The uncoordinated teachers puzzle

Michael Cohen

January 3, 2023

Consider the following distant cousin of the famous surprise exam puzzle. Aya and Bob are co-teaching a class with students Dim, Kim and Tim. On Sunday night, Aya sends the students an email stating that there won't be a quiz the following day (Monday). Few minutes later, Bob sends the students an email, stating that although he can't manage to contact Aya, if *based on what Aya told them* they know that there won't be a quiz on a given day of the week, then there won't be a quiz in the day that follows it. In addition, Bob writes, there will be a quiz some day this week.

Each student reacts differently to the two emails. Dim is reminded of the surprise exam paradox and concludes that the teachers must be uncoordinated. He argues: I know the quiz is not on Monday, so it can't be on Tuesday. But now I know it's not on Tuesday so it can't be on Wednesday, and so on. Tim argues that Dim's reasoning is invalid, as it requires the second-order knowledge that there is no quiz on Monday, which the students just don't have. According to Tim, the students know the quiz won't happen on Monday, but nothing more. Kim argues that both students are wrong: we know that the quiz won't be on Monday or Tuesday, but it can happen any day after.

On Wednesday morning, Aya and Bob give the students a quiz. They say that they gave the students enough time to prepare, letting them know that the quiz won't happen on Monday or Tuesday. Kim was right. Aya informed the students that the quiz will not happen on Monday. The students could have used the information given by Bob to conclude that there will be no quiz on Tuesday. But since the knowledge that the quiz will not happen on Tuesday is *not* entirely based on what Aya told them, they cannot use Bob's information again to conclude anything stronger from it. Wednesday is still possible.

This little story is not here to offer a new perspective on the traditional surprise exam paradox, but on Williamson's margin for error argument against the KK principle (Williamson 1992, 2000, 2014).[1] According to the KK principle, if an agent knows p,

[1]For a few other responses to Williamson's argument against the KK principle, see Neta and Rohrbaugh (2004), Conee (2005), Dutant (2007), Greco (2014), Halpern (2008), Docik and Égré (2009), Egré and Bonnay (2008, 2009), Sharon and Spectre (2008) and Stalnaker (2015).

then the agent is in a position to know that they know p.[2] Consider the following version of the Williamsonian argument against the KK principle. Mr. Magoo has imperfect eyesight. He looks at an unmarked analog clock that points to 12:17. With his bad eyesight, Magoo cannot see that it points to 12:17, but he does see (and so come to know) that the minute hand does not point to the 00 minute. Bob, Magoo's friend who happens to be a psychophysics expert, informs Magoo that even though he does not know what Magoo saw, he knows that Magoo's eyesight—under the conditions of that particular observation event—follows a *margin for error* of one minute. This means that Magoo cannot reliably distinguish between the clock pointing to minute i and minute $i \pm 1$ with his imperfect eyesight. It further implies that if Magoo does come to know (with his imperfect eyesight) that the minute hand does not point to minute i, then it cannot be pointing to minute $i + 1$.

Magoo's story and the uncoordinated teachers' story are analogous in two important respects. First, in both cases, the epistemic agent is confronted with two distinct sources of information. One source gives the agent information about the environment (the quiz won't be on Monday; the hand does not point to 00). The second source specifies what can be deduced from the knowledge obtained from the first source (if you know the quiz won't happen on day i, it won't happen on day $i + 1$; if you know the minute hand does not point to i, it does not point to $i + 1$). Second, in both cases, the second source refers to the knowledge state which the first source produced (Bob the teacher refers to the knowledge obtained from Aya's email; Bob the psychophysicist refers to the knowledge obtained from the inexact observation).[3]

For complete clarity, we can lay out the knowledge elements of the analogy:

- The students' knowledge that the quiz will not happen on Monday, based on Aya's testimony, is analogous to Magoo's knowledge that the hand does not point to 00, based on his eyesight.

- The students' knowledge that if they know that the quiz will not happen on a given day, based on what Aya told them, it will not happen in the following day, is analogous to Magoo's knowledge that if he knows that the hand does not point to i, based on his eyesight, it does not point to $i + 1$.

Williamson believes that Magoo's situation is inconsistent with the KK principle, and concludes, given the plausibility of Magoo's situation, that the KK principle is false. Magoo knows that the minute hand does not point to 00, and he knows that if he knows that

[2] See Das and Salow (2016), Dorst (2019), Goodman and Salow (2018), Greco (2014a,2014b, 2015, 2017) and Stalnaker (2015) for recent defenders of the KK principle.

[3] In Williamson's versions of the argument, the agent's knowledge of their margin for error is the result of reflection, not expert testimony like in my version. This difference is not essential to the point I am making here. In both versions, knowledge of margin for error does not come from the perceptual event in question.

the minute hand does not point to 00, it cannot point to 01. But according to Williamson, Magoo cannot conclude by modus ponens that the minute hand does not point to 01, because he does not know that he knows that the minute hand does not point to 00. Supposing the KK principle, Williamson presents the following argument: Magoo knows that the clock does not point to 00. By the KK principle, he knows that he knows that. Applying his knowledge of his margin for error, Magoo can come to know that the clock does not point to 01. By the KK principle, he knows that he knows that. Applying his knowledge of his margin for error again, Magoo can come to know that the clock does not point to 02. According to Williamson, Magoo can repeat this reasoning pattern again and again, ruling out every position of the clock, including the actual one. Since this is absurd, Williamson concludes that we have to reject the KK principle, in pain of contradiction.

The analogy between the stories can be extended to the responses stage. The Dim response to Magoo's situation is that there is a contradiction between Magoo's perceptual knowledge and the margin for error information received from Bob, because it allows Magoo to rule out every possible position of the clock. Recall that in the quiz case, Dim likewise came to the conclusion that he can rule out every day of the week. And the Tim response is to reject second-order knowledge in order to block the Dim response.

I think that the Kim response is appropriate in both stories. According to Kim, the information received from Bob can only be used once. In the quiz case, we can only use it to eliminate the Tuesday possibility; in Magoo's case, we can only use it to eliminate the 01 possibility. After that, we cannot rule out anything further. In the quiz case, Bob's information only pertains to the knowledge the students got from Aya; in Magoo's case, Bob's information only pertains to the knowledge Magoo got from his imperfect eyesight. In both cases, once we use the information given by Bob, the resulting knowledge state is not the type of knowledge for which Bob's information is applicable. Thus, the margin for error information received from Bob cannot be reused.

Suppose that in the quiz story, Bob's email stated that if the students know, *no matter how*, that the quiz will not happen on day i, then it won't happen on day $i + 1$. In this *unrestricted* case, Dim's response would be appropriate: Aya and Bob are not coordinated. This version of the story, however, is not analogous to the Williamsonian margin for error stories. The Margin for error principle in Magoo's story is a feature of the imperfect *perceptual* capacities of epistemic agents, not their *general* epistemic capacities. Magoo has a perceptual margin for error of one minute; if God comes and tells Magoo that the clock is not pointing to minute 16, then Magoo will come to know that, even though he cannot rule out that the clock is pointing to 17. This scenario does not violate Magoo's perceptual margin for error, since Magoo's knowledge of the clock is not perceptual anymore. The margin for error principle does not state that if Magoo knows, *no matter how*, that the minute hand is pointing to i, then it cannot be pointing to $i + 1$. It is restricted to the knowledge obtained from the inexact perceptual event. In both stories, Bob's information is *restricted* to a particular source of information (Aya's email; Magoo's eyesight), which makes the Kim response the appropriate one.

It might be argued that even one application of Magoo's knowledge of his margin for error is contradictory. We assume that after his inexact observation, Magoo's total knowledge includes the fact that the clock is not pointing to 00, but nothing stronger. By one application of his margin for error, Magoo comes to know that the clock is not pointing to 01. This conclusion seems to contradict the assumption that the strongest thing Magoo knows is that the clock is not pointing to 00. This apparent contradiction can be easily resolved: Magoo's total *perceptual* knowledge includes the fact that the clock is not pointing to 00, but nothing stronger. By using his margin for error once, Magoo's total knowledge includes the fact that the clock is not pointing to 01. Since this latter knowledge is not merely perceptual (rather it was obtained by combining Magoo's perceptual knowledge with his knowledge of his margin for error), this conclusion does not contradict the fact that Magoo's total *perceptual* knowledge includes the fact that the clock is not pointing to 00, but nothing stronger.

I have emphasized that the margin for error principle associated with the agent's knowledge is restricted to a particular source of information. A defender of Williamson's argument might note that this appears to contradict Williamson's more general commitment to epistemic safety and its relation to inexact knowledge. According to Williamson, inexact knowledge is a widespread phenomenon that arises from our fallible nature. Fallibility affects our perceptual knowledge, our memory, our inferential and our testimonial knowledge. Correspondingly, most, if not all, of our knowledge, is inexact, and is therefore governed by some margin for error. If this is the case (and let us grant for the sake of argument that this is so), then it opens the door for a repeated application of the margin for error reasoning. If every knowledge state is governed by a margin for error, then, in principle, for every state we can reuse the margin for error reasoning. I have argued that margin for error reasoning cannot be repeatedly applied, so there must be something wrong with my argument.

This line of reasoning is worth a closer analysis. First, we should be advised to avoid a simplistic quantifier shift fallacy: the assumption that every state of knowledge is governed by some margin for error principle does not imply that there is a particular margin for error principle that governs every state of knowledge. Even if Magoo's knowledge state *after* hearing Bob's testimony is governed by some margin for error principle, there is no reason to assume it has the form "if you know that the clock is not pointing to p_i, then it is not pointing to $p_{i\pm1}$." The latter margin for error principle, when applied to perception, is motivated by (and can be tested in) experimental psychophysics. It also has intuitive appeal. But what is the structure of margin for error principles for other knowledge states, those that involve a combination of perception, testimony, inference or reflection?

The idea that knowledge, in general, requires margin for error stems from the safety condition on knowledge. According to safety, to be in a state of knowledge requires being safe from error: one could not have been easily wrong in similar cases (Williamson 2000: 147). The case where the clock points to 16 is very similar to the case where the clock points to 17. Suppose that the clock points to 16; in that case Magoo could easily think it points to 17 (given his perceptual limitations). Since Magoo could be easily wrong in a similar case (the 16 case) Magoo does not know that the clock points to 17 when it does, according to safety. The margin for error principle, applied to Magoo's clock perception, can be seen as a concrete example of safety. More generally, the safety requirement can generate a margin for error principle for any knowledge state, once the notion of similarity is made precise in a given context.

When focusing on human perceptual limitations and perceptual knowledge, it is sometimes easy to make precise the relevant notion of similarity: in the clock example, similarity just amounts to physical distance on the face of the clock. It is much harder to make precise the relevant notion of similarity when we consider not perceptual knowledge, but inferential, memory, or testimonial knowledge. I have argued that after taking into account Bob's testimony, Magoo cannot reuse his perceptual margin for error. After taking that into account, Magoo's knowledge about the clock is a mix of perceptual, testimonial and inferential knowledge. Let us grant that this new knowledge state is subject to some margin for error due to safety; what similarity means for that new state is not a simple matter of physical distance anymore.

I am not arguing that it is impossible to try and construct an argument against the KK principle that involves multiple applications of distinct margin for error principles. In such an argument an agent starts in knowledge state 1 and a corresponding margin for error for that state (call it MFE-1). Using that margin for error results in knowledge state 2 and a different margin for error, MFE-2. Applying MFE-2 (assuming that it is known to the agent and KK), results in yet another knowledge state, and so on. This chain of reasoning might result in a contradiction, but this is not what is going on in Williamson's argument (and I doubt that such a chain of reasoning will have the elegance and appeal of Williamson's original presentation.)

I conclude that neither the quiz story nor Magoo's story are inconsistent with the KK principle. In particular, Magoo's situation does not offer a counterexample to the KK principle. In both cases, a contradiction arises if we allow the agents to reuse their knowledge of margin for error. But in both cases, this knowledge can be used only once. The most sensible response to Williamson's contradictory conclusion is not to reject the KK principle, nor is it to reject the margin for error principle, rather it is to reject the repeated application of the latter principle. Although by now quite old, Williamson's argument has generated an entire margin for error framework that is continuously being used within epistemology, and which goes far beyond the KK principle.[4] Since this Williamsonian framework is built on the idea that the margin for error principle can be iterated in an unrestricted fashion, the entire framework needs to be reconsidered, if my argument is correct.

In the remainder of this paper, I wish to compare my analysis to Williamson's margin for error argument with previous critiques. My analysis reveals a fault that is common to both Williamson's original argument and to some of its previous objectors.

Docik and Égré (2009) and Sharon and Spectre (2008) try to block Williamson's argument by distinguishing between two types of knowledge operators, perceptual and reflective. They use propositional epistemic logic to formulate Magoo's knowledge state; in that they merely follow Williamson, who, at various occasions, assumes that propositional epistemic logic can capture reasoning patterns involving margin for error (see e.g. Williamson 2013, 2014). The appeal to epistemic logic is intuitive: after all, Williamson's margin for error is set up as a counter example to the KK principle, which can be expressed in the epistemic logic formula $K\varphi \rightarrow KK\varphi$. Docik and Égré and Sharon and Spectre argue that knowledge of margin for error should be explicated as something like $K_{reflective}(K_{precetual}\neg p_i \rightarrow \neg p_{i+1})$, as opposed to Williamson's simpler $K(K\neg p_i \rightarrow \neg p_{i+1})$ formulation (Williamson 2014). With this bifurcated conception of knowledge they conclude, like me, that knowledge of margin for error cannot be used repeatedly. While I agree that Williamson's argument does not go through under this bifurcated explication, I do not believe that this explication is necessary, nor that it is entirely compelling, for the following reasons.

First, the bifurcation approach does not generalize well. Given the structural similarities between Magoo's story and the teacher's story, we should expect the same diagnosis to both. But the teacher's story does not involve perceptual and reflective knowledge, rather just testimonial knowledge. In the uncoordinated teachers story, there is an apparent tension between the knowledge obtained by testimony from Aya, and that obtained from Bob. While there are principled reasons to distinguish between perceptual and reflective knowledge, there are no principled reasons to distinguish between the testimonial knowledge obtained from Aya and that obtained from Bob. The analogy I draw between Magoo and the teachers strongly suggests that the underlying argument pattern that mo-

[4]See Hawthorne & Magidor (2009), Goodman (2013), Lederman (2018), Salow & Ahmed (2017), Holguín (2019), and Carter (2019), for a few recent and diverse applications of the Williamsonian margin for error framework.

tivates Williamson's story has nothing to do with perceptual and reflective knowledge in particular.[5]

Secondly, a syntactical distinction between different types of knowledge operators rapidly generates complexities that significantly burden the plausibility of the approach. Williamson's argument is directed against the KK principle, and those who criticize the argument might very well wish to defend it. Note, however, that by distinguishing between n types of knowledge operators, one has to consider n^3 syntactically distinct iteration principles of the form $K_{x_h}\varphi \rightarrow K_{x_i}K_{x_j}\varphi$ (where x_h, x_i, x_j are anyone of the n different knowledge operators). In general, the number of epistemic bridge principles connecting the different operators that one has to consider grows polynomially with the number of types of knowledge operators one considers. While it is certainly possible to offer a general theory for typed knowledge operators (and steps towards such theory are made in Docik and Égré (2009)), the assumptions that come with such theory weaken its plausibility as a response to Williamson's argument, and make it easily susceptible to *ad-hoc* accusations (e.g., choosing and rejecting particular bridge principles for the sole reason of blocking Williamson's argument).[6] Williamson's story is by all means quite plausible and simple, and—unlike the bifurcated approach—so is my response to it. While I will not pursue this here, my response is compatible with a defense of the KK principle, *simpliciter*.

Thirdly, and most importantly, I believe that the use of propositional epistemic logic to formulate margin for error reasoning, both in its original version and in its bifurcated form, misses a crucial element of the agents' reasoning, and is potentially more harmful than useful in understanding such cases. The uncoordinated teacher's puzzle (and by analogy, Magoo's puzzle) involves reasoning about the sources of information from which agents obtain knowledge, or, in other words, *how* the agents came to know what they know. In both stories, Bob is offering information about a distinct source of information (Aya's testimony, Magoo's eyesight). Reasoning about Bob's piece of information involves reasoning about the sources of our knowledge. As I have argued, it is only legitimate to rule out a quiz on Tuesday under the assumption that *Aya is the source* of our knowledge that the quiz is not on Monday. Likewise, it is only legitimate to rule out that the hand does not point to 01 under the assumption that *Magoo's eyesight is the source* of Magoo's knowledge that the hand does not point 00. reasoning about margin for error, we crucially take into consideration the source of knowledge, or equivalently, *how* did we

[5]It is clearly formally possible to syntactically distinguish between two types of knowledge operators $K_{testimony:Bob}$, and $K_{testimony:Aya}$ and to reconstruct my argument based on this distinction, but such reconstruction is both unnecessary and obscures the source of my objection.

[6]Different *ad-hoc* worries are addressed in Docik and Égré (2009:19) and in Sharon and Spectre (2008:297).

obtain that knowledge.

Standard epistemic logic is not concerned with the agent's sources of information, or with *how* the agent's knowledge state came about. It aims to model the knowledge state of an agent at a given moment (as a 'freeze-frame' of the epistemic state of the agent). In epistemic logic, $K\varphi$ expresses that the agent knows φ. It remains entirely silent as to *how* the agent came to know φ, or what is the source of that knowledge. Since it remains silent about that matter, one might implicitly regard $K\varphi$ as ignoring *how* the agent came to know φ: that the agent know φ, *no matter how*. But the puzzles we have considered here show that sometimes it does matter how an agent has obtained their knowledge. Likewise, one might erroneously assume that the epistemic logic expression $KK\varphi$, stating that the agent know that they know φ, implies or captures the assumption that the agent knows how they know φ, but this is mistaken. For it is entirely possible that Bob knows that Paris is the capital of France, that he knows that he knows that, and that, by now, he forgot how he knows this fact (or what was the source of that knowledge). So higher order expressions such as $KK\varphi$ are irrelevant to the question how an agent came to know φ.

This is not a criticism of epistemic logic; there are ways of extending epistemic logic to model the progression of an agent's knowledge via different sources.[7] The criticism is this: I have argued that the two puzzles crucially rely on tracking the different sources of the agents' knowledge. Epistemic logic, in its basic, standard form, does not capture such reasoning. Therefore, the use of epistemic logic in such scenarios is, at best, very limited. The bifurcation approach attempts to answer the source problem by assigning different knowledge operators to different sources of information, which are left implicit. But, as I have just argued, such approach raises problems irrelevant to the reasoning pattern at hand. I have argued that we can more easily understand what goes wrong in Williamson's reasoning without an appeal to standard epistemic logic. Arguably, the use of epistemic logic has done more to obscure than to clarify the reasoning patterns I am concerned with.

Both Williamson and its previous objectors seem to ignore an important moral that stories about margin for error reasoning put to light: reasoning about *how* (i.e. via what source) an agent comes to know the things they know is as crucial, while independent from, reasoning just about *what* they know. Moreover, the analysis offered here, if correct, can serve as a cautionary tale for those using epistemic logic in epistemology. The application of epistemic logic in such reasoning, in its simple 'freeze-frame' form, can lead us away from the heart of the epistemological puzzle at hand.

[7]Two such extensions are temporal epistemic logic and dynamic epistemic logic. For attempts to analyze Williamson's margin for error reasoning using dynamic epistemic logic, see Bonnay and Egré (2011) and Cohen (2021).

References

– Bonnay, D. & Egré, P. (2008). Margins for Error in Context. in M. Garcia-Carpintero & M. Kölbel (eds.), *Relative Truth*, pp. 103-127, Oxford University Press.

– Bonnay, D. & Egré, P. (2009). Inexact Knowledge with Introspection. *Journal of Philosophical Logic*. 38 (2), pp. 179-228.

– Bonnay, D., Egré, P. (2011). Knowing One's Limits: An Analysis in Centered Dynamic Epistemic Logic. In: Girard, P., Roy, O., Marion, M. (eds) *Dynamic Formal Epistemology.* Synthese Library, vol 351. Springer, Dordrecht.

– Carter, S. (2019). Higher order ignorance inside the margins. *Philosophical Studies* 176, 1789–1806.

– Cohen, M. (2021) Inexact knolwedge and dynamic introspection. *Synthese*

– Conee, E. (2005). The comforts of home. *Philosophy and Phenomenological Research* 70: 444-451.

– Das, Nilanjan & Salow, Bernhard (2018). Transparency and the KK Principle. *Noûs* 52 (1):3-23.

– Dokic, J., Égré, P.(2009) Margin for error and the transparency of knowledge. *Synthese* 166, 1–20 .

– Dorst, Kevin (2019). Abominable KK Failures. *Mind*.Vol. 128, 512, Pages 1227–1259.

– Dutant, J. (2007). Inexact Knowledge, Margin for Error and Positive Introspection. *Proceedings of Tark XI*.

– Goodman, J. and Salow, B. (2018). Taking a chance on kk. *Philosophical Studies*, 175(1):183–196.

– Greco, Daniel (2014a). Could KK Be OK? *Journal of Philosophy* 111 (4):169-197.

– Greco, Daniel (2014b). Iteration and Fragmentation. *Philosophy and Phenomenological Research*. 88 (1):656-673.

– Greco, D. (2015). Iteration principles in epistemology i: Arguments for. *Philosophy Compass*, 10(11):754– 764.

– Greco, Daniel (2016). Safety, Explanation, Iteration. *Philosophical Issues* 26 (1):187-208.

– Goodman, Jeremy (2013). Inexact Knowledge without Improbable Knowing. *Inquiry: An Interdisciplinary Journal of Philosophy* 56 (1):30-53.

– Halpern, J. Y. (2008). Intransitivity and vagueness. *Review of Symbolic Logic*. 1 (4):530-547.

– Hawthorne, John & Magidor, Ofra (2009). Assertion, Context, and Epistemic Access - ibility. *Mind*. 118 (470):377-397.

– Holguín, Ben (forthcoming). Indicative Conditionals Without Iterative Epistemology. *Noûs*.

– Lederman, Harvey (2018). Uncommon Knowledge. *Mind* 127 (508):1069-1105.

– Neta, R. & Rohrbaugh, G. (2004). Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly* 85 (4):396–406.

– Sharon, A. & Spectre, L. (2008). Mr. Magoo's mistake. *Philosophical Studies* 139 (2):289-306.

– Stalnaker, R. (2015). Luminosity and the KK thesis. in Goldberg, Sanford C. (ed.) *Externalism, Self-Knowledge, and Skepticism: New Essays*. Cambridge University Press.

– Williamson, T. (1992). Inexact knowledge. *Mind* 101 (402):217-242.

– Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.

– Williamson, T. (2013) Gettier Cases in Epistemic Logic, Inquiry, 56:1, 1-14

– Williamson, T. (2014). Very Improbable Knowing. *Erkenntnis* 79 (5):971-999.