DOI: 10.1111/phpr.70001

ORIGINAL ARTICLE

Philosophy and
Phenomenological Research

Knowing to infinity: Full knowledge and the margin-for-error principle

Yonathan Fiat

MIT

Correspondence
Yonathan Fiat, MIT
Email: jonathan.fiat@gmail.com

Abstract

Let's say that *I fully know that p* if I know that *p*, I know that I know that *p*, I know that I know that I know that *p*, and so on. Let's say that *I partially know that p* if I know that *p* but I don't fully know that *p*. What, if anything, do I fully know? What, if anything, do I partially know? One response in the literature is that I fully know everything that I know; partial knowledge is impossible. This response is in tension with a plausible margin-for-error principle on knowledge. A different response in the literature is that I don't fully know anything; everything that I know, I partially know. Recently, Goldstein (forthcoming, 2024) defended a third view, according to which I fully know some things and I partially know other things. While this seems plausible, Goldstein's account is based on denying the margin-for-error principle. In this paper, I show that the possibility of both full knowledge and partial knowledge is consistent with the margin-for-error principle. I also argue that the resulting picture of knowledge is well-motivated.

KEYWORDS
full knowledge, KK, knowledge iterations, margin-for-error, omega knowledge

1 INTRODUCTION

It's natural to think that there are at least two categories of knowledge. There are some things that I know, but I can reasonably doubt whether it's true that I know them. Suppose I have a meeting at 9am in my office. If I leave my house at 7:45am and arrive to my office at 8:45am, I can reflect back and wonder whether I knew that I will make on time. It seems plausible to think that I did know; I make this way five times a week, I know roughly how long it takes, I know on what days there's more traffic, etc. But it also seems plausible that I didn't know that I knew it.

But there are other things about which this seems less plausible. I look around my office and I don't see anyone else. So, I know that I'm alone in my office. And there's no uncertainty about it. It's completely obvious to me that I'm alone in my office. Of course, there are weird skeptical scenarios where assassins are hiding in plain sight, or where my brain was taken out of my body and attached to a virtual reality machine, or where someone slipped hallucinogenic drugs into my coffee. But none of those things is true, or was ever close to being true. So it's not clear how they can affect my state of knowledge.

We can try to draw the distinction between the two kinds of cases by saying that in the first kind of cases, I don't know that I know, and in the second kind of cases, I know that I know. And it seems right, but it's also unsatisfying. We can imagine variations on the first case such that I did know that I knew, but I didn't know that I knew that I knew; or maybe I didn't know that I knew that I knew that I knew. But the second case seems resistant to that. I simply know, or better yet, I fully know. There's no ignorance in the vicinity.

The following definitions can help our discussion:

> 1**-knowledge**: One 1-knows that p if one knows that p.
>
> n **-knowledge (**$n \geq 2$**)**: One n-knows that p if one knows that one $(n-1)$-knows that p.
>
> **Full knowledge**: One fully knows that p if one n-knows that p for every natural number n.
>
> **Partial knowledge**: One partially knows that p if one knows that p but doesn't fully know that p.

Cases like the ones above suggest that both partial knowledge and full knowledge are possible and indeed common. I partially know many things, such that my neighbor's kid was born in May (or was it June?), that this shirt will fit my partner, or that I will get home in time for the game. I fully know that my best friend's name is David, that I hate beige, or that the sun won't set for at least another five hours (after all, it's only 9am).

Are there any reasons to deny the claim the both partial and full knowledge are possible? On some views, partial knowledge is impossible. Whenever I know that p, I fully know that p. This is equivalent to the KK principle: the principle that whenever I know something, I know that I know it. I don't have anything new to add to the already existing extensive discussion of KK in the literature.[1] It seems clear to me that I can sometimes discover that I know something, and be surprised by it, and that if this is true then the KK principle is false.

[1] A non-exhaustive list of recent papers discussing the KK principle: Bird & Pettigrew, 2021; Cohen & Comesaña, 2013; Das & Salow, 2018; Dorr et al., 2014; Dorst, 2019; Fraser, 2022; Goldstein & Hawthorne, 2024; Goodman & Salow, 2018; Greco, 2014, 2015, 2016; Liu, 2020; San, 2023; Stalnaker, 2009, 2015; Williamson, 2021.

If we assume that the KK principle is false, then partial knowledge is possible. Is full knowledge possible? There's no obvious reason to think that the answer is no.[2] Nonetheless, in the epistemic logic literature, it's typically assumed that if the KK principle is false, full knowledge is impossible. It's not completely obvious why, but it's easy to come up with a plausible hypothesis. One of the main reasons to reject the KK principle is the margin-for-error principle: for a belief to be knowledge, is must be appropriately safe from error. If the margin-for-error principle is true, then the KK principle is false (Williamson, 2002). And on the simplest formal models of knowledge where the margin-for-error principle is true, full knowledge is impossible.[3]

The fact that full knowledge is impossible in simple margin-for-error models might be a reason to think that it's indeed impossible, but it's not a conclusive one. We should at least explore other models of the margin-for-error principle and see if full knowledge is possible in them. Maybe they are not overly complicated, and maybe they enjoy some other theoretical virtues, which might make us reconsider the possibility of full knowledge.

So this is what I plan to do. I'll show that the margin-for-error principle is consistent with the possibility of full knowledge. Moreover, I'll argue that the combination yields an intuitively plausible picture. The basic idea behind my view is simple. I can cross the ten feet from my chair to the door of my office in a finite amount of time, even though I must first walk the first five feet, then the next two and a half feet, then the next 1.25 feet, and so on *ad infinitum*. Similarly, I can have full knowledge even though I need some margin-for-error for 1-knowledge, and some further margin-for-error for 2-knowledge, and some further margin-for-error for 3-knowledge, and so on *ad infinitum*.

Recently, Goldstein (2024, forthcoming) argued for a similar position. He argues that both full knowledge and partial knowledge are possible. However, he achieved this result at the cost of giving up on the margin-for-error principle. This is, I think, an overreaction. As we shall see in the next section, it's hard to give up on the margin-for-error principle. Moreover, there's no reason to do it; we can get both the possibility of full knowledge and the margin-for-error principle.[4]

2 | THE MARGIN-FOR-ERROR PRINCIPLE

Mr Magoo (Williamson, 2002) glances at a nearby tree, and attempts to estimate its height. Unfortunately for him, his eyesight is quite poor, so there's a limit on how much he can know. Let T be a variable representing the height of the tree in inches. Margin-for-error principles have the following form:

> **Margin-for-Error**$_\epsilon$ (MFE$_\epsilon$): For every x, if $T = x$, then for all Mr Magoo knows, $T = x + \epsilon$.[5]

Williamson suggests that there's some number $\epsilon > 0$ for which MFE$_\epsilon$ is true. MFE$_\epsilon$, if true, is not supposed to be *merely* true. It's not just that it so happens that Mr Magoo doesn't know that $T \neq x + \epsilon$. It follows from the nature of knowledge: given Mr Magoo limitations, he can't have knowledge that is stronger than what MFE$_\epsilon$ allows.

[2] Williamson (2002, pp. 121–123) considers the question and makes some suggestive comments about why it might be impossible, but they are not conclusive.

[3] Some examples: Goodman, 2013; Weatherson, 2013; Williamson, 2011, 2013a, 2014, 2021.

[4] For more information on Goldstein's theory, see §A.3.

[5] Here and throughout, I use the formulation "for all Mr Magoo knows, p" as a shorthand for "Mr Magoo doesn't know that not p."

However, MFE_ϵ can't be a necessary truth. Mr Magoo can easily learn that $T \neq x + \epsilon$ if someone who measured it tells him so. And he could have had better eyesight and a knack for estimating heights, and then he could have known that $T \neq x + \epsilon$ merely on the basis of sight. So at best, MFE_ϵ is true as long as we keep Mr Magoo's way of knowing, namely by sight, fixed, and as long as we keep his visual abilities fixed. If we change one of those, then MFE_ϵ might be false. But, there might still be some other $\epsilon' > 0$ such that $\text{MFE}_{\epsilon'}$ is true.

2.1 | Motivating the margin-for-error principle

MFE_ϵ is an *anti-luck* principle.[6] If $T = 100$, and Mr Magoo believed $T < 100 + \epsilon$, this belief would be lucky in a way that's inconsistent with knowledge. The following cases demonstrate this point:

> **Truck**. Alice drives a truck and Bob sits in the passenger seat. They see the following sign: "Danger! Low bridge ahead. Maximum height: 9 feet." Knowing that the truck is more than nine feet tall, Bob anxiously asks Alice if they should find another way. Alice reassures him that those signs are always underestimating the height of the bridge and that it'll be fine. They drive under the bridge and the truck just passes underneath it. The truck makes a horrible screeching noise as it grazes the bridge. Alice tells the shocked Bob: "Do you see? I knew it would be fine."
>
> **Cookies**. Charlie and Diana co-teach a class, and they plan to bring cookies to the final lecture. The number of students who attend the last lecture varies widely: it can be anything between 50 and 150. They estimate that roughly 100 students will attend. They decide that Charlie will bake the cookies. Before the lecture, Diana looks at the bag of cookies and asks worryingly if it'll be enough for everyone. Charlie reassures her that every student will get a cookie. As Diana distributes the cookies, she notices that the bag seems to empty very quickly while there are still many students who didn't get a cookie. When she gets to the last row of students, she becomes increasingly anxious that it won't be enough. As she gets to the last student, she reaches into the bag and finds the last cookie. She hands the empty bag back to Charlie, who says "Do you see? I knew it would be enough."
>
> **Coin**. A fair coin is about to be flipped again and again until it lands tails. Edward and Felicia will lose their house if the coin is flipped more than ten times. Edward tells Felicia not to worry, because the coin won't be flipped more than ten times. The coin lands nine times heads and then tails. Edward tells Felicia "Do you see? I knew it wouldn't be flipped more than ten times."[7]

Those three cases end with someone saying "I knew that p." In each case, we can assume that the speaker had believed that p, and had good reasons to believe it. And it's part of the description of the cases that p is true. And yet, there's a strong feeling that the speaker speaks falsely.

There are many potential explanations for what went wrong in those cases. But the natural hypothesis is that the speakers didn't know what they claim to have known. And a natural explanation of this is that there's some margin-for-error principle on knowledge: to know that p, it's not enough that p is true. Something stronger must be true, that allows for some margin for error.

[6] See Williamson (2013b) on *cliff-edge knowledge*.

[7] Adapted from Dorr et al. (2014).

2.2 | Modeling margin-for-error

It's widely believed that if Mr Magoo is in a position to know that ϕ, and he's in a position to know that ψ, and χ is a logical consequence of ϕ and ψ, then he's also in a position to know that χ. For example, he can come to know that ϕ, come to know that ψ, and then infer χ, and thereby come to know that χ.

If this is true, we can model the operator "Mr Magoo is in a position to know" using *Kripke frames*. A Kripke frame is built from a set of nodes, often called "worlds," and an accessibility relation between them. Each node represents a potential state of affairs, and decides all the non-epistemic propositions. A proposition p is known in a node w if it's true in all the nodes accessible from w.

For example, here is a simple model where MFE$_\epsilon$ is true (based on Williamson, 2013a):

PURE MFE MODEL

Let $\epsilon > 0$.

The set of nodes: $\{(a, r) : a, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2. $|a' - r'| \leq |a - r| + \epsilon$

This model is based on the idea that Mr Magoo has some specific *estimate* of the height of the tree; the tree *appears* to him to be some specific height. Let A be a variable representing the apparent height of the tree. The intended interpretation of the model is that in the node (a, r), $A = a$ and $T = r$. MFE$_\epsilon$ is true in every node of the model.[8]

3 | FULL KNOWLEDGE SKEPTICISM

MFE$_\epsilon$ is true in every node of PURE MFE MODEL. Moreover, a stronger principle is also true in every node:

> **Higher-Order Margin-for-Error**$_\epsilon$ (HO-MFE$_\epsilon$): For every n, for every x, if $T = x$, then for all Mr Magoo n-knows, $T = x + n \times \epsilon$.

If there's some $\epsilon > 0$ such that HO-MFE$_\epsilon$ is true, then there's no number x such that Mr Magoo fully knows that $T \leq x$. There's some n such that for all Mr Magoo n-knows, the tree is as tall as a fully grown giant sequoia; there's some n such that for all he n-knows, the tree is as tall as Mount Everest; and there's some n such that for all he n-knows the tree reaches all they way to the moon.

It's natural to ask whether HO-MFE$_\epsilon$ is a mere artifact of the model, or whether it follows from MFE$_\epsilon$. After all, it's a very simple model; many false things are true in it. So it's interesting to see what assumptions can lead to this conclusion.

[8] Goodman (2013) presents a different kind of model where the margin-for-error principle is true. In Goodman's model, there are additional limits on knowledge, so it's easier to see that the margin-for-error principle is an anti-luck principle. See §A.1.2 for details.

TABLE 1 A single induction step in the first argument for full knowledge skepticism.

$\mathbf{P}_n$	For all Mr Magoo n-knows, $T = x + n \times \epsilon$.
$\mathbf{Q}_n$	Mr Magoo n-knows that if $T = x + n \times \epsilon$, then for all he knows, $T = x + (n+1) \times \epsilon$.
$\mathbf{P}_{n+1}$	For all Mr Magoo $(n+1)$-knows, $T = x + (n+1) \times \epsilon$.

3.1 | The argument from full knowledge of MFE$_\epsilon$

The first argument for HO-MFE$_\epsilon$ comes from Williamson (2002). In addition to some knowledge of the height of the tree, Mr Magoo also knows something about his own limitations. For example, He might know that MFE$_1$ is true. If for some $\epsilon > 0$, Mr Magoo fully knows MFE$_\epsilon$, and if knowledge is closed under deduction, then HO-MFE$_\epsilon$ is true.

We can prove it for any value of x independently by induction on n. Suppose that $T = x$; then, we can repeat the argument in table 1 any number of times.

$\mathbf{P}_1$ is true because, by assumption, MFE$_\epsilon$ is true. $\mathbf{Q}_n$ is true because Mr Magoo fully knows MFE$_\epsilon$, so he also n-knows it. $\mathbf{P}_{n+1}$ follows from $\mathbf{P}_n$ and $\mathbf{Q}_n$ and the assumption that knowledge is closed under deduction.[9]

This argument explains why HO-MFE$_\epsilon$ is true in every node of PURE MFE MODEL. In every node, Mr Magoo fully knows MFE$_\epsilon$, and so HO-MFE$_\epsilon$ follows.

3.2 | The modal argument

A second argument for HO-MFE$_\epsilon$ is based on a specific way to think about knowledge. The basic thought is that knowing that p means something like being able to exclude all the nearby worlds where p is false. This way of thinking seems to go hand-in-hand with the margin-for-error principle: knowledge must be safe.[10]

Here's an instructive example. I know that I have hands. By using my vision, I can rule out all the nearby worlds where I don't have hands. Even if I've just barely escaped a chainsaw that would have cut off my hands, and so there's a nearby world where I don't have hands, I can exclude that world. In that world, I would be screaming in pain, covered in blood, etc.

There are far away worlds that I can't exclude where I don't have hands. For example, maybe my hands where cut off while I'm connected to a virtual reality machine that gives me the experience of having hands. But, the basic thought is that I can know that I'm not in one of them, because they are sufficiently far from the actual world.

Using this ideology, we can reformulate MFE$_\epsilon$:

---

[9] If knowledge is closed under deduction, the following argument form is valid:

A For all a n-knows, p.
B a n-knows that if p then q.
C Therefore, for all a n-knows, q.

To see that it's valid, first note that if knowledge is closed under deduction, n-knowledge is also closed under deduction. Now, suppose that the conclusion is false: a n-knows that not q. Then, by **B**, a n-knows both not q and that if p then q. So he n-knows that not p, contradicting **A**.

[10] See, for example, Sosa (1999).

MFE$^*_\epsilon$: For every x, if $T = x$ (in the current world, w), then there's a non-excluded nearby world w' where $T = x + \epsilon$.

To get HO-MFE$_\epsilon$ we need the following stronger principle:

MFE$_\epsilon$-CHAINING: For every x, if $T = x$ (in the current world, w), then there's a non-excluded nearby world w' where MFE$_\epsilon$-CHAINING is true and $T = x + \epsilon$.

The difference between MFE$^*_\epsilon$ and MFE$_\epsilon$-CHAINING is that MFE$^*_\epsilon$ leaves open the possibility that MFE$^*_\epsilon$ is false in the world w'. MFE$_\epsilon$-CHAINING states that there's a world w' where $T = x + \epsilon$ and the principle itself is true.

Just like MFE$_\epsilon$, MFE$_\epsilon$-CHAINING is not supposed to be a necessary truth, because Mr Magoo could have had better eyesight. But the idea is that given his actual eyesight, there's some $\epsilon > 0$ for which MFE$_\epsilon$-CHAINING is true.[11]

Why should we accept that for some $\epsilon > 0$, MFE$_\epsilon$-CHAINING is true? MFE$_\epsilon$ is an anti-luck principle, but anti-luck considerations don't support the stronger MFE$_\epsilon$-CHAINING. We can instead try to motivate it by appealing the idea that knowledge is related to modal distance.

Suppose that MFE$^*_\epsilon$ is true and $T = 100$. Then, there's a nearby possible world that Mr Magoo can't exclude where $T = 100 + \epsilon$. Let's call this world w'. It's a world where the tree is slightly taller, perhaps because it got a bit more sunlight. But w' is not a world where Mr Magoo has better eyesight: while there are worlds where Mr Magoo has better eyesight, they are much further away. So in w', $T = 100 + \epsilon$ and Mr Magoo has the same eyesight. And, because MFE$^*_\epsilon$ is true in the actual world because of Mr Magoo's eyesight, it's also true in w'. If this reasoning is valid, then MFE$_\epsilon$-CHAINING is true.

If there's some $\epsilon > 0$ such that MFE$_\epsilon$-CHAINING is true, then HO-MFE$_\epsilon$ is true as well. To see why, let w_0 be the actual world, and assume $T = x$. MFE$_\epsilon$-CHAINING guarantees that there's a sequence $w_1, w_2, w_3, \ldots$ such that for every n, in w_n, $T = x + n \times \epsilon$ and w_{n+1} is a non-excluded nearby world. And this means that for every n, for all Mr Magoo n-knows, he is in w_n, which means that for all he n-knows, $T = x + n \times \epsilon$.

This argument gives a different explanation why HO-MFE$_\epsilon$ is true in all the nodes of PURE MFE MODEL. If we interpret the nodes as possible worlds, and "w can access w'" as "in w, w' is a non-excluded nearby world," we get that MFE$_\epsilon$-CHAINING is true in all the nodes in the model. And this is enough to imply that HO-MFE$_\epsilon$ is true.

4 | FULL SAFETY

So it seems like the margin-for-error principle leads to the conclusion that full knowledge is impossible. But this is false. To see why, we need to first talk about physical safety. Physical safety must be governed by something like the margin-for-error principle. But as we'll see, it doesn't mean that full safety is impossible.

[11] Williamson (2021) endorses something very similar to MFE$_\epsilon$-CHAINING.

4.1 | Iterations of safety

A child, Sarah, stands on a cliff. Typically, she is safe if she is far enough from the edge. In simple cases, there's some distance d_1 such that Sarah is safe iff her distance from the edge is at least d_1.

In addition to being safe from falling off the cliff, Sarah can be *safely safe* from falling off the cliff. Suppose, for example, that ten feet from the cliff edge there's a sign saying "Danger! Do not cross." And suppose further that the sign is perfectly accurate: Sarah is safe iff she is in front of the sign, and unsafe if she is behind it.

Issac, Sarah's father, wants her to stay sufficiently far from the sign. He doesn't want her to be too close to the dangerous zone, behind the sign. In other words, he wants her to be safely in the safe zone. It's a very natural thing to desire; there's nothing weird about wanting her to be safely safe.

So there's some distance, d_2, such that Sarah is safely safe iff her distance from the sign is at least d_2. Which means that she is safely safe iff her distance from the edge is at least $d_1 + d_2$.

This idea generalizes. There's some distance d_3 such that Sarah is safely safely safe iff her distance from the cliff edge is $d_1 + d_2 + d_3$. And in general, there are more distances, $d_4, d_5, \ldots$ such that Sarah is n-safe iff her distance from the edge is at least $d_1 + d_2 + \cdots + d_n$.

4.2 | Modeling safety

We can reason about safety using Kripke frames.[12] For simplicity, we'll assume that Sarah is limited to movement along a single line, perpendicular to the cliff edge. By choosing an arbitrary origin, and orienting the numbers to point away from the edge, we can represent a location using a single number. Let X denote Sarah's location and Y the edge's location. Sarah is not falling off the cliff iff $X > Y$.

Let S be the *safety* operator, such that "Sarah is safe from falling" can be written as $S(X > Y)$, and "Sarah is safely safe" can be written as $S(S(X > Y))$.

Here is a simple model for the situation:

PURE MFS MODEL

Let $\epsilon > 0$.

The set of nodes: $\{(x, y) : x \in \mathbb{R}, y \in \mathbb{R}\}$.

(x, y) can access (x', y') iff:

1. $y' = y$; and
2. $|x' - x| \leq \epsilon$

A node in this model is a pair (x, y) where x is Sarah's location and y is the cliff location. $S(X > Y)$ is true in a node (x, y) if $X > Y$ is true in all the nodes accessible from it.

PURE MFS MODEL is perfectly analogous to PURE MFE MODEL. The following principle is true in PURE MFS MODEL:

[12] Doing so presupposes that safety is closed under deduction: if it's safe that p and it's safe that q, then it's safe that p and q. See Williamson (2009) for a defense of this claim.

Margin-for-Safety$_\epsilon$ (MFS$_\epsilon$): Sarah is safe iff $X > Y + \epsilon$.

Just like MFE$_\epsilon$, it seems undeniable that there's some $\epsilon > 0$ such that MFS$_\epsilon$ is true. Sarah is only safe if she is sufficiently far from the edge.

The following principle is also true in PURE MFS MODEL:

Higher-Order Margin-for-Safety$_\epsilon$ (HO-MFS$_\epsilon$): For every n, Sarah is n-safe iff $X > Y + n \times \epsilon$.

HO-MFS$_\epsilon$ is the safety version of HO-MFE$_\epsilon$. If it's true for some $\epsilon > 0$, then each iteration of safety requires the *same* distance. It means that full safety is impossible: no matter how far Sarah is from the cliff edge, there's some n such that Sarah is not n-safe.

We asked whether HO-MFE$_\epsilon$ is a mere artifact of PURE MFE MODEL or it's independently motivated. Similarly, we can ask whether HO-MFS$_\epsilon$ is a mere artifact of PURE MFS MODEL or it's independently motivated. Very minimal variations on the arguments we've seen can be used to argue for it.[13]

4.3 | Against the higher-order margin-for-safety principle

It turns out that for every $\epsilon > 0$, HO-MFS$_\epsilon$ is false. Let's suppose that Sarah is safe iff and only if she is at least ten feet from the cliff edge. It means that for every $\epsilon \neq 10$, HO-MFS$_\epsilon$ is false. So we need to show that HO-MFS$_{10}$ is false.

Next to my home, there's a playground which is a bit elevated above its surrounding. Parents often watch their kids play there, and can get distressed if the children are too close to the edge, judging them to be unsafe.

Imagine a child playing next to the cliff at Mount Thor (4,100 feet), the greatest vertical drop on earth. If a parent watched the child playing, they would demand the child would be much further from the edge than what they would demand in the playground. A distance of three feet might be enough for safety in the playground, but it's definitely not enough on the top of Mount Thor.

This difference can be easily explained by the fact that the falling off the edge of the playground, while potentially painful, isn't very dangerous. But falling off the cliff at Mount Thor is fatal. And it teaches us something important about safety: the distance required for safety varies with the danger involved. If we keep all other factors fixed, the greater the danger, the greater the distance required for safety.

The same idea applies for higher-order safety. Recall that Issac, Sarah's father, wants her to be safely far from the danger sign. But it would be bizarre for him to think that, if the sign's distance from the cliff edge is ten feet, then to be safely far from the sign, the distance between Sarah and the sign needs to also be at least ten feet. Even though it'll be bad if Sarah crosses into the unsafe zone, it'll be less bad than if she falls off the cliff.

And this means that Sarah is safely safe at a distance smaller than twenty feet. So HO-MFS$_{10}$ is false. So for every $\epsilon > 0$, HO-MFS$_\epsilon$ is false.

[13] Williamson (2002) seems to endorse the inference from MFS$_\epsilon$ to HO-MFS$_\epsilon$.

4.4 | Upshot: Safety isn't a modal stability operator

Even though it's a simple observation, there's something surprising about it. It's common to think that "p is safe" is true if p is true in all the nearby worlds. And PURE MFS MODEL captures this idea. But this automatically leads to HO-MFS$_\epsilon$, which is false.

To see why, let's suppose that safety is indeed a modal stability operator. Assume that $\epsilon = 10$. In $(0,0)$, $X > Y$ is false. $(0,0)$ is close to the node $(9.9,0)$, so in $(9.9,0)$, $S(X > Y)$ is false. $(9.9,0)$ is close to $(19.8,0)$, so in $(19.8,0)$ $S(S(X > Y))$ is false. But as we saw, this conclusion shouldn't follow from the assumption that $\epsilon = 10$.

It means that we really can't think about safety as a kind of modal stability. Maybe it's true that for $X > Y$ to be safe, it must be true in all nearby possible worlds. But for $S(X > Y)$ to be safe, it's not the case it must be true in all the nearby possible worlds. If it had been true, HO-MFS$_\epsilon$ would have been true, but it isn't.

4.5 | An alternative model of safety

Luckily, HO-MFS$_\epsilon$ doesn't follow from MFS$_\epsilon$. To see this, consider the following model:

FULL SAFETY MODEL

The set of nodes: $\{(\epsilon, x, y) : \epsilon > 0\}$.

(ϵ, x, y) can access (ϵ', x', y') iff:

1. $y' = y$; and
2. $\frac{\epsilon}{2} \leq \epsilon' \leq \epsilon$; and
3. $|x' - x| \leq \epsilon - \epsilon'$

Suppose again that Sarah is safe iff she is at least ten feet from the cliff edge, so $\epsilon = 20$. According to FULL SAFETY MODEL, Sarah is safely safe if she is at least 15 feet from the edge, she is safely safely safe if she is at least 17.5 feet from the edge, and so on. It means that according to FULL SAFETY MODEL, full safety is possible. If Sarah is at least 20 feet from the edge, she is n-safe for every n, which means that she is fully safe.

FULL SAFETY MODEL captures a natural way to think about safety. To see why, consider the following case:

> **Party**. George organizes a party. He invited one thousand people, and allowed them to invite others, but asked that they respond with how many people they intend to bring. One hundred people confirmed that they would come. George knows that people might come without sending a confirmation.

If George only gets enough food for one hundred people, then even if exactly one hundred people show up, it's not safely enough food. Much like the cases from §2, the fact that there's enough food is lucky.

In order to be safe, George wants a *safety buffer*: he knows he needs to get enough food for more people. For example, maybe he can only be safe if he has a safety buffer of 50%, which means that to be safe he needs to get enough food for a hundred and fifty people

Suppose however that George is still anxious. He is not sure if a safety buffer of fifty people is enough. Then he might decide to increase it: to have a second safety buffer on top of the first one. He might, for example, opt for an additional safety buffer of 50% of the first safety buffer. That is, he might order food for additional twenty-five people, so overall he orders food for one hundred and seventy-five people.

The general point is that when we think about safety, we think about safety buffers. And we can make sense of multiple safety buffers. But each additional safety buffer is thought of as an additional buffer on top of the previous one, not as a completely independent safety buffer.

4.6 | Safety and interests

The fact that HO-MFS$_\epsilon$ is false for every $\epsilon > 0$ doesn't show that full safety is possible. Full safety is possible according to FULL SAFETY MODEL, but there are many alternative ways safety could work. The failure of HO-MFS$_\epsilon$ merely opens up the possibility of full safety, as it shows that at least one very simple way to think about safety doesn't work.

However, the same considerations that lead us to deny HO-MFS$_\epsilon$ lend some support to the possibility of full safety. As we saw, the minimal distance for safety is different between the playground and Mount Thor, and it's different between the first iteration of safety and the second one. And the explanation for that is that the distance required for safety changes with the danger. It's more dangerous to fall off Mount Thor than it's to fall off the edge of the playground, and it's more dangerous to fall off the cliff than to cross into the unsafe zone.

Those observations suggest that the distance required to safety always varies with the interest we have in being protected from the potentially bad result. Sarah's interest in not falling off the cliff is quite large; her interest in not crossing into the unsafe zone is substantial, but significantly less. And this explains why more distance is required for the first iteration of safety.

It seems plausible that Sarah's interest in being far away from the cliff edge runs out at some point. It's possible to question it: maybe there's always some interest at more distance. For example, one could think even if she's one mile away from the edge, she has some interest in being two miles away, but this interest is so small that it's trumped by other interests. However, as soon as we try to apply it to real world cases, it becomes a very weird view. For example:

> **Driving**. Hannah plans to drive from New York to Los Angeles. She worries about falling into the Grand Canyon, and therefore takes it into account when deciding on the route.

It would be very strange to think that even if Hannah chooses a route that passes one hundred miles away from the Grand Canyon, she still has some interest in choosing a route that passes further away from the Grand Canyon.

So let's suppose that our interests in more distance run out at some point. It still doesn't automatically mean that full safety is possible. It's possible to think that while our interests run out, there's always room for more iterations of safety: it's just that we have no interest in them.

This view makes safety very mysterious. Suppose that at one hundred miles from the Grand Canyon, Hannah is 1000-safe, and that to be 1001-safe she needs additional ten feet. What explains the fact that ten feet are required rather than twenty? By assumption, her interests play no role at this point. And we know that at close distances, the interests determine the minimal distance for safety. So it turns out that after some point, something else takes over in explaining it. And this is rather mysterious.

Those considerations suggest that full safety is indeed possible. And if it's possible, it can teach us something about knowledge.

5 | FULL KNOWLEDGE

Now that we see how full safety could be possible, let's turn back to knowledge. Recall Isaac, Sarah's father, who looks at her and worries that she'll cross the ten feet mark. The margin-for-error principle says that Issac can only know that she is safe if Sarah is sufficiently far from the ten feet mark. If she is just about ten feet from the cliff edge, Isaac can't know that she is at least ten feet from it. If he believed it, it would be a true belief, but only luckily so; so it wouldn't be knowledge.

Suppose that Isaac can only know that Sarah is at least ten feet from the edge if she is at least eleven feet from it. The margin-for-error principle again tells us that Isaac can't know that he knows it if Sarah is just about eleven feet from the edge. If he believed that he knew it, this belief would be true, but luckily so. Sarah needs to be further away from the edge for Isaac to know that he knows that she's at least ten feet from the edge.

How much further away? HO-MFE$_\epsilon$ gives a concrete answer: another foot away. If one foot was required for the belief that she is at least ten feet from the edge to be safe enough to be knowledge, then one foot is required for the belief that Isaac knows that Sarah is at least ten feet from the cliff to be knowledge.

But as we saw, the safety version of HO-MFE$_\epsilon$, HO-MFS$_\epsilon$, is false. To be safely safe, Sarah doesn't need to be twenty feet from the edge. And this suggests that for Isaac to know that he knows that Sarah is at least ten feet from the edge, Sarah doesn't need to be at least twelve feet from the edge.

5.1 | Full knowledge model

The comparison with safety is suggestive, but we might worry that it can't work for knowledge. Knowledge has its own logic, and HO-MFE$_\epsilon$ might be forced on us, maybe because of the arguments in §3.

To alleviate this worry, we can construct a model for knowledge that is based on the ideas behind FULL SAFETY MODEL:

FULL KNOWLEDGE + MFE MODEL

The set of nodes: $\{(\epsilon, a, r) : a, r, \epsilon \geq 0\}$.

(ϵ, a, r) can access (ϵ', a', r') iff:

1. $a' = a$; and
2. $\frac{\epsilon}{2} \leq \epsilon' \leq \epsilon$; and
3. $|a' - r'| \leq |a - r| + \epsilon - \epsilon'$

On the intended interpretation of this model, in the node (ϵ, a, r), $A = a, T = r$, and Mr Magoo's margin is $\frac{\epsilon}{2}$. Every node $(2\epsilon, a, r)$ agrees with corresponding node in PURE MFE MODEL on first-order knowledge. But it also allows for full knowledge.[14]

5.2 | Is it a margin-for-error model?

I promised to present a model where knowledge is governed by the margin-for-error principle, but full knowledge is still possible. And then I presented FULL KNOWLEDGE + MFE MODEL. But some might complain that this model doesn't fulfill the promise. While it's true that in nodes of the form $(2\epsilon, *, *)$, if the height of tree is x inches, Mr Magoo doesn't know that it's less than $x + \epsilon$ inches, there are other nodes in the model. For example, in some nodes of the form $(\frac{\epsilon}{100}, *, *)$, the height of the tree is x and yet Mr Magoo knows that the height of the tree is less than $x + \frac{\epsilon}{2}$.

To make this worry more concrete, let's consider the models from Goldstein (2024) (see §A.3). In Goldstein's models, full knowledge is possible and the margin-for-error principle is false. But naturally, it's true in some nodes and false in others. In some of Goldstein's models, it's only false when *conditions are abnormal*, or *in the bad case*. So, one might think that Goldstein had already shown that it's possible to have full knowledge, at the price of having some nodes where the margin-for-error principle is false.

This is, however, a misunderstanding. As we discussed in §2, for any concrete value of ϵ, MFE_ϵ is not a necessary truth. The value of ϵ is partly an empirical question: how bad Mr Magoo's eyesight is. It's possible to be ignorant of it.[15] In order to allow for ignorance of the margin, we must introduce nodes where the margin is different from what it actually is. But it doesn't mean that the margin-for-error principle is false.

The margin-for-error principle states that given Mr Magoo's actual limitations, there's some $\epsilon > 0$ such that MFE_ϵ is true. For a model to allow it, it must have, for every possible height of the tree and every possible apparent height of the tree, a node that represents Mr Magoo's situation. And this is what FULL KNOWLEDGE + MFE MODEL achieves.

Goldstein's models don't have this property. According to him, even if we keep Mr Magoo's quality of vision fixed, there are cases (perhaps abnormal or bad) where the height of the tree is x inches, and Mr Magoo knows that it's x inches or less.[16]

[14] This is a paper about artifacts in modeling, so it's important to acknowledge that this model is also not perfect. The point of this model is to show that once we give up on full knowledge of the margin, the possibility of full knowledge becomes consistent with the margin-for-error principle. The model also gives us some understanding of how they can both be possible. However, this model also has some questionable features. For example, in this model, if Mr Magoo's margin is ϵ, then he knows that his own margin is less or equal to ϵ, which might be problematic. But this isn't an essential feature of the model; see §B.

[15] Maybe it's necessary that there's some $\epsilon > 0$ such that MFE_ϵ is true, and therefore $\exists\epsilon > 0 : \text{MFE}_\epsilon$ is always known. But this claim is true in FULL KNOWLEDGE + MFE MODEL.

[16] FULL KNOWLEDGE + MFE MODEL has some superficial similarities to Goldstein's VARIABLE MARGINS MODEL. See §A for a comparison of the models. In addition, Goldstein considers two objections to his own VARIABLE MARGINS MODEL, which, if successful, are also objections to FULL KNOWLEDGE + MFE MODEL. See §C for a discussion of those two objections.

5.3 | Comparing knowledge and physical safety

In the previous section, I showed that the possibility of full safety is consistent with the thought that there's a required margin-for-safety, and then argued that we have reasons to think that it's indeed possible. Then, in this section, I showed how to use the same idea to get a model of knowledge where margin-for-error is true and yet full knowledge is possible. In the next section, I will show how to use this model to respond to the skeptical arguments from §3. But first, I want to address a worry; maybe there are some important disanalogies between knowledge and safety that undermine my claim that the two can be treated in the same way.

5.3.1 | Gradibility

Safety is *gradable*, knowledge isn't. Alice can be *safe*, and she can also be *safer than* Bob. However, while Alice can *know* that *p*, she can't *knower* that *p* than Bob. There are maybe relevant related notions (for example, she can have stronger evidence for *p*), but knowledge itself is not gradable.

On its own, this is not yet an objection. But, it's very natural to hear "safely safe" as "safer than safe," which can explain why, to be safely safe, Sarah needs to be further away from the cliff edge. However, a similar move can't be made about knowledge. When we say that Isaac knows that he knows that *p*, we are not saying that he knows that *p more*. We are not saying that his belief that *p* is safer. We are saying that he has another belief — the belief that he knows that *p* — and that this belief is safe.

To an extent, I want to concede this point. There are potential skeptical worries about full knowledge that don't apply to full safety. For example, one might worry that full knowledge requires infinite representational capacity: to fully know that *p*, one must mentally represent infinitely many propositions: *p*, "I know that *p*," "I know that I know that *p*," and so on. And maybe humans just can't do it.[17]

However, this point can't be used to defend skepticism that is motivated by the arguments from §3. The motivation for those arguments was that each iteration of knowledge requires further margin *for the original proposition*, such that infinitely many iterations are impossible. And this is an exact parallel of the physical safety case. Merely pointing out that those are different beliefs that should be safe doesn't pose any threat to the analogy, as it's also true in the physical safety case. "Sarah is safe," "Sarah is safely safe," and so on are all different propositions, but we still saw that full safety is possible.

5.3.2 | Knowledge and interests

The argument for the possibility of full safety was based on the observation that our interest in being further away from the cliff edge runs out at some point, and the claim that when it runs out, all the iterations of safety has been achieved. This interest is best understood as a kind of practical interest, our interest in not falling to our death.

To make a similar argument about knowledge, we need a similar notion: something like an *epistemic interest*. But, even if there is such a thing, it's a bit peculiar. Sometimes, we have a practical interest in knowing things, such as in knowing whether a medication is effective. However,

[17] See Greco (2023, pp. 130-132), for a relevant discussion.

this can't be the relevant sense of *epistemic interests*. There are things that matter nothing to me, like the number of blades of grass in my lawn. And still the margin-for-error principle applies to my knowledge of this number. My knowledge of the number of blades of grass is limited by some margin for error, even though the answer matters nothing to me. So the relevant sense of *epistemic interests* must be different.

It's not clear what epistemic interests might be. Nonetheless, if we accept that knowledge is governed by the margin-for-error principle, and that it's relevantly similar to physical safety, there must be something that plays the role that interests play for physical safety. In the same way that the margin-for-safety partly depends on the practical interests, the margin-for-error must partly depend on a similar notion. We might as well call it *epistemic interests*, while keeping in mind that it's a peculiar notion.

One possibility is that every belief comes with something like an interest in not getting it wrong, even if the believer doesn't care. Maybe it's a *constitutive feature* of beliefs. If an agent forms a belief at all, it comes with an epistemic interest in being right.

Even with only a minimal grasp of what epistemic interests are, we can repeat the argument for the possibility of full safety, this time for knowledge. Knowledge matters. It makes a difference whether Isaac knows that Sarah is at least ten feet from the cliff edge, and if he knows, it matters how many orders of knowledge he has. In addition, at some point further epistemic safety doesn't matter. If Sarah is one mile from the edge, being further away doesn't increase the safety of Isaac's belief that she is at least ten feet from the edge. His epistemic interest runs out. And at this point, he achieved all the iterations of knowledge, and he fully knows that she is at least ten feet from the edge.

As in the case of safety, there are two alternatives: one is that our epistemic interests never run out, and one is that they do run out, but we still don't have full knowledge at this point. And both seem problematic, for the same reasons the corresponding views about safety seem problematic.

The first option, that our epistemic interests never run out, implies that Issac, insofar as he has the belief that Sarah is on the safe side, always has some unsatisfied epistemic interest. Even when Sarah is a mile away from the edge, there's still something to be gained, epistemically speaking, in her being further away. And this sounds false.[18]

The other option is that there's a disconnect between knowledge and our epistemic interests. If Sarah is one mile away from the edge, Isaac may only, for example, 100-know that she is ten feet from the edge. She needs to be a foot further away from the edge for him to 101-know it. And this is even though he has no more epistemic interests. This make the margin-for-error seems mysterious. When she is closed to the edge, the required margin-for-error stems from Isaac's epistemic interests. But once Sarah is far enough, something else explains the required margin-for-error.

6 | RESPONDING TO THE SKEPTICAL ARGUMENTS

It's time to return to the skeptical arguments from §3. The tools we've developed will help us see where those arguments go wrong.

[18] This seems to be the view in Carter and Hawthorne (2024), although they are not committal on this question.

6.1 | Responding to the argument from higher-order knowledge of MFE$_\varepsilon$

In §3.1 we saw an argument to the conclusion that Mr Magoo has only trivial full knowledge of the height of the tree. As stated, the full knowledge skeptic can't use this argument: it uses the premise that Mr Magoo has nontrivial full knowledge of his margin. But if nontrivial full knowledge of the height of the tree is impossible, then it's hard to see how nontrivial full knowledge of the margin can be possible.

However, like many skeptical arguments, the argument still poses a challenge for the non-skeptic. The non-skeptic might worry that if full knowledge is possible at all, then it's possible for Mr Magoo to have full knowledge of MFE$_\epsilon$ for some $\epsilon > 0$. And if he has such full knowledge, then he doesn't have nontrivial full knowledge of the height of the tree.

To answer this argument, we have to deny the possibility of nontrivial full knowledge of the margin. This is, I think, a standard technique in responding to skeptical arguments. As a general rule, we ought to prefer theories that give us more knowledge of ordinary things, even if they entail some gaps in our knowledge. For an example, consider the following point from Williamson:

> [S]ceptical arguments may go wrong by assuming too much knowledge; by sacrificing something in self-knowledge to the sceptic, we stand to gain far more in knowledge of the world [...] Once we relax our claims to self-knowledge, we strengthen our claim to knowledge of the external world. (Williamson, 2002, pp. 164, 183)

Williamson says that we can be led to skepticism about the external world if we assume that we have too much knowledge of what our evidence is. If we allow for some ignorance of what our evidence is, we can resist skepticism about the external world.

The same response can be used on behalf of full knowledge. We are led to skepticism about full knowledge by assuming too much knowledge: full knowledge of the margin. By denying the possibility of full knowledge of the margin, we can resist full knowledge skepticism of the external world.

We can get some insight into why full knowledge of MFE$_\epsilon$ is impossible for any $\epsilon > 0$ by comparing it to the physical safety case. In the physical safety case, MFS$_{10}$ (for example) is true regardless of the position of the cliff edge and Sarah: Sarah is only safe if she is at least ten feet from the edge. However, it's false that MFS$_{10}$ is safely true.

On a first look, it might seem surprising. After all, it's a hypothetical case, and nothing stops us from assuming that MFS$_{10}$ is not only true, but extremely modally robust: we can assume that nothing can change it. But, upon reflection, it turns out that in order to assume that it's safely true, we need to make a surprising assumption.

For MFS$_{10}$ to be safely true, it must be true that $S(S(X > Y) \leftrightarrow (X > Y + 10))$. And, assuming safety is closed under deduction, it means that $S(S(X > Y)) \leftrightarrow S(X > Y + 10)$. And because we assume that MFS$_{10}$ is true, it means that $S(S(X > Y)) \leftrightarrow (X > Y + 20)$. Which means that Sarah needs to be least twenty feet from the cliff edge to be safely safe. And we already know that this is false.

So it turns out that in order to stipulate that MFS$_{10}$ is safely true, we must make some surprising and unnatural stipulations. It's not enough to assume that MFS$_{10}$ is modally robust. Safety, as we saw in §4.4, is not simply a modal stability operator.

Similarly, it turns out that the stipulation that Mr Magoo knows MFE$_1$ has other effects. What it means is that it's harder to make MFE$_1$ known, in the same way that it's harder to make MFS$_{10}$ safe.

6.2 | Responding to the modal argument

The argument we saw in §3.2 was based on the idea that knowledge is a modal stability operator: to know something, it must be true in all the non-excluded nearby worlds. The argument relied on the following strengthening of MFE$_\epsilon$:

> MFE$_\epsilon$-CHAINING: For every x, if w is a world where $T = x$, then there's a non-excluded nearby world w' where MFE$_\epsilon$-CHAINING is true and $T = x + \epsilon$.

As I already said, it's not obvious how to motivate MFE$_\epsilon$-CHAINING. We considered the suggestion that while the height of the tree varies at nearby worlds, Mr Magoo margin doesn't. And therefore there's a nearby world where the height of the tree is different, but the margin is the same.

However, this motivation doesn't work. While the height of trees varies considerably, we can imagine that Mr Magoo instead glances at a nearby titanium statue, whose height was very deliberately set to be exactly 100 inches, without Mr Magoo knowing it. We can make the height of the statue as modally robust as possible. We can also make Mr Magoo's margin very modally fragile, where in a lot of nearby worlds his margin is different (without him being any wiser). But none of this seem to change what Mr Magoo can know, 2-know, or n-know about the height of the statue.

The lesson is that, just like physical safety, knowledge isn't a modal stability operator. What Mr Magoo can know doesn't depend on which worlds are nearby. And this fact, while perhaps surprising, doesn't mean that the margin-for-error principle is false. Think again about the case **Truck** from §2.1. We can imagine that the exact heights of the bridge and the car are extremely modally robust, and in all nearby worlds they are the same; so in all nearby worlds the car can pass under bridge. But it's still not the case that Alice knew that the truck could pass under the bridge. She could have easily been wrong. So we can to accept MFE$_\epsilon$ while rejecting the motivation for MFE$_\epsilon$-CHAINING.

6.3 | The secret argument: Responding to the argument from models of epistemic logic

In addition to the two arguments from §3, I think that there's a third, less obvious, skeptical argument. It's based on the observation that in simple models, MFE leads to full knowledge skepticism. And if it's true in simple models, it gives us reasons to think that it's generally true, even if it's false in more complex models.[19]

As a general rule, it's a good argument form. We can learn a lot from studying simple models. And if the only way to accommodate some potential feature of knowledge is to use very complex models, then it's a reason to be suspicious of this feature. However, I don't think that the simplicity of PURE MFE MODEL gives any support for full knowledge skepticism. Mr Magoo has nontrivial full knowledge in it: he has nontrivial full knowledge of his own margin. Moreover, he has exact full knowledge of the margin. And no one thinks that nontrivial full knowledge of the margin is possible while nontrivial full knowledge of the height of the tree is not.

This is not a small problem. Any model where Mr Magoo has only trivial full knowledge must be much more complicated than the models we've considered. To really model full

[19] See Williamson (2024) for a defense of this form of argument.

knowledge skepticism, it's not enough to have nodes for any possible margin. For example, adding a second-order margin-for-error, a margin-for-error for the margin-for-error, is not enough, because Mr Magoo shouldn't have full knowledge of that margin. And similarly for a third-order margin-for-error, and any finite order of margin-for-error.

In addition, in the models we considered, we accepted a principle called "appearances centering:" the apparent height of the tree is always in the middle of range of what Mr Magoo knows about the height of the tree. And this principle was true in all the nodes, and therefore fully known in all the nodes. But if full knowledge skepticism is true, Mr Magoo shouldn't have full knowledge of that as well. After all, it's possible that he is clairvoyant, and he can know things about the height of the tree despite misleading appearances.

If simplicity considerations count in favor of anything, they count in favor of the possibility of full knowledge. By granting Mr Magoo full knowledge, at least of some things, we can consider much simpler models.

7 | CONCLUSIONS

Here is the resulting picture. I fully know some things. I fully know that the tallest sequoia tree doesn't reach the moon. I fully know that there are more than three people in New York City, and less than ten billion of them. But I don't fully know everything that I know. I know that my neighbor's kid is taller than three feet, but I don't fully know it; he is too close to three feet for me know that I know it. If I believed that I know it, this belief would be true, but due to luck, in a way that's inconsistent with knowledge.

This seems to me to be a very plausible picture of knowledge. Theoretical commitments might force us to conclude that it's false, but only on the basis of reasonably strong arguments. I've argued the arguments on offer don't give us reasons to deny this picture.[20]

ACKNOWLEDGEMENTS

Open Access funding enabled and organized by MIT Hybrid 2025.

ORCID

Yonathan Fiat https://orcid.org/0000-0003-3601-6373

REFERENCES

Bird, A., & Pettigrew, R. (2021). Internalism, externalism, and the KK principle. *Erkenntnis*, *86*(6), 1713–1732.

Carter, S., & Hawthorne, J. (2024). Dogmatism and inquiry. *Mind*, 651–676.

Cohen, S., & Comesaña, J. (2013). Williamson on gettier cases and epistemic logic. *Inquiry*, *56*(1), 15–29.

Das, N., & Salow, B. (2018). Transparency and the KK principle. *Noûs*, *52*(1), 3–23.

Dorr, C., Goodman, J., & Hawthorne, J. (2014). Knowing against the odds. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *170*(2), 277–287.

Dorst, K. (2019). Abominable KK failures. *Mind*, *128*(512), 1227–1259.

Fraser, R. E. (2022). KK failures are not abominable. *Mind*, *131*(522), 575–584.

[20] I'd like to thank Roger White, Jack Spencer, Kevin Dorst, Richard Pettigrew, Alexander Bird, Branden Fitelson, Jonathan Dixon, Jeremy Goodman, Caspar Hare, Alex Byrne, and the audiences at the 16th *Annual Cambridge Graduate Conference in Logic and the Philosophy of Mathematics* (2023), the *Full Knowledge and Common Knowledge* symposium at the 121st meeting of the APA Central Division (2024), the 2022/23 MIT dissertation seminar, and the 2023/24 MIT MATTI work-in-progress group.

Goldstein, S. (2024). *Iterated knowledge*. Oxford University Press.
Goldstein, S. (forthcoming). Omega knowledge matters. *Oxford Studies in Epistemology*.
Goldstein, S., & Hawthorne, J. (2024). KK is wrong because we say so. *Mind*, 1–27.
Goodman, J. (2013). Inexact knowledge without improbable knowing. *Inquiry*, *56*(1), 30–53.
Goodman, J., & Salow, B. (2018). Taking a chance on KK. *Philosophical Studies*, *175*(1), 183–196.
Greco, D. (2014). Could KK be ok? *The Journal of Philosophy*, *111*(4), 169–197.
Greco, D. (2015). Iteration and fragmentation. *Philosophy and Phenomenological Research*, *91*(3), 656–673.
Greco, D. (2016). Safety, explanation, iteration. *Philosophical Issues*, *26*(1), 187–208.
Greco, D. (2023). Inter-level coherence. *Idealization in epistemology: A modest modeling approach* (pp. 117–142). Oxford University Press.
Liu, S. (2020). (Un)Knowability and knowledge iteration. *Analysis*, *80*(3), 474–486.
San, W. K. (2023). KK, knowledge, knowability. *Mind*, 605–630.
Sosa, E. (1999). How must knowledge be modally related to what is known? *Philosophical Topics*, *26*(1/2), 373–384.
Stalnaker, R. (2009). On hawthorne and magidor on assertion, context, and epistemic accessibility. *Mind*, *118*(470), 399–409.
Stalnaker, R. (2015). Luminosity and the KK thesis. In S. C. Goldberg (Ed.), *Externalism, self-knowledge, and Skepticism* (pp. 19–40). Cambridge University Press.
Weatherson, B. (2013). Margins and errors. *Inquiry*, *56*(1), 63–76.
Williamson, T. (2002). *Knowledge and its limits*. Oxford University Press.
Williamson, T. (2009). Probability and danger. Amherst lecture in philosophy.
Williamson, T. (2011). Improbable knowing. In T. Dougherty (Ed.), *Evidentialism and its discontents* (pp. 147–164). Oxford University Press.
Williamson, T. (2013a). Gettier cases in epistemic logic. *Inquiry*, *56*(1), 1–14.
Williamson, T. (2013b). Response to Cohen, Comesaña, Goodman, Nagel, and Weatherson on Gettier Cases in Epistemic Logic. *Inquiry*, *56*(1), 77–96.
Williamson, T. (2014). Very improbable knowing. *Erkenntnis*, *79*(5), 971–999.
Williamson, T. (2021). The KK principle and rotational symmetry. *Analytic Philosophy*, *62*(2), 107–124.
Williamson, T. (2024). Overfitting and degrees of freedom. *Overfitting and heuristics in philosophy* (pp. 51–104). Oxford University Press.

How to cite this article: Fiat, Y. (2025). Knowing to infinity: Full knowledge and the margin-for-error principle. *Philosophy and Phenomenological Research*, 1–31. https://doi.org/10.1111/phpr.70001

APPENDIX A: A visual comparison of knowledge models

The discussion of how knowledge works and iterates is heavily informed by the construction and analysis of Kripke models for knowledge. It's easy to get lost in the variety of models in the literature. And it's often hard to interpret the formal definitions of those models.

Williamson (2013b) suggests an easy-to-understand way to visualize such models. Williamson intended it mostly for first-order knowledge, but his method can be generalized to higher-order knowledge as well.

In this appendix, I'll quickly survey a wide variety of models that has been used to answer questions about knowledge iterations, and present their *Williamson-graphs*. I'll explain how to read the visualization in the next section.

A.1 | Full knowledge skepticism models

A.1.1 | PURE MFE MODEL: Williamson (2013a)

We've encountered this model in §2.2.

PURE MFE MODEL

Let $\epsilon > 0$.

The set of nodes: $\{(a, r) : a, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2. $|a' - r'| \leq |a - r| + \epsilon$

Figure A1a is the Williamson-graph of this model. To understand how to read it, let a be Mr Magoo's best estimate of the height of the tree (in inches). Suppose that he underestimates the height of the tree by x inches, so the tree is $a + x$ inches tall. For every value of x, we can ask what is the strongest thing Mr Magoo knows of the form "the tree is at most $a + y$ inches tall."

Different models give different answers to this question. The Williamson-graph of a model shows the answer; for every value of x, it tells us the value of y according to the model. For example, Figure A1a shows us that according to PURE MFE MODEL, if Mr Magoo underestimates the height of the tree by x inches, the strongest thing he knows of the form "the tree is at most $a + y$ inches tall" is "the tree is at most $a + (x + \epsilon)$ inches tall."

The dark dashed line in the figure is the $y = x$ line. Because Mr Magoo can only know true things, every model must always satisfy $y \geq x$. It means that the Williamson-graph of any model is always above (or on) the $y = x$ line. We can see in the figure that MFE_ϵ is true in PURE MFE MODEL, because its Williamson-graph is always at least ϵ above the $y = x$ line.

Figure A1b shows the Williamson-graph of PURE MFE MODEL for different orders of knowledge. The idea it exactly the same: assuming that Mr Magoo's best estimate is that the tree is a inches

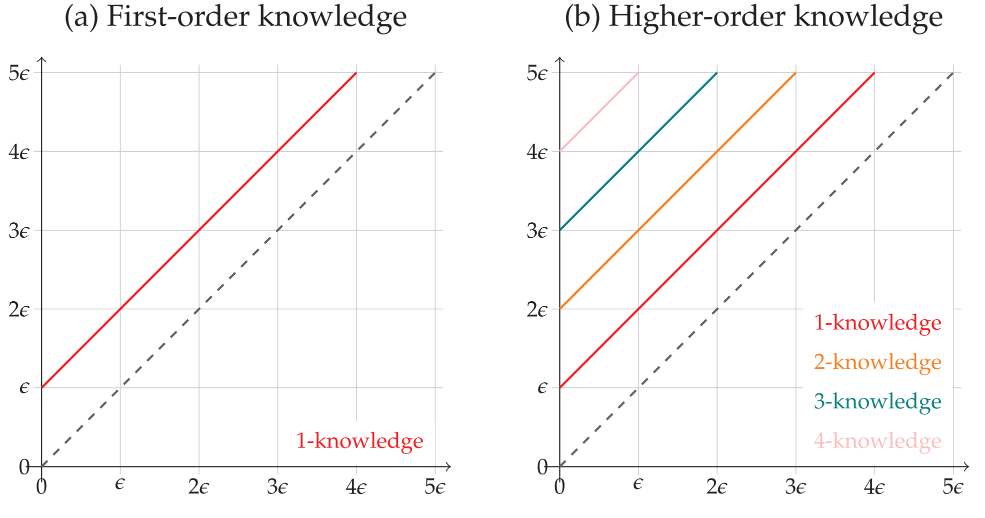

FIGURE A1 PURE MFE MODEL.

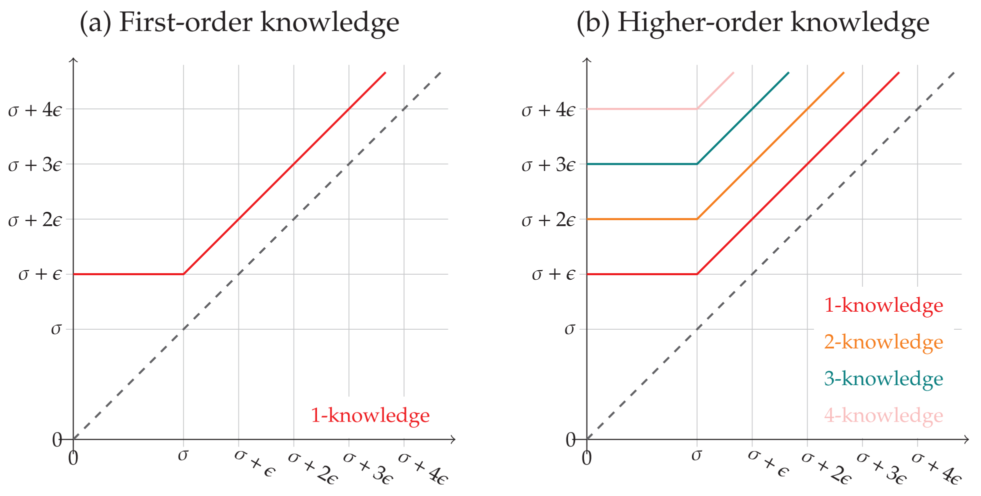

FIGURE A2 SAFE HAVENS MODEL.

tall, and that tree is $a + x$ inches tall, we can ask what is the strongest thing Mr Magoo n-knows of the form "the tree is at most $a + y$ inches tall." We can see in the figure that HO-MFE$_\epsilon$ (§3) is true in PURE MFE MODEL. It's represented by the fact that the gaps between the lines are of the same size.

A.1.2 | SAFE HAVENS MODEL: Goodman (2013)

Goodman (2013) presents a different model. In this model, MFE$_\epsilon$ is true, but it's not the only restriction on Mr Magoo's knowledge. One benefit of this model is that it clearly shows that the margin-for-error in an anti-luck principle, and that there could be additional restrictions on knowledge.

In Goodman's model, Mr Magoo's knowledge has what Goodman calls a *safe haven*: as long as Mr Magoo's estimate is close enough to the actual height of the tree, what he knows doesn't depend on the exact height of tree.[21]

SAFE HAVENS MODEL

Let $\epsilon > 0, \sigma > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2. $|a' - r'| \leq \max\{|a - r|, \sigma\} + \epsilon$

As we can see in Figure A2a, MFE$_\epsilon$ is true in SAFE HAVENS MODEL. The gap between the Williamson-graph and the $y = x$ line is always at least ϵ. And we can see in Figure A2b that HO-MFE$_\epsilon$ is also true in SAFE HAVENS MODEL, because the gaps between the lines are equal. Table A1 shows some differences between SAFE HAVENS MODEL and PURE MFE MODEL.

[21] See Goodman (2013) for some arguments to the conclusion that we should prefer this model over PURE MFE MODEL. See Williamson (2013b) for some responses.

TABLE A1 Comparison of PURE MFE MODEL and SAFE HAVENS MODEL.

A	T	What Mr Magoo knows in PURE MFE MODEL ($\epsilon = 5$)	What Mr Magoo knows in SAFE HAVENS MODEL ($\epsilon = 5, \sigma = 25$)
100	100	$95 \leq T \leq 105$	$70 \leq T \leq 130$
100	99	$94 \leq T \leq 106$	$70 \leq T \leq 130$
99	100	$93 \leq T \leq 105$	$69 \leq T \leq 129$
130	100	$95 \leq T \leq 165$	$95 \leq T \leq 165$

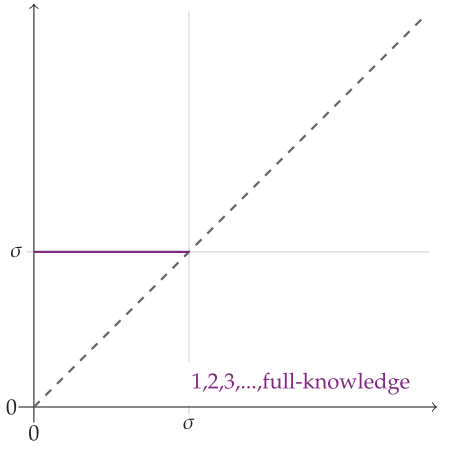

FIGURE A3 BASIC KK MODEL. It's only defined in *normal cases*: cases where $|a - r| \leq \sigma$. In KK models, all the orders of knowledge coincide, which means that they also coincide with full knowledge.

A.2 | KK models

KK models are models that validate the KK principle: "if a knows that p, then she knows that she knows that p." The KK principle is defended by, for example, Cohen and Comesaña (2013); Dorst (2019); Greco (2014, 2016), and Stalnaker (2009, 2015). In my understanding of the literature, all KK models are built on the following BASIC KK MODEL:

> **BASIC KK MODEL**
>
> Let $\sigma > 0$.
>
> The set of nodes: $\{(a, r) : a \geq 0, r \geq 0, |a - r| \leq \sigma\}$.
>
> (a, r) can access (a', r') iff $a' = a$.

This model is only defined for cases where Mr Magoo's estimate of the height of the tree is reasonably close to the actual height of the tree. Such cases are often called *normal cases*. The model doesn't say what happens in *abnormal* cases.

MFE_ϵ is false in this model for every $\epsilon > 0$. We can see this in Figure A3; the Williamson-graph of the model touches the $y = x$ line at $x = \sigma$. It means that according to this basic model, cases like the ones discussed in §2.1 can be cases of knowledge.

There are at least two ways to extend this model to abnormal cases while preserving the KK principle. The first is shown in the following model:

EXTENDED KK MODEL: HIERARCHY

Let $0 < \sigma_1 < \sigma_2 < \sigma_3 < ....$

The set of nodes: $\{(a,r) : a \geq 0, r \geq 0\}$.

(a,r) can access (a',r') iff:

1. $a' = a$; and
2.

If $\|a - r\|$ is in	Then
$[0, \sigma_1]$	$\|a' - r'\| \leq \sigma_1$
$(\sigma_1, \sigma_2]$	$\|a' - r'\| \leq \sigma_2$
$(\sigma_2, \sigma_3]$	$\|a' - r'\| \leq \sigma_3$
...	...

The idea behind it is that there's a hierarchy of normality conditions, such that when one fails, there's a fallback. Basically, different normality conditions behave a bit like different sources of information.

For example, suppose that the tree is 187 inches tall. Alice tells Mr Magoo that the tree is at most 200 inches tall, and Bob tells him that the tree is at most 250 inches tall. Mr Magoo doesn't know it, but Alice just pulled the number out of the top of her head, while Bob actually measured the height of the tree. In this case, Mr Magoo doesn't know that the tree is at most 200 inches tall, even though he believes it, and it's true. But he does know that the tree is at least 250 inches tall. Even though Alice is an unreliable source of information, Mr Magoo has a *fallback option* of relying on Bob.

According to EXTENDED KK MODEL: HIERARCHY, something similar happens when Mr Magoo underestimates the height of the tree by more than σ_1. His 'first source' of visual information, characterized by σ_1, failed. But he has a *fallback option*, a 'second source' of visual information, characterized by σ_2, to rely on.

As we can see in Figure A4a, the Williamson-graph of this model is discontinuous. It means that small differences in Mr Magoo's estimate can cause huge differences in what he knows. For

(a) EXTENDED KK MODEL: HIERARCHY (b) EXTENDED KK MODEL: CONTINUUM

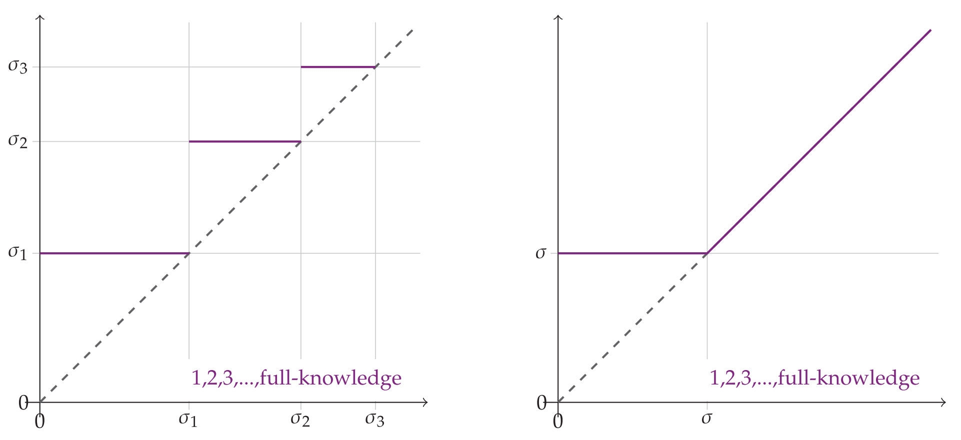

FIGURE A4 Extended KK models.

example, suppose that Mr Magoo and Ms Magoo both glance at the tree. The quality of their eyesight is equally poor. Mr Magoo estimates that the tree is 700 inches. Ms Magoo estimates that the tree is 700.1 inches. According to this model, it's possible that Mr Magoo knows that the tress is at most 700 inches, but Ms Magoo only knows that it's at most 800 inches.

Cohen and Comesaña (2013) suggest a different way to extend BASIC KK MODEL. We can think about this second way as the limit of adding more and more normality conditions to the hierarchy.

EXTENDED KK MODEL: CONTINUUM

Let $\sigma > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2. $|a' - r'| \leq \max\{\sigma, |a - r|\}$

Just like in BASIC KK MODEL and EXTENDED KK MODEL: HIERARCHY, MFE_ϵ is false for every $\epsilon > 0$. But as we can see in Figure A4b, is fails in a much more dramatic way. All the cases where conditions are abnormal are counterexamples to MFE_ϵ for every ϵ. We can see in Figure A4b by noting that the Williamson-graph overlaps with the $y = x$ line for all the values $x \geq \sigma$. It means that knowledge in cases like the ones in §2.1 is not only be possible, but is relatively common.

On the other hand, the Williamson-graph of this model has no discontinuity points. It means that if Mr and Ms Magoo's estimates are close to each other, they know similar things.

A.3 | Goldstein (2024): Models of full knowledge without KK and without margin-for-error

Goldstein (2024) presents four models that allow for full knowledge. His discussion of the models is very rich and detailed, but I can only present the models and offer a brief commentary here. In all of Goldstein's models, MFE_ϵ is false for every $\epsilon > 0$. In addition, the Williamson-graphs of all his models have discontinuity points, just as in EXTENDED KK MODEL: HIERARCHY.

A.3.1 | REFLECTIVE LUMINOUSITY MODEL

Reflective Luminousity is the claim that if a 2-knows that p, then a 3-knows that p. It implies that 2-knowledge is identical to full knowledge.

REFLECTIVE LUMINOUSITY MODEL

Let $\sigma > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$. (a, r) can access (a', r') iff:

1. $a' = a$; and
2.

If $\lvert a - r\rvert$ is in	Then
$[0, \sigma]$	$\lvert a' - r'\rvert \leq \lvert a - r\rvert + \sigma$
$(\sigma, 2\sigma]$	$\lvert a' - r'\rvert \leq 2\sigma$
$(2\sigma, 3\sigma]$	$\lvert a' - r'\rvert \leq \lvert a - r\rvert + \sigma$
$(3\sigma, 4\sigma]$	$\lvert a' - r'\rvert \leq 4\sigma$
...	...

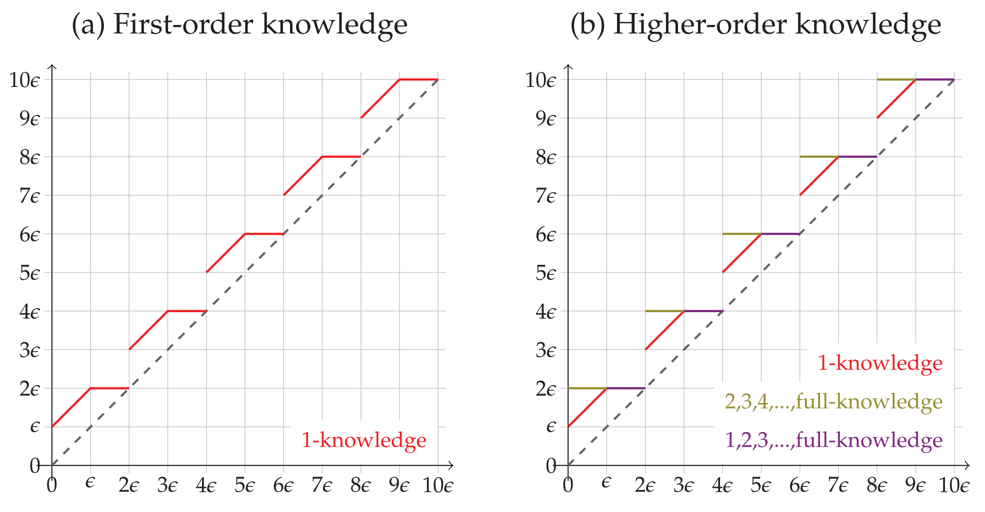

FIGURE A5 REFLECTIVE LUMINOUSITY MODEL. In this model, second-order knowledge coincides with all the higher-orders of knowledge.

As we can see in Figure A5, this model shares some features with to the EXTENDED KK MODEL: HIERARCHY (Figure A4a).

A.3.2 | FRAGILITY MODEL

Fragility is the following principle: "If you know that p, then for all you know, you fully know that p." The details of the principle are not relevant to us here, except that it allows Goldstein to build a model with full knowledge. I made one small modification to Goldstein's model: I used two parameters instead of one, because they play two distinct roles in the model.

FRAGILITY MODEL

Let $\sigma > 0, \epsilon > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2.

If $\lvert a - r \rvert$ is in	Then
$[0, \sigma]$	$\lvert a' - r' \rvert \leq \sigma$
(σ, ∞)	$\lvert a' - r' \rvert \leq \lvert a - r \rvert + \epsilon$

This model is an extension of BASIC KK MODEL (Figure A3), which means that when conditions are normal, KK is true. But in abnormal conditions, KK fails. As we can see in Figure A6, when conditions are abnormal, the model agrees with the skeptical PURE MFE MODEL and SAFE HAVENS MODEL (§2.2, Figure A1, Figure A2).

A.3.3 | REFLECTIVE LUMINOUSITY + FRAGILITY MODEL

Reflective Luminousity is consistent with *Fragility*. Goldstein presents a third model to show their consistency. Figure A7 show its Williamson-graph.

REFLECTIVE LUMINOUSITY + FRAGILITY MODEL

Let $\sigma > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2.

If $\lvert a - r \rvert$ is in	Then
$[0, \sigma]$	$\lvert a' - r' \rvert \leq \sigma$
$(\sigma, 2\sigma]$	$\lvert a' - r' \rvert \leq \lvert a - r \rvert + \sigma$
$(2\sigma, 3\sigma]$	$\lvert a' - r' \rvert \leq 3\sigma$
$(3\sigma, 4\sigma]$	$\lvert a' - r' \rvert \leq \lvert a - r \rvert + +\sigma$
...	...

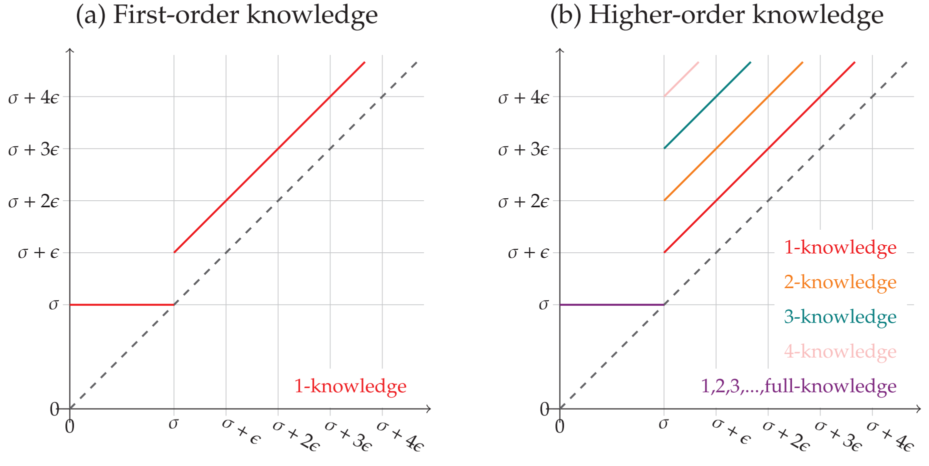

FIGURE A6 FRAGILITY MODEL. In this model, full knowledge is only possible when conditions are normal. And when conditions are normal, all the orders of knowledge coincide.

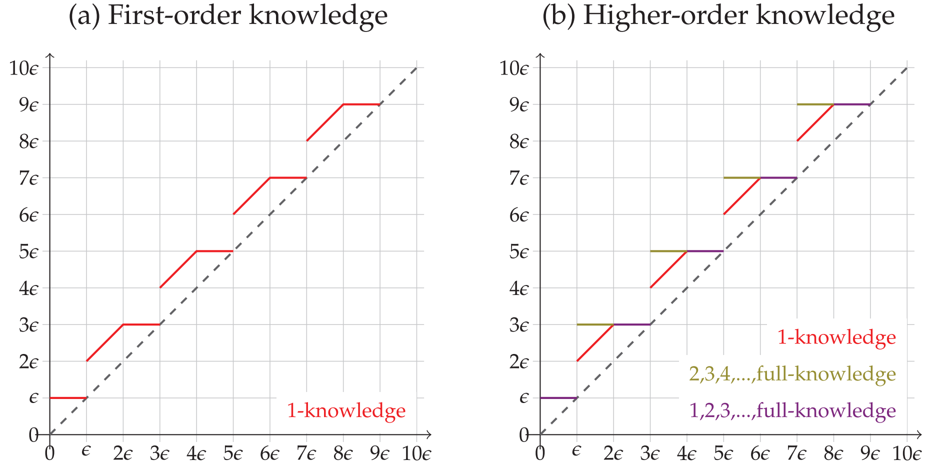

FIGURE A7 REFLECTIVE LUMINOUSITY + FRAGILITY MODEL. As in REFLECTIVE LUMINOUSITY MODEL, second-order knowledge coincides with all the higher-orders of knowledge.

Just like FRAGILITY MODEL (Figure A6), this model is an extension of BASIC KK MODEL, where KK fails in abnormal conditions. And like REFLECTIVE LUMINOUSITY MODEL (Figure A5), this model is an instance of EXTENDED KK MODEL: HIERARCHY (Figure A4a) if we disregard first-order knowledge.

A.3.4 | VARIABLE MARGINS MODEL

Goldstein's VARIABLE MARGINS MODEL is of the greatest interest to us, as it's superficially similar to FULL KNOWLEDGE + MFE MODEL.

VARIABLE MARGINS MODEL

Let $\sigma > 0$.

The set of nodes: $\{(a, r) : a \geq 0, r \geq 0\}$.

(a, r) can access (a', r') iff:

1. $a' = a$; and
2.

If $\lvert a - r \rvert$ is in	Then
$[0, 2\sigma]$	$\lvert a' - r' \rvert \leq \frac{1}{2} \lvert a - r \rvert + \sigma$
$(2\sigma, 4\sigma]$	$\lvert a' - r' \rvert \leq \frac{1}{2} \lvert a - r \rvert + 2\sigma$
$(4\sigma, 6\sigma]$	$\lvert a' - r' \rvert \leq \frac{1}{2} \lvert a - r \rvert + 3\sigma$
$(6\sigma, 8\sigma]$	$\lvert a' - r' \rvert \leq \frac{1}{2} \lvert a - r \rvert + 4\sigma$
...	...

Like REFLECTIVE LUMINOUSITY MODEL and REFLECTIVE LUMINOUSITY + FRAGILITY MODEL (Figures A5, A7), this model is somewhat similar to EXTENDED KK MODEL: HIERARCHY (Figure A4a). Unlike the two *Reflective Luminousity* models, all finite orders of knowledge are distinct from full knowledge. But as n goes to infinity, n-knowledge becomes more and more similar to full knowledge, and full knowledge behaves as in EXTENDED KK MODEL: HIERARCHY. Figure A8 shows the Williamson-graph of VARIABLE MARGINS MODEL.

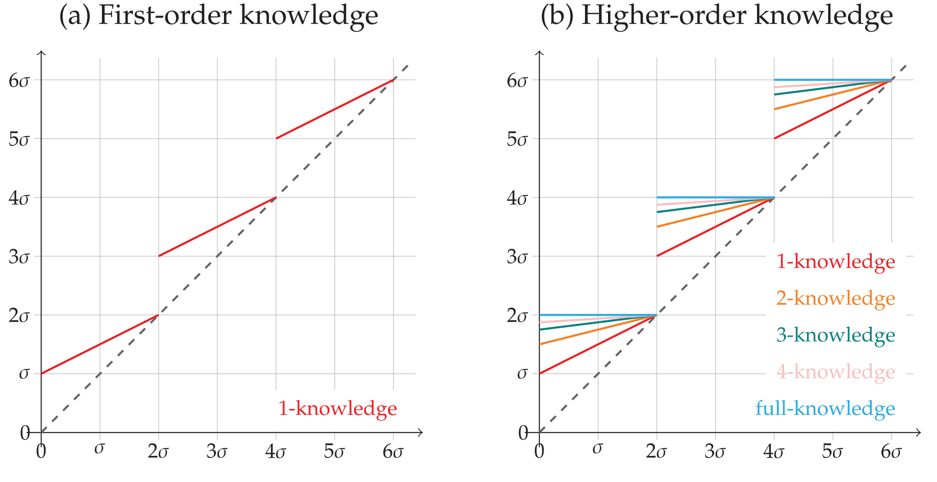

FIGURE A8 VARIABLE MARGINS MODEL. In this model, full knowledge is distinct from n-knowledge for every n.

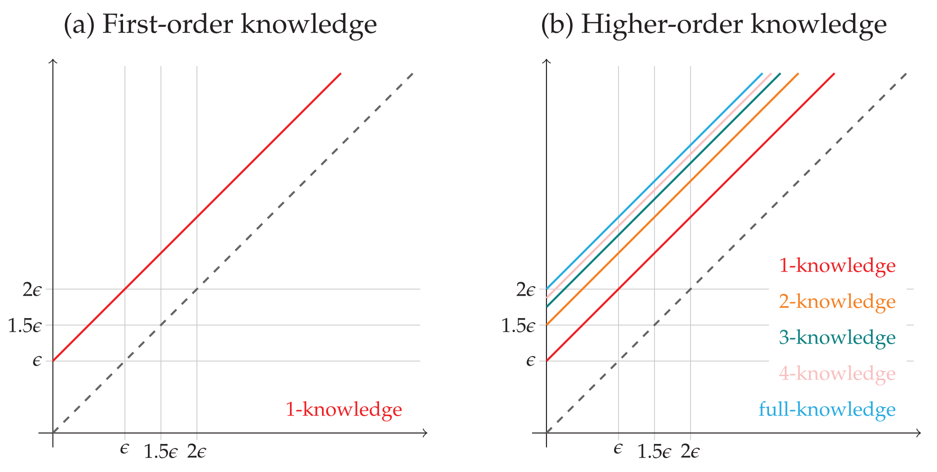

FIGURE A9 FULL KNOWLEDGE + MFE MODEL.

A.4 | Full knowledge without KK and with the margin-for-error principle

In §5.1 I presented a new model: full knowledge is possible in it, and MFE_ϵ is true.

A.4.1 | FULL KNOWLEDGE + MFE MODEL

FULL KNOWLEDGE + MFE MODEL

The set of nodes: $\{(\epsilon, a, r) : a, r, \epsilon \geq 0\}$.

(ϵ, a, r) can access (ϵ', a', r') iff:

1. $a' = a$; and
2. $\frac{\epsilon}{2} \leq \epsilon' \leq \epsilon$; and
3. $|a' - r'| \leq |a - r| + \epsilon - \epsilon'$

As we can see in Figure A9, in this model full knowledge is distinct from n-knowledge for every n, but it's still possible. Moreover, it behaves just like knowledge, simply with a wider margin.

A.4.2 | FULL KNOWLEDGE + MFE + SAFE HAVENS MODEL

FULL KNOWLEDGE + MFE MODEL is a version of PURE MFE MODEL, but where full knowledge is possible. We can use the same technique to get a version of SAFE HAVENS MODEL where full knowledge is possible.

FULL KNOWLEDGE + MFE + SAFE HAVENS MODEL

Let $\sigma > 0$.

The set of nodes: $\{(\epsilon, a, r) : a, r, \epsilon \geq 0\}$.

(ϵ, a, r) can access (ϵ', a', r') iff:

1. $a' = a$; and
2. $\frac{\epsilon}{2} \leq \epsilon' \leq \epsilon$; and
3. $|a' - r'| \leq \max\{|a - r|, \sigma\} + \epsilon - \epsilon'$

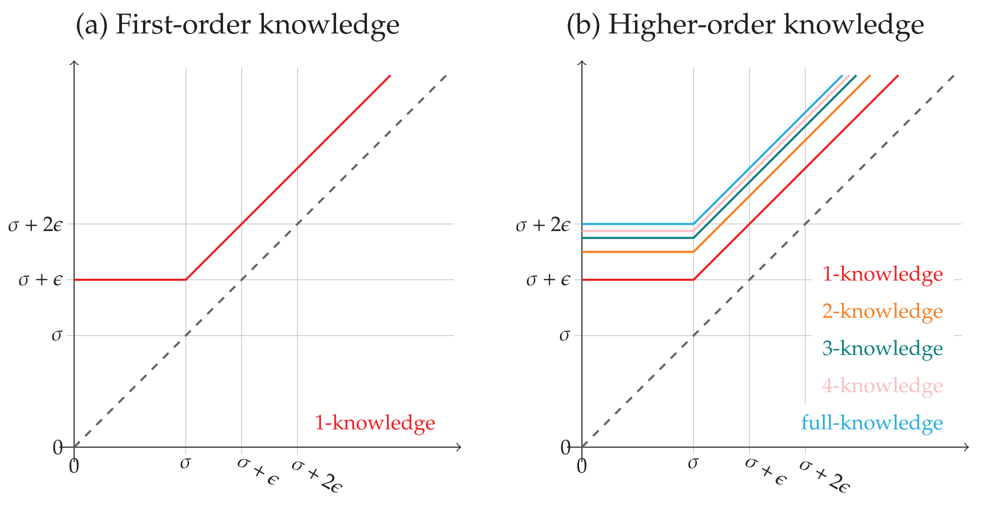

FIGURE A10 FULL KNOWLEDGE + MFE + SAFE HAVENS MODEL.

Just as in FULL KNOWLEDGE + MFE MODEL, in this model full knowledge is distinct from n-knowledge for every n, but it's still possible. Moreover, it behaves just like knowledge, simply with a wider margin (Figure A10).[22]

APPENDIX B: A different look at the full knowledge model

FULL KNOWLEDGE + MFE MODEL (§§5.1,A.4.1) shows how full knowledge is consistent with the margin-for-error principle once we allow ignorance of the margin. However, it also hides some important facts about how it's achieved. It can be useful to compare it to the following variation:

FULL KNOWLEDGE + MFE LEVELS MODEL

Let $\epsilon > 0$. Let $\epsilon_n = \frac{\epsilon}{2^{(n-1)}}$

The set of nodes: $\{(l, a, r) : a, r \geq 0, l \in \mathbb{N}\}$.

(l, a, r) can access (l', a', r') iff:

1. Either ($l' = l$ and $a' = a$ and $|a' - r'| \leq |a - r|$);
2. or ($l' = l + 1$ and $a' = a$ and $|a' - r'| \leq |a - r| + \epsilon_l$)

This model is very similar to FULL KNOWLEDGE + MFE MODEL, but it uses a different parameterization. Instead of defining the nodes of the model explicitly using the margin, it uses a new *level* parameter l. This parameter doesn't directly represent any feature of the world. We can give is a functional interpretation: the margin at a node (l, a, r) is ϵ_l.

One downside of this variation of the model is that the margin only takes values in the discrete set $\left\{\epsilon, \frac{\epsilon}{2}, \frac{\epsilon}{4}, \ldots, \frac{\epsilon}{2^n}, \ldots\right\}$. The original FULL KNOWLEDGE + MFE MODEL shows how we can convert it into a model where the margin can get any positive value.

An upside of this model is that it clearly shows how we can introduce stronger ignorance of the margin. In the original FULL KNOWLEDGE + MFE MODEL, if the margin is ϵ, Mr Magoo fully

[22] FULL KNOWLEDGE + MFE + SAFE HAVENS MODEL is less than fully satisfying because it still depends on a parameter σ, and it allows for no ignorance of it. Ideally, σ should be yet another thing about which the agent can be ignorant.

TABLE C1 Higher-order knowledge in FULL KNOWLEDGE + MFE MODEL, after Mr Magoo learns that the tree is built from one inch tall blocks.

n	What Mr Magoo n-knows at node $(10, 100, 100)$
1	$95 \leq T \leq 105$
2	$93 \leq T \leq 107$
3	$92 \leq T \leq 108$
4	$92 \leq T \leq 108$
5	$92 \leq T \leq 108$
⋮	⋮
Fully	$92 \leq T \leq 108$

knows that the margin is at most ϵ. While this is also true in FULL KNOWLEDGE + MFE LEVELS MODEL, it can be easily fixed. For example, if we define $\epsilon_n = \frac{n^2}{2^n}\epsilon$, then the model still allows for full knowledge. But in addition, if Mr Magoo's margin is ϵ, he doesn't fully know (or even know) that it is at most ϵ.

APPENDIX C: Goldstein's objections to VARIABLE MARGINS MODEL

FULL KNOWLEDGE + MFE MODEL shares some features of Goldstein's VARIABLE MARGINS MODEL (see §A.3.4). The models are different, as can be clearly seen by comparing their Williamson-graphs, and also from the fact that FULL KNOWLEDGE + MFE MODEL is a margin-for-error model (see §5.2) and VARIABLE MARGINS MODEL is not.

However, (Goldstein, 2024, chapter 4) considers two objections to VARIABLE MARGINS MODEL that, if successful, are also objections to FULL KNOWLEDGE + MFE MODEL. I want to briefly explain the objections and answer them.

C.1 | "Digital" knowledge

The height of the tree is a continuous magnitude. But in many cases, we care about knowledge of discrete magnitudes. For example, the number of students in **Cookies** from §2.1. Goldstein worries that models that are based on nodes with different margins can't be used to model such cases. For an example, imagine that after looking at the tree, Mr Magoo learns that the tree is actually built from blocks, each block 1 inch tall.

My answer to this worry is that it doesn't pose any problem. For example, let's consider the node (10, 100, 100) in FULL KNOWLEDGE + MFE MODEL: Mr Magoo's margin is five inches, and the tree is and appears to be one hundred inches tall. Suppose that Mr Magoo learns, and comes to fully know, that the tree is built from individual one inch tall blocks. Then, to figure out what Mr Magoo knows now, we need to remove from the model all the nodes with fractional heights. Table C1 shows what Mr Magoo n-knows for different values of n. Removing the nodes with fractional heights poses no special difficulty.

Perhaps a different way to present this worry is to ask what if Mr Magoo learns, and comes to fully know, that the tree is built from blocks, but this time the blocks are exactly *twenty inches* tall. Then, Mr Magoo comes to know, and fully know, that the tree is exactly one hundred inches tall; any other possible height is inconsistent with what he fully knows on the basis of sight. Does it pose any problem for the margin-for-error principle?

It does not. In this case, the source of Mr Magoo extremely precise knowledge is not merely his vision. It is the combination of his vision with a very strong additional source of knowledge: that the tree is built from blocks, and that each block is exactly twenty inches tall. If it's possible for him to get such information, it's not surprising that he can get very precise knowledge.

Realistically speaking, even if the blocks are created using the most precise scientific instruments, there'll still be some margin-for-error. Even the most precise instruments have some errors, and this is even without mentioning problems of vagueness. So it might be impossible for Mr Magoo to learn that the tree is built from such blocks. Maybe the strongest thing he can learn is that the tree is built from blocks whose height is twenty inches plus-minus a few fractions of an inch. And then his knowledge will be governed by this margin, as it should.

C.2 | Open intervals

In both VARIABLE MARGINS MODEL and FULL KNOWLEDGE + MFE MODEL, the strongest thing Mr Magoo fully knows is always an open interval. For example, in the node $(10, 100, 100)$ of FULL KNOWLEDGE + MFE MODEL, the strongest thing Mr Magoo fully knows is that $90 < T < 110$. And this is not a coincidence: the strongest thing he fully knows always have the form $x < T < y$, and never $x \leq T \leq y$, $x \leq T < y$, or $x < T \leq y$.

Goldstein suggests that this is a problem, and maybe it is for those who, like Goldstein, deny the margin-for-error principle. But I don't think it's a problem for those who accept it. The margin-for-error principle suggests that for *every* n, the strongest thing Mr Magoo is a position to n-know about a continuous magnitude is that it belongs to some open interval. And if this is true, it is not surprising that the same is true for full knowledge.

To deny that is to think that sometimes Mr Magoo is in a position to know $x \leq T \leq y$, but he is not in a position to know that $x < T < y$. Margin-for-error considerations mean that there is some $\epsilon > 0$ such that it is true that $x + \epsilon \leq T \leq y - \epsilon$, so T is neither x nor y. The only reason to think that Mr Magoo isn't in a position to know that $x < T < y$ is to think that this belief isn't safe enough, while $x \leq T \leq y$ is safe enough. And it's very hard to see why this would be the case: for both, there is a margin-for-error of at least ϵ.

Perhaps the source of this worry is that in FULL KNOWLEDGE + MFE MODEL, for every n, the strongest thing that Mr Magoo n-knows about the height of the tree is that it's in some closed interval. But this is a non-essential feature of the model that can be easily fixed by changing the "≤" signs to "<" (special care should be made to ensure that the model is still reflexive).

I didn't define the model this way because it's easier to talk about the model without making this change. In the current form, I can truly say "in the node $(10, 100, 100)$, Mr Magoo doesn't know that the tree is less than 105 inches tall." If I made the change, I would have to say "in the node $(10, 100, 100)$, for every $\delta > 0$, Mr Magoo doesn't know that the tree is less than $105 - \delta$ inches tall." This would have allowed me to make the same arguments, but they would be much harder to understand. To avoid this added complexity, I kept the model as it is now.