Ew \cap A \neq \emptyset$. So $w \models MA$. Conversely, suppose that E is not jump transitiveⁿ. Then $\forall v \in Ew : \exists z \in E^nv : z \notin Ew$. Let $A = \{w \mid \exists v \in Ew : z \in E^nv \& z \notin Ew\}$. $w \models KM^nA$, since $\forall v \in Ew : v \models M^nA$. But $w \not\models MA$, since $\neg \exists v \in Ew : v \in A$.

interesting theory of knowledge and justification. In addition, I'll show how a particular structure leads to striking predictions about what is normally true.

In developing these models, my goal is to redeem various promissory notes from Chapters 2-4 about the benefits and costs of Reflective Luminosity, Fragility, and Variable Margins. Without a full model of knowledge, it is difficult to do this. In addition, several of the benefits of Reflective Luminosity and Fragility involved each principle being compatible with further principles governing knowledge, justification, and normality. The theories I'll develop in this chapter explain how all these principles can fit together. In addition, my hope is that even if I can't motivate various structural constraints on normality in the abstract, the work in this chapter will offer some support for the structural features of normality precisely by showing how these structural features leads to a striking array of interesting consequences for the theory of knowledge, justification, and normality.

Above all, I will focus on a trade-off between Reflective Luminosity and Fragility. I'll show how each principle fits with a different conception of what is normally true. Reflective Luminosity will fit with the idea that Margin for Error is normally true. Fragility will fit with the idea that KK is normally true. While my models of normality at first may seem to appeal to ad hoc structural constraints on normality, ultimately the models are just ways of fleshing out this basic difference in the theory of what is normally true. In addition, I'll show how these differing theses about what is normally true also lead to differing predictions about the situations in which you possess cliff-edge knowledge; that is, knowledge that is so maximally discerning that it allows you to know that a tree is at least or at most as tall as it actually is. I'll show that Reflective Luminosity fits smoothly with a model of knowledge where you only possess cliff-edge knowledge when your perceptual faculties are sufficiently unreliable; by contrast, Fragility fits with a model of knowledge where you only possess cliff-edge knowledge when your perceptual faculties are sufficiently reliable.

To make precise predictions, I focus on the case of perceptual knowledge, and in particular on inexact knowledge of tree height. You are looking at the tree, and it appears to be 100 feet tall. What you know about the tree's height depends on both how tall the tree appears to be and how tall the tree actually is. Williamson 2013a offers a simple model of this case that I'll expand on in this chapter. This 'appearance / reality' model measures how much you know about a quantity like tree height or temperature. To do so, I represent each possible world as a pair of an apparent height a and a real height r. Various principles about knowledge then correspond to constraints on what you know about the tree's appearance and height, as a function of a and r.

As before, what you know about the tree's height is represented by an accessibility relation, E. But now E relates pairs (a, r) of the tree's apparent and real heights. E(a, r) is the set of worlds E-accessible from (a, r). This is the strongest proposition known at (a, r).

I'll assume for simplicity that the tree's apparent height is luminous.⁶³

⁶³But see Rosenkranz and Dutant 2020 for a generalization of these models without this

(60) **Appearance Luminosity** For any height x, if the tree appears to be x feet tall, then you know that the tree appears to be x feet tall.

Appearance Luminosity corresponds to the constraint that if (a,r)E(a',r'), then a=a'. This principle reflects a conception of normality. When I assess the normality of your belief forming process, I hold fixed how the world seems phenomenally to you. I wonder what you know, given these seemings. To do so, I consider how normal various real heights would be, supposing that the tree appears to be a certain height. (Relatedly, in Chapter 1 I considered an interpretation of normality in terms of the accuracy of your evidence: the more accurate your evidence is, the more normal the situation. In that setting, Appearance Luminosity is related to the choice to identify the evidence with how things appear, and only compare two situations in which you have the same evidence, when considering whether your evidence is more or less accurate.)

I assume that the facts of knowledge depend only on the distance between appearance and reality, |a-r|. Williamson 2013a proposes that you know that the tree's height is not r' iff the distance between the tree's apparent height a and r' is significantly larger than the distance between a and the tree's actual height r. Let the margin for error constant m model significant differences in distance. This proposal then says:

(61)
$$(a,r)E(a',r')$$
 iff $a = a'$ and $r' \in [a - (|a-r|+m), a + (|a-r|+m)]$

For example, consider Figure 1. Figure 1 represents what you know about the tree when it appears to be 100 feet tall, and when the margin for error is 10 feet. (I'll use this as a running example throughout the chapter.)




Figure 1: a = 100, m = 10

assumption.

100 is the good case, where conditions are most normal, reality matches appearance, and the tree is actually 100 feet tall. In the good case, you know that the tree's height is between 90 and 110 feet. 80 is a bad case, where reality departs from appearance considerably, and the tree is 20 feet less than it appears to be. In the bad case, you know that the tree's height is between 70 and 130 feet. Whenever the tree is in that range, the distance between the tree's apparent and real height is no greater than 30. This value of 30 is found by adding the margin for error of 10 to the actual distance between appearance and reality (|100 - 80|).

KK fails. In the good case where the tree is 100 feet tall, you know that it is between 90 and 110 feet tall. But in the worse case where the tree is 90 feet tall, you know only that the tree is between 80 and 120 feet tall. When the tree is 100 feet tall you know, but do not know that you know, that the tree is between 90 and 110 feet tall.

Not only KK but also Reflective Luminosity and Fragility fail. When the tree is 100 feet tall, you know that you know the tree is between 80 and 120 feet tall. But you do not know that you know that you know this. After all, for all you know that you know, the tree is 80 feet tall; but if the tree is 80 feet tall, then you do not know that it is between 80 and 120 feet tall. Similarly with Fragility. When the tree is 100 feet tall, you know that it is between 90 and 110 feet tall. But you know that you do not know that you know this. After all, at every height between 90 and 110 feet, you fail to know that you know the tree is between 90 and 110 feet tall.

This model embraces Omega Skepticism. It says that omega knowledge is scarce. When the tree is 100 feet tall, you omega know nothing about its height. For this reason, the model predicts that it is impossible to have common knowledge about the tree's height on the basis of how the tree appears. The model likewise predicts that dubious assertions are knowable. When the tree is 100 feet tall, you know the conjunction: the tree is 90-110 feet tall, and you don't know that you know this. If assertion is governed by a knowledge norm, then this sentence should be assertible. Similarly, you know that you know this.

This model offers an interpretation of similarity and normality. Similarity depends on the tree's real height. The worlds (a, r') similar to (a, r) are those where the height r' is sufficiently close to the actual real height r. By contrast, normality depends on the distance between the tree's apparent and real height. The worlds (a, r') almost as normal as (a, r) are those where the distance between a and r' is not very different than the distance between a and r.

Consider (100, 80), where the tree seems to be 100 feet tall but is actually 80 feet tall. When the tree is 80 feet tall, the following counterfactual has a true reading: If the tree were 79 or 121 feet tall, it would be 79 feet tall. This counterfactual has a true reading because 79 is more similar to the tree's actual height than 121. In this way, similarity tracks real height. But at (100, 80), both 79 and 121 are epistemically possible heights of the tree. This is because these two heights are equally normal. Knowledge tracks normality rather than

similarity.

This model makes a few predictions that will be relevant later. First, the model is equivalent to the conjunction of Appearance Luminosity with the constraint that you know as much when the tree is r as you do when the tree is r' iff r is as close to a as r' is.

(62) **Distance.**
$$|a-r| < |a-r'|$$
 iff $E(a,r) \subseteq E(a,r')$

Distance implies that you always know less as the distance between reality and appearance increases.

Distance entails Margin for Error. In the present model, Margin for Error says:

(63) There exists some margin m > 0 where for all $r' \in [r - m, r + m]$, (a, r)E(a, r').

Suppose that there is a first height where Margin for Error fails. For example, imagine again that the tree appears 100 feet tall. Suppose that Margin for Error fails when the tree is 80 feet tall, but does not fail when the tree is 81 feet tall. In that case, when the tree is 80 feet tall you know that the tree is 80 feet tall. But when the tree is 81 feet tall, for all you know the tree is 80 feet tall. In that case, Distance fails. Your perceptual appearances are strictly less accurate when the tree is 80 as when the tree is 81 feet tall. But you know the same thing in both cases: that the tree is at least 80 feet tall. In my own models, I will reject Distance, but retain its left to right direction.

The final principle that will be relevant later is that your knowledge is centered on the appearances:

(64) **Appearance Centering.** For some
$$x > 0$$
, $(a, r)E(a, r')$ iff $r' \in [a - x, a + x]$

For example, when the tree appears 100 feet tall and is 100 feet tall, you know that the tree's real height is within 10 of 100. When the tree appears 100 feet tall and is 90 feet tall, you know that it is within 20 feet of 100. Either way, you always know that the tree's height is within some distance of 100. I accept Appearance Centering, because it connects knowledge to evidential support. Your evidence is that the tree appears 100 feet tall. This evidence supports various hypotheses about the tree's height. It most strongly supports the hypothesis that the tree is 100 feet tall; but it also supports the hypothesis that it is 90, or 110 feet tall. It is evidence against the hypothesis that the tree is 10 or 190 feet tall.

I think that your evidence in this case is symmetric: it supports taller heights than 100 to the same degree that it supports shorter heights. Then Appearance Centering says that when two heights are supported to the same degree by your evidence, either both or neither of the heights are epistemically possible. ⁶⁴

⁶⁴See Goldstein and Hawthorne forthcomingb for a precisification of this idea in terms of evidential probability. It's worth flagging that there are reasons to worry about the principle

Goodman 2013 observes that Williamson's model above is the smallest relation satisfying Appearance Luminosity, Appearance Centering, and Margin for Error. This means that the model predicts that you know as much as possible subject to these constraints. In the next section, I'll develop a model of Reflective Luminosity. This model predicts that you know as much as possible subject to Appearance Luminosity, Appearance Centering, Known Margins, and Reflective Luminosity.

Building on Stalnaker 2009, Cohen and Comesaña 2013 develops a model that validates KK. Epistemic accessibility is again defined relative to the margin for error m; but the margin has different effects in three cases. Whenever the tree's height is within m of a, you know that it is within m of a. When the tree's height r is less than a-m, you know that it is between r and a+m feet tall. When the tree's height r is more than a+m, you know that it is between r and a-m.

(65)
$$(a,r)E(a',r') \text{ if and only if } a=a' \text{ and } \begin{cases} r' \in [r,a+m] & \text{if } r < a-m \\ r' \in [a-m,r] & \text{if } r > a+m \\ r' \in [a-m,a+m] & \text{otherwise} \end{cases}$$

The situation is depicted in Figure 2. KK is valid. When the tree's height is within [a-m, a+m], you know that the tree's height is within [a-m, a+m]. When the tree is less than a-m, you know that it is no shorter than it is.

Margin for Error fails dramatically. Following Williamson 2013a, say that you have 'cliff-edge knowledge' of the tree's height when you know that the tree is at least or at most x feet tall, and the tree is x feet tall. In this model, you have cliff-edge knowledge of the tree's height whenever the tree is at least m distance from a. The models I develop later in this chapter also accept cliff-edge knowledge, but do not agree that you have cliff-edge knowledge whenever the tree is at least m distance from a.

that any two situations equally supported by your evidence are either both possible or both impossible. Here's one example from Hawthorne 2003 (p. 71; see also Goodman and Salow forthcoming for similar cases) where this idea leads to skepticism: 'Suppose there are two newspapers, The Times and The Guardian, which I trust equally well for the purposes of obtaining soccer information. With good reason: both are extremely reliable in their reporting of soccer results. I look in The Times and find a Manchester United victory reported. I trust the report. The report is in fact correct. Under such circumstances, people are inclined to say I know both that The Times said that Manchester United won and also that Manchester United won...Suppose, in fact, that, unbeknownst to me, The Guardian [made] a mistake.' In this case, I know that the Times reported correctly; but I don't know whether The Guardian did. Yet either newspaper making a mistake is equally likely on my evidence. In cases like this, my principle leads to the conclusion that you don't know on the basis of The Times that Manchester United won. This conclusion starts to seem inappropriately skeptical as the number of newspapers increases. Imagine that there are ten thousand newspapers reporting the results, and only one makes a mistake. Still, reading the correct result in a newspaper won't give me knowledge, since my evidence equally supports that any given newspaper could be mistaken.

⁶⁵Any smaller relation has E(a,r) exclude some (a,r') with $r' \in [a-(|a-r|+m),a+(|a-r|+m)]$. Suppose for simplicity that $r' \le a-(|a-r|+m)$. By Appearance Centering, I can suppose that r < a, since otherwise I could have $r' \le a+(|a-r|+m)$. But then Margin for Error fails, because you know that the tree is not r' feet tall, and yet r' is within m of r.




Figure 2: a = 100, m = 10

The model also rejects Appearance Centering: when the tree is 80 feet tall, you know that the tree is either 20 feet less than a or 10 feet greater than a. By contrast, my own models accept Appearance Centering.

Finally, the model denies Distance. When you are within m of a, you know the same thing regardless of how far from a you are. One of my goals in this book is to deny Margin for Error while preserving as much of its spirit as possible. For this reason, I also deny Distance. However, I will continue to accept the left to right direction of Distance. Sometimes when the distance between reality and appearance increases, your knowledge does not decrease. But you never know more when the appearances become less accurate.

(66) Weak Distance. If
$$|a-r| \le |a-r'|$$
, then $E(a,r) \subseteq E(a,r')$

The model of KK above rejects Distance but accepts Weak Distance. My own models of Reflective Luminosity, Fragility, and Variable Margins will make the same prediction.

Summarizing, existing models of knowledge either accept or reject KK. But the existing models that reject KK are omega skeptical. They also reject Reflective Luminosity and Fragility. And the existing models that accept KK also contain flagrant violations of Margin for Error. My goal is to reject both KK and Omega Skepticism, while accepting principles like Reflective Luminosity, Fragility, and Variable Margins, and while giving a plausible error theory about why Margin for Error seems valid. So I need to develop new models of knowledge.

Before I turn to my own models, I also summarize existing appearance/reality models of justified belief. Appearance/reality models offer a simple way of validating Possible Knowledge and Possible Omega Knowledge. Following Williamson

2013a, I introduce an accessibility relation B that models justification. wBv iff every claim true at v is consistent with what you justifiedly believe at w.

In these models, Possible Knowledge says that you are justified in believing whatever you would know in the good case. Possible Omega Knowledge says that you are justified in believing whatever you would omega know in the good case. For example, consider the omega skeptical model from earlier. There, Possible Knowledge says you are justified in believing p when p is true at any world that is not significantly less normal than the most normal worlds. The most normal worlds are those that are no less normal than any other world. Then wBv iff v is not significantly less normal than any most normal world. ⁶⁶ Alternatively stated, you are justified in believing that things are no more than a little abnormal (Leplin 2009, Smith 2016, Goodman and Salow 2018). In these particular models, normality is understood in terms of appearance and reality. Knowledge is nested around the good case where reality matches appearance. No matter how much you know, it is always possible for you that you are in the best case. You know at least as much in the best case as you do in any other case. Even in the bad case, Possible Knowledge says that you are justified in believing exactly what you know in the best case. You are justified in believing all and only what you know when reality matches appearance.

In the best case, you know that the tree's height is within m of the tree's apparent height. Therefore, Possible Knowledge says that what you justifiably believe in any case is that the tree's height is within m of the tree's apparent height.

(67)
$$(a,r)B(a',r') \text{ iff } a = a' \text{ and } r' \in [a-m, a+m]$$

On this proposal (a, r)B(a', r') iff (a, a)E(a', r'). That is: you are justified in believing whatever you know in the best case.

As an illustration of this model of justification, consider the new evil demon problem. You are justified in believing the same thing whether you are in the good case or the bad case. Return to Figure 1. Whether the tree is 100, 90, or 80 feet tall, you are justified in believing the same thing: that the tree is between 90 and 110 feet tall. In all cases, this is the strongest thing that you could possibly know. No matter how radically deceived you may be, you retain your justification for believing that things are not too abnormal.

The omega skeptical model above would produce skepticism about justified belief when combined with Possible Omega Knowledge. On that view, you would be justified in believing whatever you omega know in the good case. Where E^* is the ancestral of epistemic accessibility, Possible Omega Knowledge says that (a,r)B(a',r') iff $(a,a)E^*(a',r')$. But since you omega know nothing in the good case, you would be justified in believing nothing. By contrast, the models I develop below make room for Possible Omega Knowledge as a theory of justification, because they are not omega skeptical.

 $[\]overline{\ }^{66}$ This assumes that all worlds are ordered for normality. When this assumption fails, the notion of being most normal would be relativized to w. See Goodman and Salow 2018 for discussion.

5.3 Reflective Luminosity

I now offer a model of knowledge on which Reflective Luminosity and Known Margins are valid, but Margin for Error, KK, and Omega Skepticism are invalid.

I begin with Known Margins. Known Margins says that when appearance matches reality, you know Margin for Error. The two defining features of my model are Known Margins and Reflective Luminosity. In particular, the first feature of my model is that you know as much as possible consistent with Known Margins, Appearance Luminosity, and Appearance Centering. This entails that whenever the tree's actual height is within m of its apparent height, the model agrees with Williamson's:

(68) If
$$0 \le |a-r| \le m$$
, then $(a,r)E(a,r')$ iff $a = a'$ and $r \in [a-(|a-r|+m), a+(|a-r|+m)]$.

Reflective Luminosity says that anything that is reflectively known is omega known. I say that you know as much as possible consistent with (68) and Reflective Luminosity. This imposes the following constraint:

(69) If
$$m < |a-r| \le 2m$$
, then $(a, r)E(a, r')$ iff $a = a'$ and $r \in [a-2m, a+2m]$.

To see why (69) follows from these assumptions, note that in order to validate Reflective Luminosity, I must guarantee that whenever v is accessible from w via some series of steps, v is accessible from w in two steps. This way, whenever you know you know p, you omega know p. (68) ensures that When reality matches appearance, the strongest thing you know you know is that r is within 2m of a. So the furthest accessibility can reach in two steps is 2m. To validate Reflective Luminosity, I therefore claim that whenever you are within 2m of a, you know you are within 2m of a. For this reason, Margin for Error can fail when the tree's height is further than m and no greater than 2m from a. This is what (69) says.

In earlier chapters, I introduced the idea that principles like Margin for Error or KK could be normally true, even if they aren't universally valid. One interpretation of a 'normally true' principle is one that is true in any situation that is almost as normal as the most normal one. In this model, Margin for Error is true in any situation where the distance between appearance and reality is no more than m. In this sense, it is normally true.

To reach a full model of knowledge, I must extend the constraints above to make predictions about what you know when the tree's height is beyond 2m of a. Here, Known Margins no longer imposes any particular constraints on what you know. For simplicity, I assume that your knowledge outside of 2m from a is structurally parallel to your knowledge within 2m of a. In particular, I suppose that there are regions in which your knowledge obeys a margin for error; and then regions in which Margin for Error fails. The first assumption limits cliffedge knowledge as much as possible; the second assumption preserves Reflective Luminosity. Summarizing, I propose the following theory of knowledge:

(70)
$$(a,r)E(a',r')$$
 iff $a = a'$ and:

```
a. if 0 \le |a-r| \le m, then r' \in [a-(|a-r|+m), a+(|a-r|+m)]
b. if m < |a-r| \le 2m, then r' \in [a-2m, a+2m]
c. if 2m < |a-r| \le 3m, then r' \in [a-(|a-r|+m), [a+(|a-r|+m)]
d. if 3m < |a-r| \le 4m, then r' \in [a-4m, a+4m]
```

Figure 3 provides one example, with m=10. The tree appears to be 100 feet tall. Whenever the tree's height is between 90 and 110 feet tall (within m feet of a), your knowledge is governed by Margin for Error. You know that the tree's height is at most 10 feet more than |100-r| away from 100. In this region, you know less as the tree's height departs from 100. When the tree is 90 feet tall, you know that the tree is at least 80 feet tall. This means that when the tree is 100 feet tall, you know that you know that the tree is 80 feet tall.

To validate Reflective Luminosity, I guarantee that when the tree is 100 feet tall, you omega know that the tree is 80 feet tall. By validating Reflective Luminosity, I depart from Williamson 2013a's model. That model accepts Distance, which says that you always know less when the distance between appearance and reality increases, so that $E(a,r)\subset E(a,r')$ if |a-r|<|a-r'|. My model denies this assumption. Sometimes, when the distance between appearance and reality increases, there is no change to what you know. In particular, I claim that Distance fails when the tree is less than 90 feet tall and at least 80 feet tall. Within this region, you know that the tree is at least 80 feet tall. If you knew any less than this, then Reflective Luminosity would fail, because when the tree is 100 feet tall, you would not omega know that the tree is 80 feet tall. (On the other hand, I retain Weak Distance, the left to right direction of Distance, which says that if the tree's height is at least as close to appearance at r as at r', then you know at least as much at r as at r'.)

In this sense, I depart from Williamson's models to the minimal degree necessary to validate Reflective Luminosity, by granting the agent only as much additional knowledge as is necessary to collapse reflective and omega knowledge, while agreeing with Williamson about Margin for Error in as many worlds as possible (in particular, when the tree is between 90 and 110 feet tall).

KK fails. When the tree is 100 feet tall, you know it is between 90 and 110 feet tall. If it is 90 feet tall, you know only that it is between 80 and 120 feet tall. So when the tree is 100 feet tall, you don't know that you know it is between 90 and 110 feet tall.

Reflective Luminosity is valid. When the tree is 100 feet tall, you know you know the tree is between 80 and 120 feet tall. You also know that you know that you know this. This is because when the tree is 80 feet tall, you also know that the tree is between 80 and 120 feet tall.

Omega Skepticism is false. When the tree is 100 feet tall, you omega know that it is between 80 and 120 feet tall. Similarly, when the tree is not between 80 and 120 feet tall, but is between 60 and 140 feet tall, you reflectively know that the tree is between 60 and 140 feet tall. You also omega know this fact.

Margin for Error is invalid. Margin for Error said that for some m > 0, when you know the tree is not x - m feet tall, the tree is not x feet tall. This fails




Figure 3: a = 100, m = 10

when the tree is 80 feet tall. When the tree is 80 feet tall, you know it is not 80 - m feet tall for every m > 0. (Similarly when the tree is 60 feet tall.)

Again following Williamson 2013a, say that you have 'cliff-edge knowledge' of the tree's height when you know that the tree is at least or at most x feet tall, and the tree is x feet tall. In this model, you have cliff-edge knowledge when the tree is 80 and 60 feet tall.⁶⁷

Known Margins is valid. Known Margins says that when reality matches appearance, you know that Margin for Error is valid. This holds. When the tree is 100 feet tall, you know it is between 90 and 110 feet tall. In the worst case, the tree would be 90 feet tall. When the tree is 90 feet tall, for all you know the tree is 90 - m = 80 feet tall. So Margin for Error is true whenever the tree is between 90 and 110 feet tall. So when the tree is 100 feet tall, you know Margin for Error is true.

Given Possible Knowledge, Known Margins implies Justified Margins. You are justified in believing what you know when reality matches appearance. So no matter what the tree's height is, you are justified in believing Margin for Error. Unfortunately, however, in this model Possible Knowledge leads to Dubious Justification. You are justified in believing that: the tree is between 90 and 110 feet tall even though you don't know that you know the tree is between 90 and 110 feet tall.

In this model, Possible Omega Knowledge avoids the last result. In addition, Possible Omega Knowledge does not lead to skepticism about justified belief, as it would in Williamson 2013a's model. Even in the bad case, you are justified in

 $^{^{67}}$ Stalnaker 2009, p. 406 defends cliff-edge knowledge. Hawthorne and Magidor 2010, p. 1092 object to the defense. Goodman 2013 and Williamson 2013b discuss cliff-edge knowledge further, as does Weatherson 2013, p. 67.

believing the tree is between 80 and 120 feet tall, because you omega know this in the good case. Unfortunately, however, Possible Omega Knowledge implies that Justified Margins fails. On this theory, you would justifiably believe what you omega know when reality matches appearance. But when reality matches appearance you do not omega know Margin for Error. For all you omega know, the tree is 80 feet tall. If the tree is 80 feet tall, Margin for Error is false. Finally, in this model Possible Omega Knowledge would entail that knowledge does not imply justification. When the tree is 100 feet tall, you know it is between 90 and 110 feet tall; but you would only be justified in believing it was between 80 and 120 feet tall. (See Lasonen-Aarnio 2010 for defense of 'unreasonable knowledge'.) This view would say that you can know something even though you shouldn't believe it, in the sense that you do not have an excuse for believing it. You shouldn't believe it because you know that any knowledge you may have is imperfect, because you know you don't omega know it.

In Chapter 1, I observed that Margin for Error follows from Safety and Possibility. Safety says that you know p iff you could not easily have believed p falsely. Possibility says that for any height x, if the tree is x feet tall and you believe it is not x-m feet tall, then it could easily have been x-m feet tall while you believed it was not x-m feet tall.

In this model, Margin for Error fails. Safety is valid. But Possibility fails. Possibility fails because what could easily have happened tracks normality, not counterfactual similarity.

Counterfactual similarity tracks the real height of the tree. If the tree is 80 feet tall, it being 79 feet tall is more counterfactually similar than it being 100 feet tall. But it is significantly less normal for the tree to be 79 feet tall than for the tree to be 80 feet tall. So the tree could not easily have been 79 feet tall even when it is 80 feet tall.

When the tree is 80 feet tall, you know it is between 80 and 120 feet tall. This does not violate Safety, because you couldn't easily have believed falsely that it was between 80 and 120 feet tall. In order for this belief to be false, the tree would have to be 79 feet tall or less. But this could not easily have happened, even when the tree is actually 80 feet tall. This could not easily have happened because the tree being 79 feet tall is significantly less normal than the tree being 80 feet tall. 68

 $^{^{68}}$ Dokic and Égré 2009 offer a related analysis of inexact knowledge. They distinguish two kinds of knowledge: perceptual and reflective. They claim Margin for Error is true of perceptual knowledge; but we reflectively rather than perceptually know Margin for Error. They claim that when you perceptually know that p, you reflectively know that you perceptually know that p. The Margin for Error argument is thereby deflated: from knowing that the tree is not 50 feet tall, closure now produces reflective rather than perceptual knowledge that the tree is not 51 feet tall. Margin for Error does not apply to this knowledge.

My account also distinguishes properties of perceptual and reflective knowledge. But I do so differently. First, I define reflective knowledge as knowing that one knows, while they think of reflective knowledge as simply a different kind of knowledge. Second, I deny that when you perceptually know that p, you reflectively know that p. After all, given my analysis of reflective knowledge this would say that if you know p, then you know that you know that you know that p. Third, I allow that you can perceptually know that Margin for Error is true.

In the next sections, I develop models of Fragility and Variable Margins, and explore several of their implications.

5.4 Fragility

My model of Fragility has three key features. First, no matter what condition you are in, it is always possible for you that you are in the good case, where reality matches appearance. Second, in the good case, you know at least as much as in any other case. Third, if you know p in the good case, then you omega know p in the good case. These three features jointly imply Fragility. Suppose you know p. Then it is possible that you know p in the good case. So it is possible that you omega know p. (The first two of these features are shared by all of the models discussed in this book; the last is unique to Fragility.)

I again let a world be a pair (a,r) of an apparent and real height of the tree. I showed above that Williamson accepts Distance, so that you always know less as the distance between reality and appearance increases. In order to validate Fragility, I must deny this assumption. Sometimes, when the distance between appearance and reality increases, there is no change to what you know. Here, my model of Fragility shares a structural property with my model of Reflective Luminosity. Not every decrease in normality produced a decrease in knowledge. Fragility and Reflective Luminosity differ in where knowledge stays fixed. With Fragility, you know the same in the good case as in any case that is epistemically possible in the good case. With Reflective Luminosity, you know the same thing in epistemically worse cases where appearance departs somewhat from reality.

While I reject Distance, I retain Weak Distance (the left to right direction of Distance), which again says that if $|a-r| \leq |a-r'|$, then $E(a,r) \subseteq E(a,r')$. If the distance between a and r is no greater than the distance between a and r', then you know at least as much at (a,r) as at (a,r'); but it is possible that you know no more at (a,r) than at (a,r'). Suppose the tree seems to be 100 feet tall. Now compare what you know when the tree is 100 feet tall with what you know when the tree is 99 feet tall. In the former case, reality matches appearance exactly; in the latter case, reality is slightly different than appearance. Distance implies that when the tree is 99 feet tall, you know strictly less than when the tree is 100 feet tall. By contrast, Weak Distance allows that you may know the same thing in both cases. Weak Distance guarantees the first two features of my model: that the good case is always possible, and that you know at least as much in the good case as in any other case.

Holding fixed Weak Distance, Fragility holds at every world if and only if KK holds in the good case where reality matches appearance.⁶⁹ That means that when reality matches appearance and you know p, it follows that you omega

⁶⁹ Fragility requires (a,a) to see some (r,a) where $E^*(r,a) \subseteq E(a,a)$. Since $|a-a| \le |r-a|$, Weak Distance implies that $E(a,a) \subseteq E(r,a)$ and hence $E^*(a,a) \subseteq E^*(r,a)$. So Weak Distance implies that $E^*(a,a) \subseteq E^*(r,a) \subseteq E(a,a)$, and so whenever S knows A at (a,a), she also knows that she knows A. Conversely, suppose that the KK holds at (a,a), so that $E^*(a,a) \subseteq E(a,a)$. Weak Distance implies that $E(a,a) \subseteq E(r,a)$ for all (r,a). So every (r,a) sees (a,a), where $E^2(a,a) \subseteq E(r,a)$. So Fragility holds at every world.

know p (so that $E^*(a, a) = E(a, a)$). In the good case, knowing implies omega knowing.

In this way, Fragility says there is an inner sanctum wherever reality is close enough to appearance. For some distance around the apparent value, you know exactly as much as if reality perfectly matched appearance. This inner sanctum includes exactly the heights that are epistemically possible when reality matches appearance. Throughout the inner sanctum, you know exactly the same thing. Throughout the inner sanctum, you experience automatic iterated knowledge. In the most extreme case, you know the same thing even at the smallest height consistent with what you know when reality matches appearance (say, 90 feet). When the tree is 90 feet tall, you know that the tree is at least 90 feet tall, because this is also what you know when the tree is 100 feet tall. This condition corresponds to the idea that normally, KK holds. When conditions are normal and you know something, you also know that you know it.

The idea of an inner sanctum corresponds to a particular conception of normality. In addition to thinking about one situation being more or less normal than another, we can also think about a situation being either normal or abnormal. Then the idea behind this model is that any normal situation is significantly more normal than any abnormal situation. When things are normal, you know that things are normal, because abnormality is significantly less normal than normality.

Cohen and Comesaña 2013 (p. 27) defends the inner sanctum on the basis of ex post predictions about assertion: 'suppose that I go outside and it feels like 70 degrees. You are inside getting dressed up, and you ask me how cold it is. I reply that it feels like 70 to me, and so it is at least 65. You then acquire a thermometer and come to know that it is actually 69 degrees. If [Distance] is right, I did not really know that it was at least 65 degrees—and so there is something wrong with my telling you that it could not be as cold as 64' (Cohen and Comesaña 2013 p. 26). The idea is that you are permitted to assert what you would know in the best case, as long as the tree's height is in fact consistent with the what you know in the good case as long as the tree's height is consistent with what you know in the good case.

Again, in earlier chapters, I introduced the idea that principles like Margin for Error or KK could be normally true, even if they aren't universally valid. In this model, the KK principle is true in any situation where the distance between appearance and reality is no more than m. In this sense, it is normally true. (I'll show in a moment that Margin for Error can fail when the distance between appearance and reality is m; in this sense, Margin for Error is not normally true in the model.)

To validate Fragility without KK, I'll agree with Williamson 2013a about how knowledge behaves outside of the inner sanctum, where the tree's height is so different from its apparent height that the real height is epistemically impossible in the good case. In this way, I depart from Cohen and Comesaña 2013, who accept an inner sanctum, but also posit cliff-edge knowledge outside of the inner

 ${\rm sanctum.}^{70}$

Generalizing, the model says that knowledge works differently depending on whether the tree is within m of its apparent value. In that region, you know that you know anything that you know. Outside that region, you don't. Again, this corresponds to the idea that knowledge works differently when conditions are normal than it does when conditions are abnormal.

(71)
$$(a,r)E(a',r')$$
 iff $a = a'$ and
$$\begin{cases} r' \in [a-m, a+m] & \text{if } r \in [a-m, a+m] \\ r' \in [a-(|a-r|+m), a+(|a-r)+m)] & \text{if } r \notin [a-m, a+m] \end{cases}$$

For illustration, consider Figure 4.




Figure 4: a = 100, m = 10

The tree appears 100 feet tall, and the margin m is 10 feet. When the tree is 100 feet tall, you know that it is between 90 and 110 feet tall. This demonstrates that the margin for error is 10. In this model, you know the same thing whenever the tree is between 90 and 110 feet tall. This is the inner sanctum. Throughout this region, you know that the tree is 90 and 110 feet tall. Throughout this region, anything you know is something that you know you know. This validates Fragility. No matter how tall the tree is, anything you know is something that you would also know if the tree were 100 feet tall. No matter how tall the tree is, it is epistemically possible for you that the tree is 100 feet tall. So if you know p, then it is possible that you omega know p. Within [90, 110], epistemic

⁷⁰Here, I also depart from Stalnaker 2015. His theory agrees with mine throughout the inner sanctum. But outside the sanctum, he predicts that you know nothing. This validates KK at the cost of skepticism whenever your conditions are less normal than is possible in the good case: an all or nothing approach to knowledge.

accessibility is transitive. So the same thing is known at (100, 90), (100, 100), and (100, 110): that the tree is between 90 and 110 feet tall.

While Fragility is valid, KK fails. Accessibility is intransitive once the tree is less than 90 feet tall. When the tree is 80 feet tall, you know only that the tree is between 70 and 130 feet tall. But when the tree is 70 feet tall, you know less; simply that the tree is between 60 and 140 feet tall. So when the tree is 80 feet tall, you know it is at least 70 feet tall, but you don't know that you know this. In this region of the model, Margin for Error is true. In this way, Fragility predicts that Margin for Error is a phenomenon that only arises when your evidence is inaccurate.

Omega Skepticism is blocked. When the tree appears 100 feet tall, and is between 90 and 110 feet tall, you omega know that it is between 90 and 110 feet tall. Within this inner sanctum, rational assertion and behavior more generally is guaranteed. Outside of the inner sanctum, by contrast, you can know without omega knowing. If assertion is only governed by a knowledge norm, then within this region you will be permitted to assert p even though you will be forbidden from defending yourself if someone asks you how you know p (indefensible assertion). If assertion is governed by Omega Assertion, then within this region you will not be permitted to assert anything about the tree's height.

On the other hand, I'll show later that this theory predicts that no matter how tall the tree is, you are always justified in believing that you omega know it is within 90 and 110 feet. If justification provides an excuse for assertion and action, then rational action is excused when the appearances are inaccurate.

Like my model of Reflective Luminosity, and unlike Williamson's model, this theory is discontinuous. Small changes in how far reality departs from appearance can produce large changes to what is known. For example, as the height of the tree changes from a-m to $a-m-\epsilon$, the agent suddenly shifts from knowing that the tree is at least a-m to knowing only that the tree is at least a-2m.

One way to think of this model is in terms of barriers to knowledge. When the tree is a-m feet tall, you know only that it is within [a-m,a+m]. Here, your epistemic position is as strong as it can ever be. Further improvements in the match between reality and appearance have no effect on your epistemic position, because your powers of discrimination have already reached their limit.

Now consider Margin for Error, the principle that if you know the tree is not x feet tall, then the tree is not slightly larger than x feet tall. The validity of Margin for Error is incompatible with KK holding locally at any world, let alone in the good case where reality matches appearance. Margin for Error rules out the possibility of cliff-edge knowledge, which occurs whenever the tree's height is n and you know it is at least or at most n. Margin for Error fails in the model above because the model has cliff-edge knowledge. In particular, the model implies that there is cliff edge knowledge when the tree is m feet shorter or m feet taller than it appears. Then, and only then, you know that the tree is either at least or at most the height that it actually is.

Fragility requires cliff-edge knowledge, and invalidates Margin for Error.

We saw in Chapter 2 that Reflective Luminosity also invalidated Margin for Error. While Reflective Luminosity invalidated Margin for Error, it nonetheless allowed that you know Margin for Error in the good case where reality matched appearance. By contrast, Fragility delivers the opposite kind of view. Margin for Error fails in a case that is epistemically possible in the good case, and so Margin for Error is not known in the good case. When the tree is 90 feet tall, you know it is not slightly less than 90 feet tall. Since you don't know Margin for Error in the good case, you are never justified in believing Margin for Error. In this respect, Reflective Luminosity offers a better error theory of Margin for Error than Fragility.

On the other hand, Fragility also has several advantages here over KK and Reflective Luminosity. While Fragility requires cliff-edge knowledge, it is committed to fewer cases of cliff-edge knowledge than existing models of KK or Reflective Luminosity. Recall the model of KK. There, cliff-edge knowledge was pervasive. Whenever the tree's height was at least m feet away from its apparent height, you had cliff-edge knowledge. Or recall the model of Reflective Luminosity. There, cliff-edge knowledge was not pervasive. But it occurred many times, and could occur no matter how far from its apparent height the tree was.

Williamson 2013b objects to the model of KK from before precisely on these grounds: 'Weirdly, [the model] predict[s] that one has cliff-edge knowledge just in case appearances differ severely enough from reality ... a hefty difference between appearance and reality is just what one needs for perfect discrimination, on one side or the other. In [Weatherson 2013]'s terms, this violates the principle that you don't know more by measuring worse. Moreover, since [the model] validate[s] the KK principle, whenever one has cliff-edge knowledge that the real value is at least [r], one also knows that one knows that it is at least [r]. Fragility handles this challenge. True, it requires cliff-edge knowledge. But it says that you have cliff-edge knowledge only when you are in the relatively good case where reality is quite close to appearance. By contrast, Reflective Luminosity is structurally similar to KK in requiring that cliff-edge knowledge happens even in bad cases where reality is arbitrary far from appearance. In this way, even though Fragility rejects Margin for Error, it retains some of the spirit of the principle. More precisely, Margin for Error holds except when the tree's height is m distance away from its apparent height.

Summarizing, Reflective Luminosity and Fragility differ regarding where Margin for Error holds and where it fails. Reflective Luminosity posits that Margin for Error holds at all cases epistemically possible in the good case; but it fails at some case that is epistemically possible at one of those cases. By contrast, Fragility holds that Margin for Error fails at a case that is epistemically possible in the good case. In sum, then, Reflective Luminosity says that Margin for Error is a property that your knowledge satisfies as your belief forming processes become more reliable; while Fragility says that Margin for Error is a property your knowledge satisfies as your belief forming processes become less reliable.

With the status of Margin for Error clarified, I now turn to some features of justified belief in the model. In this model of Fragility, justified belief is

defined as possible knowledge, and equivalently as possible omega knowledge. You are justified in believing p iff you know p in the good case where reality and appearance match iff you omega know p in the good case.

In Chapter 2, I considered two theses about justified belief in KK:

- (72) a. If you are justified in believing you know p, then you are justified in believing you know that you know p.
 - b. You are justified in believing that if you know p, then you know that you know p.

My model of Fragility validates both principles. If you are justified in believing you know p, then in the good case you omega know p, and so you are justified in believing you know that you know p. In addition, since KK holds in the good case, you know in the good case that if you know p, then you know that you know p. So you are justified in believing this conditional. In this way, Fragility explains the appeal of KK. It seems valid because you are always justified in believing it is true.

My model of Fragility also validates:

(73) **Strong Belief.** If you are justified in believing p, then you are justified in believing you know p.⁷¹

Strong Belief says that when you justifiably believe p, your justification is strong enough to suggest you know p. Here is why Strong Belief is valid. If you are justified in believing p, then you know p when conditions are good. But my model of Fragility says that when conditions are good and you know something, you know that you know it. So you are justified in believing that you know it.

Strong Belief is also valid in models of KK. But my model of Reflective Luminosity rejects this principle, and rejects the thesis that you are justified in believing that if you know p, then you know that you know p. Rather, my model of Reflective Luminosity only validates the principle that if you are justified in believing you know, then you are justified in believing that you know that you know. In this way, Reflective Luminosity and Fragility build very different relationships between knowledge and justification.

In Chapter 3, I noted that Fragility entails the JK thesis, that if you know p, then you are justified in believing that you know p. The JK thesis is valid in all of the models considered in this chapter, because all of the models predict that you know the most in the good case.

Williamson 2013a introduced appearance/reality models to give an account of Gettier cases, where you have justifiedly believe something true without knowing it. As before, let B be an accessibility relation modeling justified belief. In this model, I say that you justifiedly believe what you know (and, equivalently, omega know) in the good case: (a, r)B(a', r') iff (a, a)E(a', r') iff a = a' and $r' \in [a - m, a + m]$. (By contrast, in models of Reflective Luminosity one can either define justification in this way, or define it non-equivalently in terms of

⁷¹See Stalnaker 2006 for endorsement.

what you omega know in the good case, in which case m would be replaced with 2m.) As before, let B(a,r) be the set of (a',r') that are B-related to (a,r). B(a,r) is the strongest thing you are justified in believing at (a,r).

Appearance/reality models contain Gettier cases. For example, suppose that the tree is 80 feet tall and appears 100 feet tall. In my model and in Williamson 2013a's, you know that the tree's height is between 70 and 130 feet tall. You are justified in believing it is between 90 and 110 feet tall. Now consider the claim that it is either 90-110 feet tall, or 80 feet tall. This claim is true and you are justified in believing it. But you don't know it, because it is false when the tree is 70 feet tall, and this is consistent with what you know. So this is a Gettier case.

Williamson 2013b defines a special class of Gettier cases which structurally resemble fake barn cases. In this class of 'purely veridical' Gettier case, an agent fails to know p despite having no false justified beliefs. More precisely: (a, r) has purely veridical Gettier cases just in case B(a, r) is true at (a, r) even though B(a, r) is strictly smaller than E(a, r) (more precisely: iff $(a, r) \in B(a, r)$ and $B(a, r) \subset E(a, r)$.).

In the famous fake barn case, you are traveling through the countryside when you see a barn. Unbeknownst to you, all of the other barn facades in the area are fake barns.⁷² Plausibly, you fail to know by perception that you are looking at a barn, because your environment is too unsafe. Purely veridical Gettier cases resemble fake barn cases because they do not involve reasoning from a false lemma. Return to the previous example, where you believe the tree is either 90-110 feet tall or 80 feet tall. You are only justified in believing this claim because you are justified in believing that the tree is 90-110 feet tall. But this stronger claim is false. By contrast, in purely veridical Gettier cases you would fail to know a true justified belief, even though there is no stronger justified belief which is false.

Return to Williamson 2013a's model from before, as described in Figure 1. That model had purely veridical Gettier cases. When the tree was 90 feet tall, you were justified in believing that it was between 90 and 110 feet tall. And you were justified in believing this. But you didn't know it; all you knew was that it was between 80 and 120 feet tall.

Weak Distance implies that Fragility is valid if and only if there are no purely veridical Gettier cases. For suppose Fragility is valid. Then KK holds at (a, a), and so for every (a, r) in B(a, r), B(a, r) = E(a, a) = E(a, r). Conversely, the absence of purely veridical Gettier cases implies that KK holds at (a, a), and so implies the validity of Fragility. In ruling out purely veridical Gettier cases, I agree with the theory in Cohen and Comesaña 2013, and depart from the theories in Williamson 2013a, Goodman 2013, Weatherson 2013, and Carter 2018. This shows that KK is not required in order to rule out purely veridical Gettier cases. Like the unknowability of dubious assertions, this condition is equivalent to Fragility.⁷³

 $^{^{72}}$ See Goldman 1976.

 $^{^{73}}$ Another advantage of this theory is that it predicts that 'what you justifiably believe

Do purely veridical Gettier cases exist? One reason to reject them is that it allows for an elegant statement of the connection between knowledge, justification, and truth. A true, justified belief fails to be knowledge only if it is implied by a false justified belief. On the other hand, the phenomenon of fake barn cases suggest that perhaps there are purely veridical Gettier cases. Nonetheless, fake barn cases are controversial, and so perhaps can't be used to adjudicate the status of Fragility.⁷⁴

Alternatively, the defender of Fragility could contest that fake barn cases are purely veridical. When you are in fake barn country, you are plausibly justified in believing that you are not in fake barn country. You are also plausibly justified in believing that you are looking at a real barn, and that you are not in fake barn country. This would be a false justified belief that is implied by your true justified belief that you are looking at a barn.

Before concluding, I briefly consider the question of improbable knowing. Williamson 2013a observes that Distance generates cases of improbable knowing. At (a,a), E(a,a) is the strongest known proposition. Distance implies that E(a,a) is not known at the worlds inside E(a,a) besides (a,a). Although E(a,a) is known at (a,a), you consider it unlikely conditional on what you know at (a,a) that E(a,a) is known. At every epistemic possibility other than (a,a), E(a,a) is not known.

To make this precise, I follow Williamson 2011 and Williamson 2014 and introduce an evidential probability function Pr. I let the evidential probability $Pr_{(a,r)}$ at world (a,r) come from conditionalizing a prior Pr on E(a,r), the agent's knowledge at (a,r). (In this way, I temporarily bracket concerns from earlier chapters about the thesis that knowledge is evidence.) Improbable knowing occurs at (a,r) when there is a proposition p that is known at (a,r) while the probability that it is known falls below a threshold t. For any proposition p, let $\mathbf{K}p = \{(a',r') : E(a',r') \subseteq p\}$ be the set of worlds at which p is known. Then:

(74) You have improbable_t knowledge at (a, r) if and only if $\exists p : (a, r) \in \mathbf{K}p$ and $Pr_{(a,r)}(\mathbf{K}p) \leq t$.

Distance implies that improbable knowledge is pervasive. At any world (a, r), E(a, r) is known at (a, r), but is not known at any world (a, r') where the distance between r' and a exceeds that between r and a.

Since I replace Distance with Weak Distance, I can prevent improbable

Since I define justified belief so that $B(r^*, a) = E(a, a)$, Weak Disposition to Know is valid in my model of Fragility. In particular, let m^* be the distance between a and the highest value in E(a, a). Then E(a, a) is believed and known throughout the inner sanctum within m^* distance of a, where reality and appearance are sufficiently similar. On the other hand, for criticism of Weak Disposition to Know, see Williamson 2013b p. 87.

is known in all normal worlds with the same appearances' (Goodman 2013, building on Lasonen-Aarnio 2010). Williamson 2013b formulates a weak version of this principle:

⁽i) Weak Disposition to Know. For any r^* , there is some $0 < m^* \le m$ where if $|a-r| \le m^*$, then $E(a,r) \subseteq B(a,r^*)$.

⁷⁴For discussion, see Gendler and Hawthorne 2005.

knowledge. As the distance between reality and appearance grows, the epistemic possibilities cannot diminish. But they may sometimes stay the same. To avoid improbable knowing, I can create bands of constancy. As you move from (a, r) to worlds (a, r') further from r but still inside E(a, r), you can for a while retain the same epistemic possibilities, so that E(a, r') = E(a, r). More precisely, E(a, r) has a band of constancy at E(a, r) of length E(a, r) if and only if E(a, r) = E(a, r). At any world E(a, r), the probability of knowing E(a, r) conditional on E(a, r) will increase as a function of the length of the band of constancy at E(a, r) if you know the same thing E(a, r) as the height of the tree moves further away from E(a, r), then you have a significant probability of continuing to know E(a, r).

Reflective Luminosity also has interesting consequences for improbable knowing. If you know that you know p, then the probability of p on your evidence (be it knowledge or omega knowledge) is 1. If knowledge is evidence, then my model of Reflective Luminosity also prevents many cases of improbable knowing. When the tree is 90 feet tall, you know it is between 80 and 120 feet. You continue to know this at almost all worlds in this range, except when the tree is either 80 or 120 feet. So it is very probable conditional on your knowledge that you know the tree is between 80 and 120 feet tall. By contrast, if evidence is omega knowledge, then the model allows improbable knowing. When the tree is 100 feet tall, you only know you know that the tree is between 80 and 120 feet tall. It is slightly less than 50% likely conditional on your omega knowledge that you know the tree is between 90 and 110 feet tall.

Appealing to evidential probabilities also allows the defender of Fragility to develop another error theory for Margin for Error. In my model of Fragility, Margin for Error only fails in one situation that is possible in the good case. This situation is the least normal one possible in the good case. On some theories of normality, this means that such a situation would also be the least probable one possible in the good case (Goldstein and Hawthorne forthcomingb). This all suggests that when you are in the good case, the probability that Margin for Error is true conditional on what you omega know is extremely high. (In addition, you'd be justified in having this high credence in any situation.) Perhaps this explains the appeal of Margin for Error: it seems true because it probably is true. On the other hand, this view predicts that it would be rational to wonder whether Margin for Error is true, and that it would be rational to think that Margin for Error might not be true. For those who find this inappropriately credulous, the model of Reflective Luminosity offers a more satisfying error theory.

I've now developed models of Fragility, and models of Reflective Luminosity. My model of Fragility invalidated Reflective Luminosity; my model of Reflective Luminosity invalidated Fragility. In fact, the two principles are compatible. Consider the following model, which combines structural features of m previous models:

```
(75) (a,r)E(a',r') iff a=a' and:

a. if 0 \le |a-r| \le m, then r' \in [a-m,a+m]

b. if m < |a-r| \le 2m, then r' \in [a-(|a-r|+m),a+(|a-r|+m)]
```

```
c. if 2m < |a-r| \le 3m, then r' \in [a-3m,a+3m]
d. if 3m < |a-r| \le 4m, then r' \in [a-(|a-r|+m),a+(|a-r|+m)]
e. if 4m < |a-r| \le 5m, then r' \in [a-5m,a+5m]
f. ...
```

This model predicts that KK holds in the good case where reality matches appearance. So it validates Fragility. But it also validates Reflective Luminosity. To see why, notice that outside of m from a, the model replicates the structure of my earlier model of Reflective Luminosity. Margin for Error holds when the tree's height is greater than m and less than or equal to 2m from a. In this region, KK fails: when the tree is slightly further than 2m from a, you know but do not know that you know that it is not much more than 3m from a. But failures of KK are contained: when the tree is within 3m and 4m from a, you know that the tree is at most 4m feet away from a. So when the tree is slightly more than 2m from a, you know that you know that it is at most 4m from a. You also omega know this. While the model validates Reflective Luminosity, this model does not validate Known Margins. Margin for Error is not normally true; nor is it normally known. In this way, the model loses some of the power of the earlier model of Reflective Luminosity.

5.5 Variable Margins

According to Variable Margins, knowledge is governed by a margin for error that can change in size. In the case of tree height, the idea is that the margin of error for knowledge changes as a function of the tree's actual height.

As usual, I assume Appearance Luminosity: the only epistemic possibilities are those where the apparent height of the tree is the same as it actually is. The epistemically possible tree heights are centered on the apparent height, and extend at least as far from the apparent height as the real height. The epistemically possible tree heights extend slightly beyond the real height, in order to accommodate a margin for error.

TO validate Variable Margins, I let the margin for error vary as a function of the distance between appearance and reality. I let $m^*(a,r)$ represent this variable margin, and I assume that $m^*(a,r)$ is always positive. The epistemically possible real heights include all those that are within the sum of |a-r| and $m^*(a,r)$ from the apparent value.

(76)
$$(a,r)E(a',r')$$
 iff $a = a'$ and $r' \in [a - (|a-r| + m^*(a,r)), a + (|a-r| + m^*(a,r))]$

This assures that you know as much as possible, consistent with (i) the appearances being luminous; (ii) your knowledge being centered on the apparent height of the tree; (iii) your knowledge of the height being consistent with the actual height of the tree; and (iv) your knowledge satisfying a margin for error defined by $m^*(a,r)$. In particular, (iv) says that if the real height of the tree is r and the apparent height is a (greater than r), then it is epistemically possible that the tree is $r - m^*(a,r)$ feet tall.

As long as $m^*(a, r)$ is positive, this model validates Variable Margins. To reach specific predictions about knowledge, all that is left is to specify how the margin $m^*(a, r)$ changes as a function of the distance between a and r. I'll now give one toy proposal for how this works, which illustrates the relevant structural properties of Variable Margins. The proposal is defined in terms of a fixed underlying margin m, which you can think of as the margin for error in the good case where reality matches appearance:

(77)
$$m^*(a,r) = \begin{cases} \frac{2m-|a-r|}{2} & \text{if } |a-r| < 2m\\ \frac{4m-|a-r|}{2} & \text{if } 2m < |a-r| < 4m\\ \dots \end{cases}$$

On this proposal, there are a series of 'limit points', where your knowledge approaches a cliff, and which produce omega knowledge. Where m is 10 feet and the tree appears 100 feet tall, the cliffs occur when the tree is 80, 60, 40, or 20 feet tall. In other words, the cliffs occur when the real height is 2m, 4m, 6m, or 8m away from its apparent height. As the tree's height approaches one of these cliffs, the margin for error decreases, approaching but never reaching 0. In particular, the margin for error is always half of the distance to the nearest cliff. a-r measures the distance between reality and appearance. 2m is the distance of the nearest cliff from a. So when you are near the good case, the relevant question is how far |a-r| is from 2m. When appearance and reality match at 100 feet and the nearest cliff is at 80 feet (2m from a), the margin for error is 10 feet. This margin is found by taking half of the difference between 2m and 0. As reality begins to depart from appearance, the distance to the cliff diminishes. So when the tree appears 100 feet tall and is only 90 feet tall, the cliff is only 10 feet away, and so the margin is only 5 feet. Again, this is found by dividing in half the distance between 2m and |a-r|. The closer the tree gets to 80 feet, the smaller the margin becomes. But the margin is always greater than 0, as long as the tree is above the cliff of 80 feet. When the tree is actually 80 feet tall the relevant cliff is instead 60 feet tall, and the margin is again 10 feet.

To see this model in action, consider Figure 5. Here, the margin for error in the good case, m, is 10 feet. As the tree's height departs from appearance, the margin for error shrinks, approaching but never reaching 0.75

In this model, Omega Skepticism fails. When the tree appears 100 feet tall and is so, you omega know that it is greater than 80 feet tall. You know that the tree is 90 feet or taller. If the tree had been 90 feet tall, you would have know that it was 85 feet or taller. If it had been 85 feet or taller, you would have known it was 82.5 feet or taller. As the tree's height approaches 80 feet, you continue to know that it is greater than 80 feet tall.

Open Knowledge also holds. The strongest thing you omega know is always an open interval. For example, when the tree appears 100 feet tall and is so, you omega know that it is greater than 80 feet tall. Your omega knowledge is

 $^{^{75}\}mathrm{See}$ Bonnay and Égré 2009 p. 190 for another model of knowledge that validates Variable Margins.




Figure 5: a = 100, m = 10

inconsistent with the tree being exactly 80 feet tall. If the tree had been 80 feet tall, you would have only known that the tree was 70 feet tall or greater. So if your omega knowledge was consistent with the tree being 80 feet tall, it would also be consistent with the tree being 70 feet tall. In this model, the strongest thing you know at any world is a closed interval. In addition, for any number n, the strongest thing you n-know is a closed interval. But in this model, the strongest thing you omega-know is an open interval. In this way, the strongest thing you omega know could never be the strongest thing that you know.

Since the model validates Open Knowledge, it does not generate cliff-edge knowledge. Suppose that when the tree appeared 100 feet tall and was so, you omega knew it was at least 80 feet tall. In that case, you would have cliff-edge knowledge: when the tree was 80 feet tall, you would know that it was at least 80 feet tall. By contrast, in this model you never have cliff-edge knowledge. In this way, the model respects the idea that your knowledge is inexact.

The model rejects KK, Reflective Luminosity, and Fragility. When the tree appears 100 feet tall and is so, you know that it is at least 90 feet tall. But you don't know that you know this; if the tree had been 90 feet tall, you wouldn't have known it was at least 90 feet tall. You do know that you know that the

tree is at least 85 feet tall. But you don't know that you know that you know this; if the tree had been 85 feet tall, you wouldn't have known it was at least 85 feet tall. Finally, even though you know that the tree is 90 feet tall, you also know that you don't omega know that the tree is 90 feet tall; you fail to omega know that the tree is 90 feet tall at every epistemic possibility, and so you know that you don't omega know it.

In some ways, Variable Margins is more like Reflective Luminosity than Fragility in its conception of how the inexactness of knowledge changes as a function of the normality of one's situations. According to this model, the margin for error is largest when reality matches appearance, and shrinks as reality departs from appearance. This violates the principle from Weatherson 2013 and Williamson 2013b, discussed above, that "you don't know more by measuring worse". Rather, as the quality of your measurement decreases, the fineness of your discrimination increases. On the other hand, in some way Fragility also has some version of this feature, since the margin decreases to nothing as you move from the good case to the maximum departure between reality and appearance consistent with what you know in the good case. ⁷⁶

(i)
$$m^*(a,r) = \begin{cases} \frac{2m - |a - r|}{2} & \text{if } |a - r| < 2m \\ \frac{4m - |a - r|}{1.5} & \text{if } 2m < |a - r| < 4m \\ \dots \end{cases}$$

This proposal implies that you have marginally less knowledge each time you pass a cliff. The proposal does not affect how much omega knowledge you have each time you pass a cliff.

Another proposal would increase the gaps between cliffs as the distance between reality and appearance increases:

(ii)
$$m^*(a,r) = \begin{cases} \frac{2m-|a-r|}{2} & \text{if } |a-r| < 2m \\ \frac{6m-|a-r|}{2} & \text{if } 2m < |a-r| < 6m \\ \dots \end{cases}$$

This proposal implies that you have marginally less omega knowledge each time you pass a cliff. Each of these different proposal reflect different conceptions of what it means to say that your powers of discrimination weaken as your measurement device becomes less accurate.

⁷⁶One way to mitigate this concern about Variable Margins would be to let the rate of decrease in the size of the margin itself decrease as the distance between reality and appearance increases. Consider for example the following structure:

6 Conclusion

6.1 Comparisons

In this book, I've developed three theories of how we can omega know many ordinary claims about the world. In this concluding chapter, I summarize the benefits and costs of each theory, which are laid out in the table below. This summary draws on the entire book. But most of the discussion will make sense for those who haven't read chapter 5.

| | RL | F | VM |
|--|-----|-----|-----|
| -Avoids Omega Skepticism | + | + | + |
| -Consistent with Omega Assertion and Norm Iteration | + | + | + |
| -Three iterations seems like the same as four iterations | + | - | - |
| -It is weird to ask people whether they know that | + | _ | - |
| they know | | | |
| -Unconfident examinee | + | +/- | + |
| -Margin for Error is known, justifiably believed, and | + | _ | - |
| normally true | | | |
| -Testifiers with different reliability produce different | - | + | + |
| amounts of knowledge | | | |
| -Revenge Murine Research | - | + | + |
| -Revenge unconfident examinee | - | +/- | + |
| -You can remember seeing without being certain | - | + | + |
| -Possible Omega Knowledge | - | + | - |
| -You are not justified in believing dubious assertions | - | + | - |
| -KK is justifiably believed and normally true | - | + | - |
| -Compatible with anti-luminosity | - | + | + |
| -You only have cliff-edge knowledge when conditions | - | + | +/- |
| are good | | | |
| -Strong belief: if you are justified in believing, then | - | + | - |
| you are justified in believing you know | | | |
| -No purely veridical Gettier cases | - | + | - |
| -Memory Experiment | +/- | - | + |
| -You can know without justifiably believing you know | + | - | + |
| -You can have omega knowledge outside of the good | + | - | + |
| case | | | |
| -Systematic error theory for Margin for Error | - | - | + |
| -Extends from analogue to digital knowledge | + | + | - |
| -You are permitted to assert closed claims about tree | + | + | - |
| height. | | | |

Comparisons

First, all three theories respond to the arguments from Chapter 1 of the book. All three theories deny Omega Skepticism; they are all ways of allowing that you can omega know many things about the world. For this reason, each theory

can be combined with Omega Assertion and Norm Iteration, without leading to skepticism about rationally permissible assertion or action.

But the three theories differ significantly in various predictions about what you know or justifiably believe. In addition, each theory gives a quite different error theory for Margin for Error, and a different account of the appeal of the KK principle.

Throughout the book, I've considered four main benefits of Reflective Luminosity, compared to Fragility or Variable Margins.

- 1. First, Reflective Luminosity explains why high knowledge iterations collapse. Again, it seems strange to grant that someone has three iterations of knowledge without having four. Reflective Luminosity predicts this; Fragility and Variable Margins do not.
- 2. Second, Reflective Luminosity explains why we rarely if ever ask anyone whether they know that they know. According to Reflective Luminosity and Omega Infallibilism, knowing that you know is the same state as knowing for sure, and we often ask people whether they know something for sure. By contrast, Fragility and Variable Margins predict that knowing that you know is its own unique state, and so leave it a mystery why the state is elusive.
- 3. Third, Reflective Luminosity (like Variable Margins) offers a simple account of the unconfident examinee. The agent knows, without knowing that they know. Variable Margins can also give this simple diagnosis. By contrast, Fragility requires a more complicated analysis. According to Fragility, the unconfident examinee who knows the answer must also not know that they fail to omega know. But this is strange, if the unconfident examinee knows that they don't believe that they know.
- 4. Fourth, Reflective Luminosity offers an error theory for why Margin for Error seems true. This error theory has several parts. First, Reflective Luminosity says that the principle seems true because normally it is true. In fact, it is known when conditions are sufficiently good. By contrast, the theories I've developed which validate Fragility and Variable Margins must deny these claims. In addition, Reflective Luminosity (when combined with Possible Knowledge) predicts that no matter what condition you are in, you are justified in believing that Margin for Error is true, because the principle is known in the good case. By contrast, Fragility and Variable Margins both deny that you are justified in believing Margin for Error.

On the other hand, I have also considered six potential challenges for Reflective Luminosity, compared to Fragility or Variable Margins.

1. Reflective Luminosity rules out certain chains of testifiers with a particular structure. It doesn't allow that there are a pair of Very Reliable and Less Reliable testifiers whose differing reliability of testimony transmits

- second-order and first-order knowledge respectively. By contrast, Fragility and Variable Margins both allow this.
- 2. Reflective Luminosity rules out revenge versions of Murine Research. In the revenge case, Mia the researcher knows that she knows that Accuphine causes hyperactivity in mice. But she is not certain of this, and so continues inquiry into the question, opening an e-mail about further test results. Fragility and Variable Margins allow for this kind of case, but Reflective Luminosity rules it out (when combined with Omega Infallibilism).
- 3. Third, Reflective Luminosity is also potentially vulnerable to revenge versions of unconfident examinee. Imagine that the unconfident examinee is guessing about the question of whether they know when Queen Elizabeth died, and they reliably guess that they do know, in exactly the situations where they once learned when she died. Reflective Luminosity says that if the examinee knows that they know, then the examinee omega knows. By contrast, Variable Margins allows that the unconfident examinee could know that they know, without realizing this, in a way that explains their lack of confident. (Fragility can make a similar prediction, although still faces the problem of explaining how the examinee's knowledge could be consistent with omega knowing.)
- 4. A fourth challenge for Reflective Luminosity concerns iterated factive mental states. Someone can plausibly remember seeing that p without being certain of p. Fragility and Variable Margins allow this, while Reflective Luminosity (together with Omega Infallibilism) rule it out.
- 5. A fifth challenge for Reflective Luminosity concerns justification. When combined with Possible Omega Knowledge, Reflective Luminosity predicts that knowledge does not imply justification (Variable Margins faces a similar challenge). By contrast, Fragility is tailor-made to deliver this prediction. In Chapter 1, I observed that Possible Omega Knowledge follows from a truth norm for belief along with Norm Iteration and Possible Permission, which said that you are justified in a behavior iff for all you know, the behavior is permitted. So the defender of Reflective Luminosity must give up one of these principles.
- 6. Sixth, in response to this challenge, the defender of Reflective Luminosity may accept Possible Knowledge instead (while the defender of Fragility may accept both Possible Omega Knowledge and Possible Knowledge). But in that case, the theory predicts that you can justifiably believe dubious assertions.

Throughout the book, I've considered five benefits for Fragility, compared to Reflective Luminosity or Variable Margins. (These benefits are separate from the previous costs for Reflective Luminosity compared to Fragility.)

1. First, Fragility offers a good explanation of why KK is attractive. In my model of Fragility, the KK principle is normally true. Reflective Luminosity

and Variable Margins do not make this prediction. Relatedly, Fragility predicts that you are justified in believing the KK principle. KK holds in the good case, and so it is consistent with what you know and omega know that KK is true. Reflective Luminosity and Variable Margins do not make these predictions.

- 2. Second, Fragility (and Variable Margins) is compatible with the general anti-luminosity of all mental states. By contrast, Reflective Luminosity implies that second-order knowledge is luminous.
- 3. Third, Fragility differs from Reflective Luminosity and Variable Margins in its treatment of cliff-edge knowledge. Again, you have cliff-edge knowledge if your perceptual faculties have maximum discriminatory power: the tree would be x feet tall, and you would know that it was at least x meters tall. Fragility predicts that your perceptual faculties are only this powerful in the good case where you are maximally reliable; by contrast, my model of Reflective Luminosity permits you to have cliff-edge knowledge even when your perceptual faculties are malfunctioning significantly. Strictly speaking, my model of Variable Margins doesn't admit any cliff-edge knowledge. But it does allow arbitrarily close approximations of cliff-edge knowledge even when your perceptual faculties malfunction.
- 4. Fourth, Fragility makes the controversial prediction that justified belief is strong, in the sense that if you are justified in believing, then you are justified in believing that you know. Reflective Luminosity and Variable Margins do not make this prediction.
- 5. Fifth, Fragility makes the controversial prediction that there are no purely veridical Gettier cases: if you have a true, justified belief without knowing, then you must also have a false, justified belief. Reflective Luminosity and Variable Margins deny this.

I've also considered three distinctive challenges for Fragility, that are avoided by Reflective Luminosity and Variable Margins:

- 1. First, Fragility struggles with Memory Experiment. In this case, Joan remembers and thereby knows something, even though they also know there is a 50% chance that her memories are fake. Here, Joan plausibly knows that she should not be certain. Fragility (together with Omega Infallibilism) then imply that Joan fails to know. Reflective Luminosity and Variable Margins avoid the prediction that knowledge implies justified certainty. (Although Reflective Luminosity may face a challenge from revenge versions of Memory Experiment.)
- 2. Second, Fragility makes the controversial prediction that knowledge is 'lustrous', so that whenever you know, you are justified in believing that you know. Reflective Luminosity and Variable Margins do not have this commitment.

3. Third, my model of Fragility only allowed for omega knowledge in situations that are epistemically possible in the good case. Outside of this 'inner sanctum', omega knowledge was elusive. By contrast, my models of Reflective Luminosity and Variable Margins both allowed for Omega Knowledge even when one's situation is moderately bad. (This last challenge may be an artifact of the model of Fragility, rather than of the principle itself.)

This leaves Variable Margins. Much of the appeal of Variable Margins has already been covered: Variable Margins avoids many of the challenges faced by Reflective Luminosity and Fragility. For example, it doesn't face revenge versions of the unconfident examinee, Murine Research, or Memory Experiment. It allows you to have chains of iterated factive mental states without certainty. It allows chains of testifiers with differing reliability to produce different iterations of knowledge. It is compatible with anti-luminosity, and also with the failure of knowledge to be lustrous. In addition, Variable Margins offers a genuine alternative to Margin for Error, avoiding cliff-edge knowledge in all cases.

On the other hand, we've seen that Variable Margins does not predict that Margin for Error is normally true, or that you are justified in believing it, or that you know it in the good case. And Variable Margins also does not predict that KK is normally true, or that you are justified in believing it. In this way, Variable Margins loses out on some of the theoretical benefits of Reflective Luminosity or Fragility.

Finally, in Chapter 4 I developed two more challenges for Variable Margins, not faced by Reflective Luminosity or Fragility:

- 1. First, Variable Margins only applies to 'analogue' knowledge, where you are measuring a continuous quantity like height; but it can't make sense of 'digital' knowledge of discrete quantities like calendar dates.
- 2. Second, Variable Margins predicts that the strongest thing you omega know about a tree's height must be an open interval; this suggests that no one is ever in a position to assert that a tree's height falls in a closed interval.

I am unsure where this leaves us. Reflective Luminosity and Fragility each face serious potential counterexamples, involving versions of the unconfident examinee, Murine Research, and Memory Research. Those swayed by such counterexamples may be attracted to Variable Margins. On the other hand, the failure of Variable Margins to generalize beyond 'analogue' knowledge seems serious

Other readers may be willing to bite bullets regarding particular cases, in exchange for some of the structural features of Reflective Luminosity or Fragility. Again, there are choices. Reflective Luminosity emphasizes that Margin for Error is normally true, and so is always justifiably believed. By contrast, Fragility says that KK is normally true, and so is always justifiably believed. This corresponds to a different perspective on how one's powers of perceptual discrimination are affected by error. According to Fragility, you only have cliff-edge knowledge

when conditions are good; by contrast, Reflective Luminosity allows you to have cliff-edge knowledge in worse conditions.

The other major difference between Reflective Luminosity and Fragility concerns the theory of justified belief. Fragility fits nicely with Possible Omega Knowledge, according to which you are justified in believing p iff it is possible that you omega know p. Reflective Luminosity does not fit well with this principle: accepting both principles requires denying that knowledge implies justification, and also requires denying that you are justified in believing Margin for Error. Instead, defenders of Reflective Luminosity do best to accept Possible Knowledge, according to which you are justified in believing p iff it is possible that you know p. Unfortunately, this allows that you can justifiably believe dubious assertions.

Personally, I find myself most attracted to either Reflective Luminosity or Fragility, rather than Variable Margins. For me, the failure to generalize beyond analogue knowledge is fatal. In addition, I am attracted to the structural theses about knowledge and justification that are made possible by Reflective Luminosity and Fragility. On the other hand, I find it difficult to determine which of Reflective Luminosity and Fragility is a better theory.

The last option would be to accept more than one of Reflective Luminosity, Fragility, and/or Variable Margins. I see no special advantage of going this way. Each of these three principles has the same basic advantage: defeating Omega Skepticism. But each principle has its own costs. For example, once Fragility is accepted, Margin for Error can no longer be normally true. So the combined theory would lose Reflective Luminosity's ability to give an error theory of Margin for Error. In addition, once Reflective Luminosity is accepted, the model will require cliff-edge knowledge outside of the good case. This means that even when your epistemic situation is quite bad, you can still possess incredible powers of discrimination. Alternatively, accepting Variable Margins in addition to either Reflective Luminosity or Fragility would remove cliff-edge knowledge. But cliff-edge knowledge would only be avoided for analogue knowledge. For these reasons, I believe that the best theory of knowledge should accept one of Reflective Luminosity, Fragility, or Variable Margins, but not more than one of these principles.

I don't know which principle is best. Each has its costs and benefits, and these features are difficult to compare. The arguments of this book are not conclusive. Much work remains to explore the full range of consequences of principles like Reflective Luminosity, Fragility, and Variable Margins. My aim above all has been to open up new territory in the study of knowledge. These principles offer interesting weakenings of the KK principle. They come with a rich profile of costs and benefits. My hope is that future work on knowledge will uncover new principles of greater subtlety and complexity that balance the costs and benefits more carefully.

6.2 Loose ends

6.2.1 Omega Infallibilism

One loose end is Omega Infallibilism, the principle that you are permitted to be certain of p if and only if you omega know p. Throughout the book, I've taken this assumption on board, and used it to explore the varying consequences of Reflective Luminosity and Fragility. But one possibility would be to accept Reflective Luminosity or Fragility while rejecting Omega Infallibilism.

In particular, one option would be to reject the right to left direction of Omega Infallibilism, so that omega knowledge does not imply permissible certainty. This direction of Omega Infallibilism could be rejected while retaining the opposite direction, which says that omega knowledge is necessary for permissible certainty. This thesis follows from Norm Iteration, when combined with the thesis that you are permitted to be certain of something only if it is true. If you reject the right to left direction of Omega Infallibilism, then you could allow that rational certainty is a stronger epistemic state than omega knowledge. This kind of theory would fit naturally with work from Beddor 2020a and Beddor 2020b, arguing that certainty plays a crucial role in epistemology.

By giving up Omega Infallibilism, the defender of Reflective Luminosity can solve several problems. First, they can avoid revenge versions of Murine Research: omega knowledge would no longer permit the cessation of inquiry. Second, they can make sense of remembering seeing that p without being certain of p.

The defender of Fragility would also benefit, by avoiding a bad prediction in Memory Experiment. Without Omega Infallibilism, Fragility would allow you to remember something even while knowing that you shouldn't be certain of it.

On the other hand, giving up Omega Infallibilism would also remove one of the benefits of Reflective Luminosity. Knowing that you know would be distinct from knowing for sure; this would raise the question of why it is strange to ask someone whether they know that they know.

In addition, giving up Omega Infallibilism might require defenders of either Reflective Luminosity or Fragility to rethink the nature of justified belief. In Chapter 1, I argued that Omega Infallibilism and Possible Permission imply Possible Omega Knowledge. If belief is interpreted as certainty, then Possible Permission implies that you are justified in believing something iff for all you know you are permitted to be certain of it. Omega Infallibilism then says that you are permitted to be certain iff you omega know.

Without Omega Infallibilism, you might still accept the thesis that you are permitted to believe p only if p is true. When combined with Norm Iteration, this implies that you are permitted to believe p only if you omega know p. From Possible Permission, it then follows that you are only justified in believing p if it is possible that you omega know p. At this point, two options present themselves. First, you could equate belief (in the sense necessary for knowledge) with certainty, and deny that possible omega knowledge is sufficient for justification. Second, you could distinguish belief and certainty, and hold that unlike certainty, omega

knowledge is sufficient for permissible belief, and so hold that possible omega knowledge is equivalent to justified belief. At any rate, I leave as an open question how the opponent of Omega Infallibilism should think about justified belief.

6.2.2 Epistemic Position

Some philosophers think in terms of the concept of the *strength of your epistemic position*. Maybe this tracks the folk concept of *how well you know something*.

One conception of epistemic position is iterative: the strength of your epistemic position regarding p is proportionate to the number of iterations of knowledge you possess regarding p.

Earlier in the book, I suggested that animals can know things without knowing that they know, since they lack the concept of knowledge. This would then imply that you would know better than your dog that there is food in his bowl, even if the dog is looking at the food, while you only know on the basis of inference.

An alternative conception of epistemic position would let this strength be proportionate to the number of iterations of being in a position to know. Perhaps your dog is in a position to know that he knows there is food in his bowl, because if he did possess the concept of knowledge, he could successfully deploy it in this instance.

On these theories of epistemic strength, omega knowledge would play a special role. Anything you omega know (or anything you are in an omega position to know) would be known best.

Another conception of epistemic strength appeals to evidential probability. I am tempted by a phenomenal conception of evidence, according to which your evidence is made up of the your phenomenal experiences. This theory allows you to know things that are not maximally probable on your evidence. In this case, we could say that the epistemic strength of p is proportionate to the evidential probability of p.

Alternatively, we might say that the epistemic strength of p is proportionate to the evidential probability that you know (or omega know) p. This would imply that lottery propositions and Moorean propositions have minimal epistemic strength.

The probabilistic and iterative conceptions of epistemic strength come apart. There can be two propositions p and q, where you possess more iterations of knowledge regarding p than of q, and yet the evidential probability of p is lower than q (and the evidential probability that you know p is lower the evidential probability that you know q).

As a simple example, consider a model with four worlds: 4, 3, 2, and 1. The evidential probability of the worlds is .4, .3, .2, and .1 respectively. At 4, 4 and 3 are possible. At 3, 4 and 3 and 1 are possible. At 2 and at 1, 4 and 3 and 2 and 1 are possible. In this model, at world 4 you know that you know 4 or 3 or 1. You know but do not reflectively know that 4 or 3 or 2. But the probability of 4 or 3 or 2 is .9, while the probability of 4 or 3 or 1 is only .8.

Personally, I don't find the concept of *strength of epistemic position* very useful. I don't have a good sense of what theoretical role it plays. But any defender of the concept must choose between the iterative or probabilistic conception (or offer an alternative), because the two theories are not compatible.

6.2.3 Beyond Safety

In this book, I've thought about iterated knowledge while assuming that knowledge involves belief that is safe from error. A different approach to iterated knowledge would focus on justification instead of safety. In this approach, we could think about principles like Reflective Luminosity and Fragility as flowing from principles governing justification.

As a simple example, I'll start by imagining that knowledge is true, justified belief, and consider how principles about iterated knowledge would correspond to principles about iterated justification. Then, I'll briefly consider how these ideas would generalize to more complex theories, for example where knowledge is safe and justified belief.

Imagine that knowledge is true, justified belief. In addition, assume that knowledge and justification are closed under consequence, and that this analysis of knowledge is known.

Given these assumptions, iterated knowledge reduces to a special kind of iterated justification that I'll call 'factorial'. For example, you know that you know that p iff p is true, you justifiably believe p, and you justifiably believe that you justifiably believe p. Similarly, you 3-know p iff p is true, you justifiably believe p, you 2-justifiably believe p, and you 3-justifiably believe p. Generalizing, in this framework iterated knowledge is true 'factorial' justified belief. Say that you have an n-factorial justified believe that p iff you n-justifiably believe p and you have an (n-1)-factorial justified belief that p, and say that you have a 1-factorial justified belief that p iff you justifiably believe p. Then you n-know that p iff p is true and you have an n-factorial justified belief that p.

In this setting, Reflective Luminosity and Fragility correspond to theses about iterated justification. Reflective Luminosity corresponds to the thesis that if p is true and you have a 2-factorial justified belief that p, then you also have a 3-factorial justified belief that p. This principle follows from the thesis that if you justifiably believe that you justifiably believe p, then you justifiably believe that you justifiably believe p.

Similar ideas apply to Fragility. For example, consider Weak Fragility, the thesis that if you know p, then it is possible that you know that you know p. In this framework, Weak Fragility corresponds to the thesis that if you have a true, justified belief that p, then it is possible that you have a true, 2-factorial justified belief that p. With further simplifications, this is equivalent to the following condition: if (i) p, you have a justified belief that p, and you don't have a justified belief that you have a justified belief that p, or you don't have a justified belief that p, or you don't have a justified belief that p (in symbols: if $p \land Jp \land \neg JJp$, then $\neg J[\neg Jp \lor \neg JJp]$). Among other things, this follows from

an analogue of Weak Fragility for justification: if you are justified in believing p, then you aren't justified in believing that you aren't justified in believing that you are justified in believing p (if Jp, then $\neg J \neg J J p$).

There is a rich tradition of work exploring iteration principles for justification (see for example Smithies 2019). Much of this work has focused on the idea that justification freely iterates (the analogue of KK for justification). One fruitful research project would be to explore the extent to which the various considerations in this tradition can be explained by weaker justification iteration principles.

From Gettier, we learned that knowledge is not true, justified belief. But this leaves open that knowledge could be safe and justified belief. On this proposal, you know p iff (i) you could not easily have believed p falsely; and (ii) your belief that p is justified. In this more complex framework, it is harder to identify interesting joint constraints on safety and justification would correspond to Reflective Luminosity or Fragility. But one possibility is that the safety condition itself freely iterates (as in Greco 2014a), and that whenever you safely believe p, you justifiedly believe that you safely believe p. In that case, the validity of Reflective Luminosity and Fragility would again depend essentially on the structure of justification.

6.2.4 Verbal debates

I say that KK fails, and that omega knowledge rather than knowledge is the state that governs assertion, action, inquiry, and so on. One of my opponents says that KK is true, and knowledge is the state that plays the crucial explanatory role. But maybe this is just a verbal dispute about the word 'knowledge'.

In particular, my opponent can offer the following reinterpretation of the term 'knowledge'. What they mean by 'knowledge' is just the state that I call omega knowledge. In their mouth, it is 'knowledge' that plays a crucial explanatory role, by behaving exactly as I say omega knowledge behaves. By definition, omega knowledge freely iterates: if you omega know, then you omega know that you omega know. So in their language, 'knowledge' freely iterates.

In my theory, knowledge itself does not play a central role in the explanation of assertion, action, inquiry, or whatever. So its hard for me to point to something important that is missed by choosing to use the word 'knowledge' in this way. Its not as if by my lights this will lead to a gap in some central explanatory role that knowledge fills. So it is hard to see how this dispute could be anything but verbal.

At least two issues are not verbal. First, the theories of omega knowledge that I've developed do not posit omega knowledge as a primitive, undefined state. Instead, omega knowledge is defined in terms of knowledge: you omega know when you possess every iteration of knowledge. And on each of my theories, knowledge is distinct from omega knowledge. My opponent chooses to use 'knowledge' to refer to what I call omega knowledge. But this means that my opponent should make room for another state, call it 'proto-knowledge', where (i) you know p if and only if you possess every iteration of proto-knowledge, and

(ii) proto-knowing does not imply knowing. On this proposal, proto-knowledge is something like safe belief; and then knowledge involves possessing every iteration of safe belief.

There is at least one other substantive aspect of the discussion. The theories of omega knowledge that I've developed make predictions about how much you omega know in different situations. After omega knowledge is reinterpreted as 'knowledge', these theories will still disagree about how much you 'know' in differing cases. In particular, my model of Fragility was skeptical about possessing omega knowledge when in a situation that is incompatible with what you know in the good case; my model of Reflective Luminosity instead allowed for this kind of omega knowledge. When we reinterpret omega knowledge as 'knowledge', this difference corresponds to a disagreement about skepticism. The reinterpreted model of Fragility will be skeptical about possessing knowledge in any situation that is incompatible with what you know in the good case. By contrast, my model of Reflective Luminosity will allow you to know things even when your conditions are abnormal in this way.

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.
- Robert Aumann. Agreeing to disagree. *Annals of Statistics*, 6(4):1236–1239, 1976.
- Andrew Bacon. Giving your knowledge half a chance. *Philosophical Studies*, 171 (2):1–25, 2014. doi: 10.1007/s11098-013-0276-6.
- 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.
- Bob Beddor and Carlotta Pavese. Modal virtue epistemology. *Philosophy and Phenomenological Research*, 2019.
- Matthew Benton. Dubious objections from iterated conjunctions. *Philosophical Studies*, 162:355–358, 2013.
- Selim Berker. Luminosity regained. Philosophers' Imprint, 8(2):1–22, 2008.
- Alexander Bird. Justified judging. *Philosophy and Phenomenological Research*, 74:81–110, 2007.
- Denis Bonnay and Paul Égré. Inexact knowledge with introspection. *Journal of Philosophical Logic*, 38(2):179–227, 2009. doi: 10.1007/s10992-008-9087-1.
- Jessica Brown. Subject-sensitive invariantism and the knowledge norm for practical reasoning. $No\hat{u}s$, 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. Practial 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.

- Sam Carter and Simon Goldstein. The normality of error. *Philosophical Studies*, January 2021.
- Sam Carter and Simon Goldstein. Getting accurate about knowledge. *Mind*, pages 1–28, 2022. doi: 10.1093/mind/fzac009.
- David Christensen. Putting Logic in its Place. Oxford University Press, 2005.
- Herbert H. Clark and Catherine R. Marshall. Definite knowledge and mutual knowledge. In Aravind K. Joshi, Bonnie L. Webber, and Ivan A. Sag, editors, *Elements of Discourse Understanding*, pages 10–63. Cambridge, UK: Cambridge University Press, 1981.
- 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.
- Kevin Dorst. Abominable kk failures. Mind, 128(512):1227-1259, 2019. doi: 10.1093/mind/fzy067.
- Fred I. Dretske. Knowledge and the Flow of Information. MIT Press, 1981.
- Pascal Engel. In what sense is knowledge the norm of assertion? Grazer Philosophische Studien, 77(1):45–59, 2008.
- Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Vardi. *Reasoning About Knowledge*. MIT Press, 2003.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, 2009.
- Richard Feldman. Respecting the evidence. Philosophical Perspectives, 19, 2005.
- Richard Foley. Working Without a Net. Oxford University Press, 1993.
- Jane Friedman. Question-directed attitudes. *Philosophical Perspectives*, 27(1): 145–174, 2013. doi: 10.1111/phpe.12026.

- Jane Friedman. Why suspend judging? $No\hat{u}s$, 51(2):302-326, 2017. doi: 10.1111/nous.12137.
- Tamar Szabó Gendler and John Hawthorne. The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124(3): 331–352, 2005. doi: 10.1007/s11098-005-7779-8.
- 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.
- Alvin Goldman. What is justified belief? In George Pappas, editor, *Justification and Knowledge*, pages 1–25. Boston: D. Reidel, 1979.
- Alvin I. Goldman. Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(November):771–791, 1976. doi: 10.2307/2025679.
- Simon Goldstein and John Hawthorne. Counterfactual contamination. *Australasian Journal of Philosophy*, pages 1–17, forthcominga. doi: 10.1080/00048402.2021.1886129.
- Simon Goldstein and John Hawthorne. Knowledge from multiple experiences. *Philosophical Studies*, pages 1–32, forthcomingb. doi: 10.1007/s11098-021-01710-4.
- Simon Goldstein and Daniel Waxman. Losing confidence in luminosity. *Noûs*, pages 1–30, 2020. doi: 10.1111/nous.12348.
- Jeremy Goodman. Inexact knowledge without improbable knowing. *Inquiry*, 56 (1):30–53, 2013.
- Jeremy Goodman and Ben Holguin. Thinking and being sure. August 2021.
- Jeremy Goodman and Bernhard Salow. Taking a chance on kk. *Philosophical Studies*, 175(1):183–96, 2018.
- Jeremy Goodman and Bernhard Salow. Epistemology normalized. *Philosophical Review*, forthcoming.
- 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.
- Daniel Greco. Iteration principles in epistemology 1: Arguments for. *Philosophy Compass*, 10(11):754–771, 2015.

- Daniel Greco. Safety, explanation, iteration. *Philosophical Issues*, 26(1):187–208, 2016. doi: 10.1111/phis.12067.
- Daniel Greco. Justifications and excuses in epistemology. $No\hat{u}s$, forthcoming. doi: 10.1111/nous.12309.
- 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.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is weak. *Philosophical Studies*, 173(5):1393–1404, 2016. doi: 10.1007/s11098-015-0553-7.
- Eliran Haziza. Assertion, implicature, and iterated knowledge. *Ergo: An Open Access Journal of Philosophy*, forthcoming.
- Jane Heal. Common knowledge. *Philosophical Quarterly*, 28(111):116–131, 1978. doi: 10.2307/2219358.
- Ben Holguín. Knowledge by constraint. *Philosophical Perspectives*, pages 1–28, forthcoming. doi: 10.1111/phpe.12153.
- Jonathan Jenkins Ichikawa. Justification is potential knowledge. Canadian Journal of Philosophy, 44:184–206, 2014.
- Daniel Kahneman. Thinking Fast and Slow. Farrar, Straus, and Giroux, 2011.
- 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.
- Christoph Kelp and Mona Simion. Criticism and blame in action and assertion. Journal of Philosophy, 114(2):76–93, 2017. doi: 10.5840/jphil201711426.
- 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.
- Maria Lasonen-Aarnio. Unreasonable knowledge. *Philosophical Perspectives*, pages 1–21, 2010.
- Harvey Lederman. Uncommon knowledge. Mind, 127(508):1069-1105, 2018a. doi: 10.1093/mind/fzw072.
- Harvey Lederman. Two paradoxes of common knowledge: Coordinated attack and electronic mail. *Noûs*, pages 921–945, 2018b. doi: 10.1111/nous.12186.
- Keith Lehrer and Stewart Cohen. Justification, truth, and coherence. Synthese, 55(2):191–207, 1983. doi: 10.1007/BF00485068.
- Wolfgang Lenzen. Recent work in epistemic logic. Acta Philosophical Fennica, 30, 1978.
- Nick Leonard. Epistemological Problems of Testimony. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Summer 2021 edition, 2021.
- Jarrett Leplin. A Theory of Epistemic Justification. Springer, 2009.
- David K. Lewis. Convention: A Philosophical Study. Wiley-Blackwell, 1969.
- Clayton Littlejohn and Julien Dutant. Justification, knowledge, and normality. *Philosophical Studies*, 177(6):1593–1609, 2020. doi: 10.1007/s11098-019-01276-2.
- Dustin Locke. Practical certainty. *Philosophy and Phenomenological Research*, 90(1):72–95, 2015. doi: 10.1111/phpr.12036.
- Matthew Mandelkern. Practical moore sentences. $No\hat{us}$, 55(1):39–61, 2021. doi: 10.1111/nous.12287.
- Matthew Mandelkern and Kevin Dorst. Assertion is weak. *Philosophers' Imprint*, forthcoming.
- David Manley. Safety, content, apriority, self-knowledge. *Journal of Philosophy*, 104(8):403–423, 2007. doi: 10.5840/jphil2007104813.
- Berislav Marušic. 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.

- Ruth Garrett Millikan. Language, Thought, and Other Biological Categories: New Foundations for Realism. MIT Press, 1984.
- 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.
- Dana K. Nelkin. The lottery paradox, knowledge, and rationality. *Philosophical Review*, 109(3):373–409, 2000. doi: 10.1215/00318108-109-3-373.
- Duncan Pritchard. Epistemic Luck. Oxford University Press, 2005.
- Colin Radford. Knowledge-by examples. Analysis, 27:1-11, 1966.
- Baron Reed. A defense of stable invariantism. $No\hat{u}s$, 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. $No\hat{u}s$, 52(1):171-195, 2018. doi: 10.1111/nous.12156.
- Sven Rosenkranz. The structure of justification. Mind, 127(506):309-338, April 2018.
- Sven Rosenkranz and Julien Dutant. Inexact knowledge 2.0. Inquiry: An Inter-disciplinary Journal of Philosophy, 63(8):1–19, 2020. doi: 10.1080/0020174x. 2020.1754286.
- 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.
- Ariel Rubinstein. The electronic mail game: Strategic behavior under "almost common knowledge". American Economic Review, pages 385–391, 1989.
- 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.
- Weng Kin San. Kk, knowledge, knowability. May 2019.
- 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.

- Alex Silk. What normative terms mean and why it matters for ethical theory. In Mark C. Timmons, editor, *Oxford Studies in Normative Ethics, Vol. 5*, pages 296–325. Oxford University Press, 2015.
- Martin Smith. What else justification could be 1. $No\hat{u}s$, 44(1):10-31, 2010. doi: 10.1111/j.1468-0068.2009.00729.x.
- Martin Smith. Between Probability and Certainty: What Justifies Belief. Oxford University Press UK, 2016.
- Declan Smithies. *The Epistemic Role of Consciousness*. New York, USA: Oxford University Press, 2019.
- David Sosa. Dubious assertions. Philosophical Studies, 146(2):269–272, 2009.
- Ernest Sosa. How to defeat opposition to moore. *Philosophical Perspectives*, 13: 137–49, 1999. doi: 10.1111/0029-4624.33.s13.7.
- Robert Stalnaker. Presuppositions. *Journal of Philosophical Logic*, 2:447–457, 1973.
- Robert Stalnaker. Indicative conditionals. *Philosophia*, 5(3):269–286, 1975. doi: 10.1007/bf02379021.
- Robert Stalnaker. On logics of knowledge and belief. *Philosophical Studies*, 128: 169–199, 2006.
- Robert Stalnaker. On hawthorne and magidor on assertion, context, and epistemic accessibility. *Mind*, 118(470):399–409, 2009.
- 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.
- Robert C. Stalnaker. Context and Content: Essays on Intentionality in Speech and Thought. Oxford University Press UK, 1999.
- Dennis W. Stampe. Towards a causal theory of linguistic representation. *Midwest Studies in Philosophy*, 2(1):42–63, 1977. doi: 10.1111/j.1475-4975.1977.tb00027. x.
- 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.
- Brian Weatherson. Margins and errors. Inquiry, 56(1):63–76, 2013.
- 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.
- Christopher Willard-Kyle. Being in a position to know is the norm of assertion. *Pacific Philosophical Quarterly*, 101(2):328–352, 2020. doi: 10.1111/papq. 12305.
- Malte Willer. Lessons from sobel sequences. Semantics and Pragmatics, 10(4): 1–57, 2017.
- Timothy Williamson. In exact knowledge. Mind, 101(402):217-242, 1992. doi: $10.1093/\min(101.402.217)$.
- Timothy Williamson. Vagueness. London and New York: Routledge, 1994.
- 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. Probability and danger. In Amherst Lecture in Philosophy, pages 1–35. http://www.amherstlecture.org/williamson2009/, 2009.
- Timothy Williamson. Improbable knowing. In Trent Dougherty, editor, *Evidentialism and its Discontents*. Oxford University Press, 2011.
- Timothy Williamson. Gettier cases in epistemic logic. Inquiry, 56(1):1-14, 2013a.
- Timothy Williamson. Response to cohen, comesaña, goodman, nagel, and weatherson on gettier cases in epistemic logic. *Inquiry*, 56(1):77–96, 2013b.
- Timothy Williamson. Very improbable knowing. Erkenntnis, 79:971–999, 2014.
- Timothy Williamson. Justifications, excuses, and sceptical scenarios. In Fabian Dorsch and Julien Dutant, editors, *The New Evil Demon.* Oxford: Oxford University Press, forthcoming.
- Juhani Yli-Vakkuri and John Hawthorne. Being in a position to know. *Philosophical Studies*, pages 1–17, forthcoming. doi: 10.1007/s11098-021-01709-x.