Iterated Knowledge Isn't Better Knowledge*

Bernhard Salow
Magdalen College, Oxford

Abstract

Recent work in epistemology implicitly assumes that we can measure the quality or strength of someone's knowledge (whether understood intuitively, or by its normative connections to action, inquiry, belief, or assertion) by the number of iterations it permits. I show that this idea is hopeless, because, even in set-ups that look maximally friendly, one can construct cases where someone goes from having available only a single iteration of knowledge that p to having arbitrarily many such iterations, without their knowledge that p becoming better in any way.

I know that the Second World War ended in 1945. I also know that the Hundred Year War ended in 1453, having checked on *Wikipedia* a few days ago. But my knowledge of the first claim seems importantly different. It seems stronger, more secure. It seems an appropriate thing to act on or assert in a wider range of situations. Most of what further inquiry into the second claim might hope to procure, I seem to already have for the first. To introduce a shorthand, and without prejudging whether all these considerations point to the same thing, my knowledge seems *better*.

A common, albeit rarely explicit, idea in recent epistemology is that we can measure (some aspect of) the quality of someone's knowledge by the number of iterations it permits.[1] More precisely, say that to 1-know that p is to know that p, that to $n+1$-know that p is to know that one n-knows p, and that to ω-know

*Thanks to Kevin Dorst, Jeremy Goodman, two anonymous referees, and audiences at New York Univeristy and the University of Oxford for helpful feedback.

[1]We will see that proposals from Williamson (2005), Benton (2013), Schulz (2017), and Carter and Hawthorne (2024) all require something like this to be true. Other occurrences of similar ideas include the following. Carter (2022, p.45) equates higher iterations of knowledge with greater epistemic security. Goldstein (2024, p.1) opens his discussion by asserting that "Omega knowledge is the strongest kind of knowledge." He also (2024, ch. 6.2.2) considers measuring the strength of one's epistemic position by the number of available iterations of knowledge, but sees no use for the notion. Fiat (forthcoming) refers to infinitely iterated knowledge of p as "full knowledge", by contrast to the "partial knowledge" one has when the iterations run out; and his model is explicitly motivated by the idea that "it matters how many orders of knowledge [one] has" (§ 5.3.2). Kelp (2021, ch. 1.2) presents iterated knowledge as one of the "most obvious" candidates for a state stronger than knowledge that inquiry could aim at, although he argues that it does not. By contrast, Willard-Kyle (2023) and Woodard (2022) hold that it does sometimes make sense to inquire into p when one already knows p because one is aiming at iterated knowledge.

that p is to n-know that p for every n. The idea is then that S knows p better than q if S n-knows p for some n for which they do not n-know q.

I will argue that his idea is hopeless: nothing like it can be right. I will focus on examples that are maximally friendly to the required fine-grained facts about iterations. I will set aside all complications about the psychological underpinnings of such iterated knowledge. And I will show that, in these examples, someone can go from merely knowing p to having arbitrarily many iterations of knowledge that p, without their knowledge that p becoming better in any of the ways just canvassed. So additional iterations do not improve one's knowledge.

The paper proceeds as follows. Section 1 constructs the example. Section 2 analyzes the example in more detail. Section 3 draws out how the example raises problems for some recent discussions of knowledge norms for action (Williamson, 2005, Schulz, 2017), inquiry (Carter and Hawthorne, 2024), and assertion (Benton, 2013). Section 4 argues that variants of the example also tell against Goldstein's (forthcoming, 2024) proposals about the special significance of ω-knowledge. Section 5 briefly suggests some ways in which our discussion supports the KK principle, that one 2-knows everything one knows. Section 6 concludes by considering what else better knowledge might be.

1 The Example

I will develop my example within a version of the appearance/reality models introduced by Williamson (2013, 2014). These have a number of features that make them particularly suitable for our purposes. They initially look friendly to the idea that knowledge is better when more iterations are available. They allow us to think precisely about the number of iterations of knowledge that a subject has – something that, as Bonnay and Égré (2009) and Greco (2014) have emphasized, our intuitive understanding doesn't obviously allow when the number of iterations exceeds two. And they are naturally idealized, so that we needn't worry about the psychology involved, and hence about whether the relevant iterations are achieved or merely available.

Consider, then, the following scenario:

> *Headcount*
> You are in a large lecture room. You count the attendees, getting the answer 350. But you know that you often miscount slightly with numbers this big.

Here is a Williamsonian model of this example. You have a margin for error; for concreteness, suppose it's 5. You always believe that the number of people in the room is within your margin for error of the count, i.e. between 345 and 355. If your count was, in fact, perfectly correct, this belief amounts to knowledge. If your count wasn't perfectly correct, you know less. In particular, suppose that it was out by x. Then you know only that the number of people in the room is within $x + 5$ of your count. After all, since you could easily have miscounted by up to 5 even when you get it exactly right, and you did miscount by x on this occasion, you could easily have miscounted by $x + 5$ on this occasion; so you cannot know that you didn't.[2]

To get a simple model of knowledge iterations, let us assume that you have many iterations of knowledge about all these facts about the setup; and that you have thought long and hard about your situation, so that you know everything your knowledge entails. Finally, suppose that your count was, in fact, exactly right. Then you know that the number of people is between 345 and 355. Of the cases where the number of people is between 345 and 355, you know the least about how many people there are in the ones where your count was out by 5, i.e. where the number is 345 or 355. But even in those cases, you know that the number is between 340 and 360. So you know that you know that number is between 340 and 360, but do not know that you know anything stronger. More generally, writing $K^n(x)$ for the set containing all the numbers such that, for all you n-know when you miscounted by x, there are that many people in the audience, we have that $K^n(x) = [350 - x - 5n, 350 + x + 5n]$.

This model seems friendly to the idea of measuring the strength of knowledge by the number of available iterations. Given a perfect count, you know that there are between 345 and 355 people in the room; but you know better that there are between 340 and 360, and even better that there are between 335 and 365. The model also predicts that your knowledge that there are either between 345 and 355 or exactly 300 people in the room is equally good to your knowledge that there are between 345 and 355, which seems subtle, but correct. Finally, the model predicts that your knowledge that there are between 344 and 356 people in the room is not better than your knowledge that there are between 345 and 355; this doesn't look clearly correct, but it doesn't look obviously wrong either.

However, we now add to the story:

> *Headcount with Oracle*
> You are in a large lecture room. You count the attendees, getting the answer 350. But you know that you often miscount slightly with numbers this big. On a whim, you ask an oracle whether you miscounted by 1 to 5. The oracle informs you that you did not.

I will assume that the oracle is extremely trustworthy, and that you have many iterations of knowledge that this is so. So when the oracle tells you that you did not miscount by 1 to 5, you come to m-know, for some high value of m, that you did not miscount by 1 to 5, and hence that the size of the audience is not between 345 and 349, or between 351 and 355. I will also assume that the

[2] This model of the case is not uncontroversial. Goodman (2013) proposes an alternative model to which my discussion could easily be adapted. Cohen and Comesaña (2013) and Stalnaker (2015) propose alternatives that validate the KK principle. These models would block my example, but are incompatible with the idea I am criticizing, since the KK principle entails that knowledge can always be iterated arbitrarily often. Goldstein (2024) develops a number of variant models, which I discuss in section 4 below. Independently of these controversies, it's clear that the model can't be exactly right, since it doesn't represent knowledge you have independently of your count, e.g. that there are more than 10 people but fewer than 10,000. But this simplification is harmless, as long as we focus only on scenarios where the number is reasonably close to the count, say between 300 and 400.

oracle's announcement does not cause you to lose any knowledge you previously had if there are in fact 350 people in the room (and that you m-know that this is so). I'm not convinced that this assumption is plausible in the present version of *Headcount with Oracle*, for reasons I explain in section 2.3. But as I show there, we can revise the example so that this assumption becomes plausible, without affecting the remainder of the argument.

It follows from these assumptions that, if there are 350 people in the audience, you come to m-know that there are 350 people people in the audience when you hear the oracle's announcement.

To see this, note first that the setup entails that, if there are no more than 350 people, you know that the number of people is between 345 and 355 based on your count. Since you come to know the oracle's information, you also know that the number of people is not between 345 and 349 or between 351 and 355. Putting your information together, you know that there are 350 people.

So you know that there are 350 people. But you also know the setup, and the setup entails that, if there are 350 people, then you know that the number of people is between 345 and 355 based on your count. So you 2-know that the number of people is between 345 and 355. Since you also 2-know the Oracle's information, you 2-know that the number of people is not between 345 and 349 or between 351 and 355. Since you know that you've put all your information together, you 2-know that there are 350 people.

So you 2-know that there are 350 people. But you also 2-know the setup, and the setup entails that, if there are 350 people, then you know that the number of people is between 345 and 355 based on your count. So you 3-know that the number of people is between 345 and 355. Since you also 3-know the Oracle's information, you 3-know that the number of people is not between 345 and 349 or between 351 and 355. Since you 2-know that you've put all your information together, you 3-know that there are 350 people.

We can keep repeating this reasoning for ever larger n, as long as you have the required iterations of knowledge of the setup and of the oracle's information. So you n-know that there are 350 people in the audience, for any $n > 1$ such that you $(n-1)$-know the setup and n-know the oracle's information.

Consider now your knowledge that there are at most 355 people in the audience. The Oracle's announcement in no way supports this claim. The Oracle told you that there are not 345 to 349 people, nor 351 to 355. This is exactly what (you know) it would have told you if there were more than 355 people. So the announcement cannot make it more likely that there are at most 355 people. Similarly, the closest or most normal possibilities in which you start off believing falsely that there are at most 355 people are all possibilities in which the Oracle makes the same announcement; so they are all possibilities in which you retain the belief after the announcement; and so your post-announcement belief is no safer or more reliable than your pre-announcement belief. More generally, there seems no sense in which your knowledge that there are at most 355 people could have become stronger, or more secure, as a result of the announcement.

The same seems true for any of the other ways of thinking about the quality of your knowledge. If there were any reservations about you asserting that, believing that or acting on the assumption that there are at most 355 people before the announcement, the announcement will have done nothing to change that. In so far as you are interested in whether there are at most 355 people, the case for recounting or seeking independent confirmation is no worse after the announcement than it was before. So you do not know that there are at most 355 people any better after the oracle's announcement than you did before.

This makes your knowledge that there are at most 355 people in the audience a counterexample to the claim that iterated knowledge is better knowledge. Before the announcement, you only 1-knew that there are at most 355 people. After the announcement, you m-know this claim for very large m, since it is entailed by your m-knowledge that there are 350 people. Since your knowledge that there are at most 355 people has become no better, despite the vast increase in the number of available iterations, more iterations do not, in general, make for better knowledge.[3]

In what follows, your knowledge that there are at most 355 people in the audience in *Headcount with Oracle* will be my central example of knowledge that doesn't become better when it is iterated. But before drawing out the implications, it may be useful to analyze this example in more detail.

2 Analysis and Variations

In the previous section, I argued that in *Headcount with Oracle* you can get iterated knowledge that isn't better knowledge. This section further analyzes the example in three independent (and skippable) ways. First, it reconstructs the example using epistemic logic; this illuminates the structure and makes clear that the problem is very general, and not tied to the details of *Headcount*. Second, it relates the example to the notion of physical safety sometimes used to provide an analogy for iterated knowledge. Finally, it returns to the assumption, flagged above, that, in *Headcount with Oracle*, the oracle's announcement does not cause you to lose knowledge gained from your count.

2.1 Epistemic Logic

In epistemic logic, we have a set of possibilities W, and a binary relation R between them. We say that you know, at w, that p if and only if p is true at every $v \in W$ such that wRv. This allows for an elegant representation of iterated knowledge. Say that there is an n-chain from w to v if there are $u_1, u_2, \ldots, u_{n-1}$ such that $wRu_1, u_1Ru_2, \ldots, u_{n-1}Rv$. Then you n-know that p at w if and only if there is no n-chain from w to a possibility v in which p is false.

The standard treatment of "public announcements" in dynamic epistemic logic models them as simply removing from the model all possibilities incom-

[3] The example also involves going from lacking justification to believe even that you 1-know that there are at most 355 people to having justification to believe that you m-know that there are at most 355 people. So the example also makes trouble for the idea that justification to believe that one has more iterations of knowledge that p amounts to better justification for believing p, or justification for believing p more strongly.

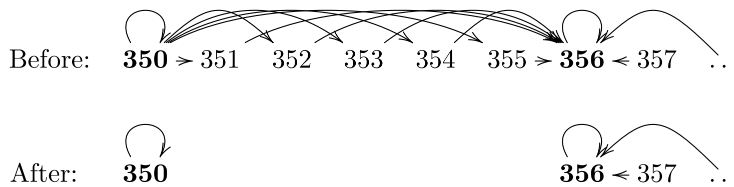

Figure 1: Epistemic accessibility from possibilities where there are 350 people in the audience and to possibilities where there are 356 people in the audience, before and after the oracle's announcement. Possibilities where the number of people is below 349 have been ommitted to avoid clutter.

patible with their content.[4] This builds in some controversial assumptions: it means that you don't just come to know, but ω-know the content; and it means that the announcement has no other effect on what you know, such as defeating knowledge you previously had. But we can think of these as idealizations that compensate for their inaccuracy with a significant increase in tractability.

The key observation is then that removing just a few possibilities from a model can dramatically affect what chains it includes. Prior to the announcement, we could get from a possibility where the number is 350 to one where it is 356 in two steps, via a possibility where the number is between 351 and 355. But with those intermediate possibilities removed from the model, there may be no way at all to get from the 350 possibility to the 356 one; see figure 1. As a result, your mere 1-knowledge that the number of people is not 356 has been upgraded to ω-knowledge – even though no possibilities in which the number of people is 356 have been eliminated.

This analysis makes clear that the construction behind the example generalizes far beyond the details of *Headcount*. In fact, consider any possibility w. We can then define $K(w) := \{v : wRv\}$, the set of possibilites consistent with what you know in w. We can also define $K_{\leq}(w) := \{v : K(v) \subseteq K(w)\}$, the set of possibilities in which you know everything you know in w (and maybe more). We then construct an announcement that is false in all and only the possibilities in $K(w)$ that are outside $K_{\leq}(w)$ – that is, in all and only the possibilities consistent with what you know in w in which you do not know something that you do know in w. Removing the possibilities in which this announcement is false cuts any n-chain from w to any possibility outside $K_{\leq}(w)$. So you now ω-know in w everything that you 1-knew in w before the announcement. But the announcement has only eliminated possibilities inside $K(w)$, so only possibilities in which the things you previously 1-knew are false. So it does not seem plausible that all these pieces of knowledge should have become better as a result of the announcement.

[4]See Baltag and Renne (2016), who trace this to Plaza (1989).

2.2 Iterated Safety

At least anecdotally, the idea that iterated knowledge is better knowledge is sometimes motivated by an analogy with physical safety, influentially drawn by Williamson (2000, ch.5). The rough idea is that knowledge requires safety from error, and that being safely safe from something makes you safer from it. Examples analogous to *Headcount with Oracle*, however, show that this second claim is not generally true: iterated safety isn't always safer.

Consider skydiving. Any good parachute comes with a backup. In order for skydiving to be safe, both your primary and your backup parachute need to be working. But having two working parachutes doesn't make skydiving safely safe: one could easily be defective (that's why you need two), and if it were, skydiving would not be safe – you would survive it only by luck. An obvious way to make skydiving safely safe is to introduce a third parachute: that way, you would be safe even if one were defective. This way of making skydiving safely safe would also make it safer.

But there are other ways of making skydiving safely safe. Suppose a guardian angel inspects all the parachutes you ever come across, and removes any in which exactly one of the primary or the backup is defective. So, when you have a parachute where both components are working, it is no longer true that one could easily have been defective. So skydiving is now safe, not just in actuality, but in any other situation that you could have easily found yourself in. So skydiving has become safely safe. But notice that the guardian angel doesn't remove parachutes where both the primary and the backup are defective. So the guardian angel does not make skydiving safer.

2.3 Defeat

In arguing that, in *Headcount with Oracle*, you m-know that there are 350 people in the room, I assumed that you know everything that you know in *Headcount*. In particular, I assumed that, if there are 350 people in the room, you know that there are between 345 and 355 people based on your count even after the Oracle has spoken. But this could be denied. Plausibly, the Oracle's announcement provides you with evidence that there are not, after all, between 345 and 355 people in the room: it eliminates some possibilities in which there are 345 and 355 people, but does not eliminate any possibilities in which there are not. So it may well be that the oracle's announcement defeats your initial knowledge that there are between 345 and 355 people in the room, undermining the argument.

Figure 2: Probability of different numbers of people in *Headcount*

≤342	343	344	345	346	347	348	349	350
.8%	1.1%	2.1%	3.7%	5.7%	8.1%	10.3%	11.9%	12.5%

351	352	353	354	355	356	357	≥ 358
11.9%	10.3%	8.1%	5.7%	3.7%	2.1%	1.1%	.8%

Fortunately, the example can be adapted to address this concern. Consider the following schematic version:

> *Headcount with n-Oracle*
> You are in a large lecture room. You count the attendees, getting the answer 350. But you know that you often miscount slightly with numbers this big. On a whim, you ask an oracle (i) whether you over-counted and (ii) whether you miscounted by 1 to n. The oracle informs you that you did neither.

When n is large enough, it is implausible that the oracle's announcement would defeat your knowledge that there are between 345 and 355 people in the audience. After all, it now eliminates many possibilities in which this claim is false. As a result, the original argument that, after the announcement, you m-know that there are 350 people in the audience goes through. So you also m-know that there are no more than n people in the audience, despite starting with very few iterations of knowledge that this is so. Yet your knowledge that there are no more than n people in the audience has not improved as a result of the oracle's announcement. The original problem thus re-arises.

To make this less abstract, let's fill in some details. It is natural to think that, if the oracle's announcement defeats your knowledge that there are between 345 and 355 people in the audience, it does so by making the probability of this claim on your evidence fall below a threshold required for knowledge. What do the probabilities on your evidence look like? Plausibly, they initially form an approximate "bell curve" over the possible number of people, peaked around 350. A natural way to generate such an approximate a bell curve is to equate the probability that there are n people in the room with the probability of getting $20 + (n - 350)$ when flipping a fair coin 40 times (where 40 is a somewhat arbitrary choice of an even number, and 20 is chosen because it is the most likely number of heads). The resulting probabilities are recorded in figure 2.

Given this probability distribution, the initial probability that there are between 345 and 355 people in the room is roughly 91.9%. If the Oracle then informs you that you didn't overcount, and that you did not miscount by 1 to 7, the probability that there are between 345 and 355 people rises to roughly 94.0%. So the oracle's announcement does not reduce the probability that there are between 345 and 355 people, hence does not defeat your knowledge that there are between 345 and 355 people in the room, and hence permits you to m-know that there are 350 people in the room. It thus also permits you to m-know that there are no more than 358 people in the room, which you previously only 1-knew; but it does not make your knowledge that there are no more than 358 people in the room any better in the process.[5]

3 Implications

3.1 Action

Many think that we can act on what we know. But a problem arises when the stakes are very high: if one of my options has disastrous consequences if p is false, it looks irrational to choose it, even if I know that p is true, and hence that the option won't have those consequences. In response, Williamson (2005) and Schulz (2017) suggest that, when the stakes are high, I can only act on p if I also n-know that p for $n > 1$, with n increasing as the stakes do.[6]

Our example sinks this proposal.[7] Suppose you are the fire-safety officer for the lecture room, and that the official safety limit prescribes an audience size of at most 355 people; it is your job to make sure that the limit isn't breached. Before the oracle makes its announcement, you had 1-knowledge that the audience contains 355 people or less, but not 2-knowledge. Once the oracle has made its announcement, you have m-knowledge that the audience contains 355 people or less for many $m \geq 2$. But there are no stakes that would have made it irrational for you to rest content before the oracle's announcement, that do not also make it irrational for you to rest content after the announcement. So additional iterations of knowledge of a proposition do not make that proposition more suitable to act on when the stakes are higher.

Note that this objection leaves intact a variant of the idea. Instead of saying that, as the stakes rise, the iterations of knowledge required to act on p rise, we could simply say that, as the stakes rise, better knowledge – understood in some other way – is required to act on p.[8] The problem highlighted by our example arises from the attempt to explain quality in terms of the number of available iterations, not from the idea that higher stakes require better knowledge.

[5] A final loose thread: Goodman (2013) and Goodman and Salow defend models on which, given probability distributions like the above, the margin for error shrinks for higher miscounts. However, given natural choices of their parameters and the probabilities described above, their models still predict that, had the number of people initially been 355, you could not have known that there were no more than 358 people in the audience. So this complication does not threaten the claim that you initially only had 1-knowledge that there are no more than 358 people in the audience.

[6] Williamson is non-committal about whether this should be treated as a theory of when someone may act on p or as a theory of when someone's acting on p would be appropriately subject to significant criticism; my problem arises on either construal.

[7] Gao (2019), Schulz (2021), and Vollet (2025) criticize the idea for other reasons.

[8] Schulz (2017, p.480) floats this idea; Schulz (2021) defends and develops it in more detail.

3.2 Inquiry

Many think that inquiry aims at knowledge. However, it sometimes seems to make sense to inquire further into p even when one is already in a position to know p. To resolve this tension, Carter and Hawthorne (2024) suggest that a norm to ϕ only if X gives rise to a derivative norm to ϕ only if you know that X; this then gives rise to an even more derivative norm to ϕ only if you 2-know that X; and the process continues indefinitely.[9] Combined with a truth or knowledge norm on belief, this implies that someone who believes that p satisfies more norms the more iterations of knowledge they have. Carter and Hawthorne conclude that epistemically ideal agents believe p only if they ω-know p. They also hold that agents like us might reasonably aim for something less, such as believing p only if we n-know p for some large n.

Against this background, Carter and Hawthorne maintain that one may be in a position to know that p, without yet being in a position to n-know that p. In such a case, one may be reluctant to close inquiry by forming the belief that p, since this belief would not satisfy all one's epistemic aims. However, by continuing to gather evidence about p, one may eventually become positioned to n-know that p, so one has reason to keep investigating.

Our example refutes this explanation. The belief that there are at most 355 people in the room is no more or less ideal after the announcement than it is before it. If you had good reason to double-check that it is correct before the announcement, say by recounting, then you have equally good reason to double-check after the announcement. Increasing the number of available iterations need not affect how epistemically ideal a belief is, nor how strong the reasons to keep investigating are.

Again, our objection leaves intact a variant of the idea. Instead of saying that the reason to keep gathering evidence is to make available additional iterations of knowledge, we could say that the reason to keep gathering evidence is to make available better knowledge, understood in some other way. Of course, this requires motivating, independently of the particular derivative norms Carter and Hawthorne appeal to, that it can be reasonable to aim to believe only when we have knowledge that is better than the minimum. It remains to be seen whether this can be done. The point here is only that, as in section 3.1, it is not the appeal to better knowledge that generates the problem, but the temptation to understand this in terms of the number of available iterations.

3.3 Assertion

Many philosophers accept that knowledge is the primary norm of assertion, but reject the KK principle, that one 2-knows everything one knows. As noted by Sosa (2008), this combination of views doesn't immediately explain the oddity of quasi-Moorean assertions of the form 'p, but I don't know whether I know

[9] See also Williamson (2005, 2009) and Goldstein (forthcoming, 2024). These derivative norms differ from ones formulated in dispositional or character-centric terms, as in Hawthorne and Stanley (2008), Lasonen-Aarnio (2010), or Williamson (forthcoming).

p'; after all, one could know p without knowing that one does. To provide an explanation, Benton (2013) suggests that asserting p when one doesn't 2-know p is careless, since it means that one doesn't know that one is acting in line with the primary norm. Quasi-Moorean assertions are then odd because they are necessarily either careless or false.

This response is refuted by our example.[10] An assertion that there are at most 355 people in the room is no more acceptable after the oracle's announcement than it is before, the additional iterations of knowledge notwithstanding. So either carelessness isn't bad; or doing something one doesn't know doesn't violate a norm needn't be careless; or there is some deep mistake in the Williamsonian model generating our example. None of these sit easily with Benton's explanatory ambitions.[11]

4 ω-Knowledge isn't Better Knowledge

Goldstein (forthcoming, 2024) does not hold that finitely iterated knowledge is better knowledge. But he holds that infinitely iterated knowledge – ω-knowledge – is better knowledge. The theses he is sympathetic to include: that we should only act on what we ω-know; that we should only assert what we ω-know; and that we should be certain of all and only what we ω-know. Since Goldstein holds that we do not ω-know everything we know, any of those would make ω-knowing p better than merely knowing p.

As it stands, the example from section 1 is not effective against Goldstein's theses. The simple reason is that the example does not involve the subject coming to ω-know anything. But a deeper reason is that the example is developed within a broadly Williamsonian model, on which any ω-knowledge is a mere artifact of the idealizations. Since the theses Goldstein discusses are only plausible if there is substantive and interesting ω-knowledge, assessing them charitably requires us to reject these Williamsonian models.

Fortunately, Goldstein presents alternatives. In fact, he presents whole families of alternatives. However, I will show that our example can be adapted to raise problems on all three of the models he presents.

The adaptation is straightforward in the models involving 'Variable Margins' (2024, chs. 4 and 5.5) and 'Reflective Luminosity' (2024, chs. 2 and 5.3). A 'Variable Margins' model, applied to *Headcount*, might go as follows. Let x represent how far off your count was. If $x < 5$, the margin is 5, so that you know the number of people is within $x + 5$ of 350. If $5 \leq x < 10$, the margin is 4, so that you know that the number of people is within $x + 4$ of 350. If

[10] Goldstein (2024) argues that this response also doesn't account for the full data surrounding the oddity of the kinds of quasi-Moorean conjunctions under discussion. Kirkpatrick (ms) extends Benton's proposal to handle this additional data, but the resulting view doesn't help with the problem identified here.

[11] Haziza (2021) also suggests that knowledge norms on assertion ultimately generate norms requiring higher iterations of knowledge, because of the norms governing conversational implicatures. However, he accepts the KK principle; and so he is not vulnerable to our example, which presupposes that the KK principle is false.

$10 \leq x < 15$, the margin is 3, so that you know that the number of people is within $x + 3$ of 350. And so on, until, when $25 \leq x$, the margin is 0 and you know that the number of people is within x of 350. In this model, you can have ω-knowledge about how many people are in the audience; for example, if you counted correctly, you ω-know that there are between 335 and 365 people.

A version of the 'Reflective Luminosity' model, applied to *Headcount*, goes as follows. When $0 \leq x \leq 5$, $10 < x \leq 15$, $20 < x \leq 25$, etc., knowledge works as it does in the Williamsonian model, so that you know that the number of people is within $x + 5$ of 350. But when $5 < x \leq 10$, $15 < x \leq 20$, $25 < x \leq 30$, and so on, you always know, respectively, that $x \leq 10$, that $x \leq 20$, that $x \leq 30$, and so on. On this model, what you ω-know when the number of people is between 340 and 360 is that it is between 340 and 360; what you ω-know when the number of people is between 330 and 340 or between 360 and 370 is that it is between 330 and 370; and so on.

Both models share with the Williamsonian model the prediction that, if your count was perfectly correct, you initially 1-know but do not 2-know that the number of people is between 345 and 355. So the only changes we make are to (a) assume that you ω-know the setup and (b) assume that you come to ω-know the content of the oracle's announcement. These would have been problematic assumptions in the context of section 1; but on the present view, ω-knowledge is possible, even common, so it is hard to deny that their legitimacy. As a result of coming to ω-know the content of the oracle's announcement, your 1-knowledge that there are at most 355 people in the room then becomes ω-knowledge. But the claim that there are at most 355 people in the room does not become a better thing to act on, a better thing to assert, or a more sensible thing to be certain of than it was before the announcement. So our example refutes the special significance of ω-knowledge on these models.

Things are a little more complicated with the models built to vindicate a principle Goldstein calls 'Fragility' (2024, chs. 3 and 5.4). Applied to *Headcount*, these work as follows. There is a privileged range of cases, where x is small, in which you ω-know everything you know. Outside of this range, knowledge works as it does in the Williamsonian model above. More precisely, we might say that if $x \leq 5$, you always know that the number of people is within 5 of 350; and that if $x > 5$, you know that the number of people is within $x + 5$ of 350. On this model, you ω-know that the number of people is between 345 and 355 whenever it is.

Since, if your count was perfectly accurate, you ω-know that there are at most 355 people in the audience from the start, we will need to proceed differently this time. Instead of supposing that your count was correct, suppose that it was off by 6: there are actually 356 people in the audience. Then you know that there are at most 361 people in the audience, but do not 2-know that there are at most 361. This time, you ask the oracle whether you miscounted by 7 to 11; again, the oracle announces that you did not, giving you ω-knowledge that you did not. Adapting the reasoning from the original case, this means that you now have ω-knowledge that there were at most 361 people in the audience, which you did not previously have. But the claim that there were at most 361 has not become a better thing to take for granted, a better thing to assert, or a more sensible thing to be certain of than it was before the announcement. So our example refutes the special significance of ω-knowledge on these models too.

I conclude that, on each of the family of models proposed by Goldstein, our example can be adapted to refute the special significance of ω-knowledge over ordinary knowledge.

5 An argument for KK?

The KK principle holds that you 2-know anything you know; repeated applications of this principle entail that you ω-know anything you know, so that all distinctions in the number of available iterations of knowledge are obliterated. In so far as these distinctions are theoretically useful, this tells against KK. By showing that such distinctions are less useful than one might have thought, the considerations raised here thus weaken an argument against KK.

Benton's (2013) ideas about assertion and recklessness are designed to undercut an argument for KK. One of Goldstein's (2024) stated ambitions is similarly to undermine arguments for KK. His insight is that many of these arguments are really arguments that we often have ω-knowledge and that ω-knowledge is significant in a variety of ways. For example, the oddity of asserting 'p, but I don't know whether I know that p' is easily explained by an ω-knowledge norm of assertion; KK's additional claim that all ordinary knowledge is ω-knowledge is irrelevant. However, our examples show that, even if it is possible to know p without ω-knowing it, the change from one to the other need not constitute an improvement in one's epistemic position with respect to p. This makes it implausible that ω-knowledge is significant in a way that knowledge is not. So arguments for the significance of ω-knowledge are arguments for KK, undermining Goldstein's response. Our examples thus strengthen existing arguments for KK, by raising problems for attempts to undercut them.

Finally, we get a new positive argument for KK. For we have seen arguments that iterated knowledge should be better knowledge. In Carter and Hawthorne's (2024) and Goldstein's (forthcoming, 2024) discussions, this is motivated by the idea that a norm to ϕ only if X gives rise to a derviative norm to ϕ only if you know that X. It can be similarly motivated from Benton's (2013) idea that a norm to ϕ only if X makes it reckless to ϕ if you don't know X, combined with a norm against recklessness. But these motivations go through only if KK can fail: if KK is true, one satisfies all these derivative norm and avoids all these forms of recklessness whenever one knows that X, and so we do not get the false prediction that iterated knowledge of X would be better. In so far as the norms and principles used in these motivations are appealing, this is a new consideration in favour of KK.

6 What is better knowledge?

I have argued that we cannot understand the quality of one's knowledge in terms of the number of iterations it permits. Assuming that this is a useful and unified notion, what are the alternatives?

One idea would be to appeal to properties of the attendant belief. We hold some of our beliefs with greater confidence than others; we also (which may or may not come to the same thing) require stronger evidence to change our minds about some of them than others. We might thus suggest that S knows p better than q if S knows p and their belief that p (or justification for believing p) is stronger. But such proposals run into trouble in a variant of *Headcount with Oracle* where one has strong but misleading evidence that the Oracle is trustworthy, when it is in fact completely unreliable. For suppose that your count was, in fact, correct, that you ask the oracle whether you miscounted by more than 5, and that it correctly but unreliably announces that you did not. Since you have every reason to trust the oracle, hearing its announcement will (quite rightly) bolster your confidence in the claim that there are between 345 and 355 people in the room, and make it harder to dislodge your belief through additional evidence. But, since the oracle is in fact completely unreliable, you do not now know this claim any better than you did beforehand.

Another idea might be to appeal to defeat, or to conditional knowledge. We might say that S knows p better than they know q, if they know that, even if p and q are not both true, p still is; or, alternatively, if compelling evidence that not both are true would defeat their knowledge of q but not their knowledge of p. Unfortunately, however, this doesn't answer cardinal questions we might have: that someone's knowledge that p is their worst knowledge in this sense doesn't tell us whether this is because it isn't very good, or because the rest is even better. So this notion couldn't play the roles envisaged here, of setting a standard higher than knowledge that one might aim for or hold others to.

A more promising idea is to look at theories about the nature of knowledge. On many such theories, there is a parameter which could be set in more or less stringent ways. On a safety conception, we can ask how close a possibility needs to be before it could have easily obtained. On some normality-based accounts, we can ask how much more abnormal than actuality an evidential possibility needs to be for us to know it doesn't obtain; on other versions, we might be able to set the parameters that determine when a possibility is less normal than actuality in different ways. On a virtue theoretic account, we can ask how much of the success of a belief should be attributed to the capacities exercised, which we might understand in terms of how (much more) unfavourable background conditions would have to become for those capacities not to result in success. Given such parameters, we can then rank pieces of knowledge according to whether they would still have qualified as knowledge if those parameters had been set in more stringent ways. The devil will, of course, be in the details. But accounts along these lines seem well-placed to predict that, in *Headcount*, your knowledge that there are between 340 and 360 people is better than your knowledge that there are between 345 and 355, while also predicting that the oracle's announcement in *Headcount with Oracle* does not improve your knowledge that the number is at most 355. If there is a good account of what makes some knowledge better knowledge, I suspect that it runs along these sorts of lines.

References

Alexandru Baltag and Bryan Renne. Dynamic Epistemic Logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2016 edition, 2016.

Matthew Benton. Dubious objections from iterated conjunctions. *Philosophical Studies*, 162(2):355–358, 2013. doi: 10.1007/s11098-011-9769-3.

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.

Sam Carter. Degrees of assertability. *Philosophy and Phenomenological Research*, 104(1):19–49, 2022. doi: 10.1111/phpr.12725.

Sam Carter and John Hawthorne. Dogmatism and inquiry. *Mind*, 133(531): 651–676, 2024. doi: 10.1093/mind/fzad070.

Stewart Cohen and Juan Comesaña. Williamson on gettier cases and epistemic logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):15–29, 2013. doi: 10.1080/0020174x.2013.775012.

Yonathan Fiat. Knowing to infinity: Full knowledge and the margin-for-error principle. *Philosophy and Phenomenological Research*, forthcoming. doi: 10.1111/phpr.70001.

Jie Gao. Against the iterated knowledge account of high-stakes cases. *Episteme*, 16(1):92–107, 2019. doi: 10.1017/epi.2017.30.

Simon Goldstein. *Iterated Knowledge*. Oxford University Press, Oxford, 2024.

Simon Goldstein. Omega knowledge matters. *Oxford Studies in Epistemology*, forthcoming.

Jeremy Goodman. Inexact knowledge without improbable knowing. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):30–53, 2013. doi: 10.1080/0020174x.2013.775013.

Jeremy Goodman and Bernhard Salow. In Joseph Y. Halpern and Andres Perea, editors, *Theoretical Aspects of Rationality and Knowledge 2021 (TARK 2021)*, pages 171–186. doi: 10.4204/EPTCS.335.15.

Daniel Greco. Could KK be OK? *Journal of Philosophy*, 111(4):169–197, 2014. doi: 10.5840/jphil2014111411.

John Hawthorne and Jason Stanley. Knowledge and action. *Journal of Philosophy*, 105(10):571–590, 2008. doi: 10.5840/jphil20081051022.

Eliran Haziza. Assertion, implicature, and iterated knowledge. *Ergo: An Open Access Journal of Philosophy*, 8(n/a), 2021. doi: 10.3998/ergo.2236.

Christoph Kelp. *Inquiry, Knowledge, and Understanding*. Oxford University Press, Oxford, 2021.

James Ravi Kirkpatrick. Assertion at the limits: In defence of norm contextualism. ms.

Maria Lasonen-Aarnio. Unreasonable knowledge. *Philosophical Perspectives*, 24 (1):1–21, 2010. doi: 10.1111/j.1520-8583.2010.00183.x.

Jan Plaza. Logics of public communications. In M.L. Emrich, M. S. Pfeifer, M. Hadzikadic, and Z.W. Ras, editors, *Proceedings of the fourth international symposium on methodologies for intelligent systems: Poster session program*, pages 201–216. Oak Ridge National Labratory, 1989.

Moritz Schulz. Decisions and higher-order knowledge. *Noûs*, 51(3):463–483, 2017. doi: 10.1111/nous.12097.

Moritz Schulz. Practical reasoning and degrees of outright belief. *Synthese*, 199 (3-4):8069–8090, 2021. doi: 10.1007/s11229-021-03153-0.

David Sosa. Dubious assertions. *Philosophical Studies*, 146(2):269–272, 2008. doi: 10.1007/s11098-008-9255-8.

Robert Stalnaker. Luminosity and the KK thesis. In S. Goldberg, editor, *Externalism, Self-Knowledge, and Skepticism*. Cambridge UP, 2015.

Jacques-Henri Vollet. Antiluminosity, excuses and the sufficiency of knowledge for rational action. *Erkenntnis*, 90(1):21–42, 2025. doi: 10.1007/s10670-023-00682-7.

Christopher Willard-Kyle. Valuable ignorance: Delayed epistemic gratification. *Philosophical Studies*, 180(1):363–84, 2023. doi: 10.1007/s11098-022-01902-6.

Timothy Williamson. *Knowledge and its Limits*. Oxford UP, 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. Reply to Stephen Schiffer. In Patrick Greenough and Duncan Pritchard, editors, *Williamson on Knowledge*. Oxford UP, 2009.

Timothy Williamson. Gettier cases in epistemic logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1):1–14, 2013. doi: 10.1080/0020174x.2013.775010.

Timothy Williamson. Very improbable knowing. *Erkenntnis*, 79(5):971–999, 2014. doi: 10.1007/s10670-013-9590-9.

Timothy Williamson. Justifications, excuses, and sceptical scenarios. In Julien Dutant, editor, *The New Evil Demon: New Essays on Knowledge, Justification and Rationality*. Oxford University PRess, forthcoming.

Elise Woodard. Why double-check? *Episteme*, pages 1–24, 2022. doi: 10.1017/epi.2022.22.