

Some Thoughts on the Logic of Imprecise Observation

Alexandru Baltag and Johan van Benthem

Abstract

We propose a dynamic-epistemic analysis of perceptual knowledge placing structured observational events at center stage. Our starting point is the specific entanglement of epistemic accessibility with the relation of ‘closeness’ or similarity, that is claimed in the so-called Margin of Error principle. While we find the standard formulation of this principle to be defective, we think that it contains valuable intuitions for a joint logic of closeness and uncertainty, whose contours we develop. In particular, our analysis explains how imprecise observation can lead to higher-precision knowledge.

1 Introduction

Perceptual knowledge and its inherent imprecision are an essential feature of human cognition and action. In our daily lives, but also with measurement regimes in the exact sciences, each observation comes with a margin of error or imprecision.¹ This ubiquitous imprecision is reflected in the vocabulary of natural language which contains many hedges that serve us well in communication. Moreover, there is the striking fact that by piling up single observations that are imprecise, we can usually reach any desired degree of precision and certainty.

These phenomena give rise to questions for logicians interested in information flow and information-driven agency. What is the nature of observational events with imprecision, and how do they increase our information? How are the ranges of epistemic options that determine knowledge related to the closeness structure of perceptual or other spaces where observation takes place? What is the resulting logic for reasoning with, and about, imprecision? Such questions have been historically neglected by epistemic logicians².

One of the few formal epistemologists who did take seriously the vagueness of perception and put it at the very center of his philosophical reflection is Timothy Williamson. His epistemic view on vagueness and its logic is one of the leading approaches to this topic in the literature, and led him to an extensive investigation of ‘inexact knowledge’ over the years [29, 32], including his celebrated Margin of Error argument purporting to show that epistemic introspection fails for perception-based knowledge.

¹The term ‘imprecision’ seems preferable, as there is no possibly blameworthy ‘error’ involved.

²Closely related topics, such as the Sorites paradox, fuzzy sets, rough sets, many-valued logics, etc., were explored by logicians engaged in studying vagueness and ambiguity in natural language and science, [11].

In this paper, we start with a brief analysis of this argument, also known as the Epistemic Sorites paradox, since it contains valuable triggers for a further study of imprecise observation. Our discussion fits within the tradition of ‘signal-based’ solutions to the paradox, following in the steps of Bennet [4], Mott [21], Halpern [17], Dutant [13] and Spector [26]. Such approaches focus on the fact that knowledge is not uniquely determined by the objective value of the ontic feature under observation, as a prima facie reading of the Epistemic Sorites might suggest, but it also depends on the subjective signals perceived by the observer, or the subjective comparisons made by him against given benchmarks. The conclusion of this analysis is that the Margin-of-Error principle is not tenable in its standard form.

At this stage, we encounter a fork in the road. One further path leads toward the dialectics of epistemological discussion of the Epistemic Sorites, on which there is a rich literature. We might explore to which extent our analysis presupposes epistemic introspection, or whether our rejection of Margin-of-Error rests on even more deeply hidden debatable philosophical assumptions. While these are important issues, we do not pursue them in this paper.

Instead, we choose another road at the fork: a study of what lies implicitly behind the signal-based approach, namely various types of observational events, and the ways in which these events update information. To the best of our knowledge, this dynamic-epistemic road has been taken only once before, in the recent, still unpublished paper Cohen [10]. However, we will give our own version of the epistemic dynamics involved, and with that framework in place, we explore a new logic for reasoning about knowledge and perceptual closeness. In particular, this logic can take on board various modifications of the Margin-of-Error principle that we do think are true as regulating imprecise observation and knowledge.

The eventual test that we see for the dynamic framework of this paper is whether it makes intuitive sense, and generates a logic that can link up with what we take to be the most sophisticated current account of imprecise observation, namely, Measurement Theory.

2 An intuitive problem

It seems that our intuitions about the fundamental imprecision of perceptual evidence are incompatible with the well-known epistemic KK Principle $K\varphi \rightarrow KK\varphi$ saying that agents know that they know when they know (the Positive Introspection law of epistemic logic). The following ingenious example was proposed in [32] as an argument against the KK Principle.

Example 1, Tree in the distance. Looking through the window, Mr Magoo can see a tree far off. Suppose that in fact the tree is 666 inches tall. But, given his poor eyesight and limited ability to judge heights, Mr Magoo cannot tell the tree’s height to the nearest inch just by looking. For all he knows, it can be 665, 666 or 667 inches tall. More generally, he cannot tell apart any height of m inches from a height of $m + 1$ or $m - 1$ inches: these heights are perceptually indistinguishable to him. But on the other hand, Mr Magoo *can* see the tree, so he definitely does know, just by looking, that the tree is *not* 0 inches tall.

A surprising argument The preceding description seems natural and uncontroversial. Moreover, in this scenario, the following *Margin of Error* statement seems quite plausible:

“For all natural numbers m , if the tree is in fact $m + 1$ inches tall, then Mr Magoo does not know that it is not m inches tall”

which we can write formally as

$$p_{m+1} \rightarrow \neg K \neg p_m \text{ holds for all } m \quad (\mathbf{MoE})$$

or equivalently as the implication

$$K \neg p_m \rightarrow \neg p_{m+1}, \text{ for all } m. \quad (1)$$

Moreover, Mr Magoo is aware of the limits of his eyesight, so he *knows* this statement:

$$K(K \neg p_m \rightarrow \neg p_{m+1}), \text{ for all } m. \quad (\mathbf{K-MoE})$$

Then, if we assume Closure of Knowledge (embodied in the distribution law of epistemic logic, whose debatability is not the issue here), a surprising chain of reasoning starts with

$$K K \neg p_m \rightarrow K \neg p_{m+1} \quad (2)$$

If we now apply just one concrete instance of Positive Introspection, we obtain

$$K \neg p_m \rightarrow K \neg p_{m+1}. \quad (3)$$

Continuing by induction, we can then show that

$$K \neg p_0 \rightarrow K \neg p_m \text{ for all } m. \quad (4)$$

But this conclusion seems absurd. Since we agreed that in our situation $K \neg p_0$ holds (since Mr Magoo can see the tree), it follows that the tree can have no height at all for Mr. Magoo. Introspection implies the existence of heightless trees.

This is the gist of Williamson’s argument [32], and his conclusion is that, since **MoE** is correct, the KK principle $K\varphi \rightarrow KK\varphi$ fails for perceptual knowledge. In other words, perception and well-known epistemic principles are at odds. While we have no special wish to defend KK, there is an issue here of whether this particular quandary is convincing.

A spate of follow-up work, such as [28, 25, 10], has tried to elucidate more general philosophical intuitions that might underwrite the Margin of Error principle **MoE**.³ In contrast, other authors [4, 21, 17, 13, 8, 9, 26] have found the Margin of Error principle defective, and proposed alternative analyses in the earlier-mentioned line of carefully analyzing the ‘signals’ provided by imperfect observation. There is a sequence of responses by Williamson to these criticisms, some of which are mentioned below. Still, to us, the signal-based approach captures an important insight about perception, and we now proceed to analyze more precisely what is going on, with the aim of learning more about the logic of perceptual observation.

³One such general intuition is ‘perceptual robustness’ or ‘safety’: if I know a fact based on perception in some world w , then I must also know it in worlds that are close to w .

3 Analyzing observational events

3.1 Margin of Error Principle: Weak and Strong Version

The **MoE** principle seems natural, but let us look at what it says exactly. What is the relevance of the 1-inch differences in the story? The intuition is that an agent’s perception has a minimal margin of error or ‘tolerance’ $\epsilon > 0$. This seems to pose a fundamental limit to her powers of discrimination: no perceptual observation by this agent can have an accuracy that goes beyond ϵ . We can formulate this as a “weak” version of **MoE**:

“If the agent’s minimal margin of error is ϵ , then features of the world whose values differ by less than or equal to ϵ cannot be told apart by the agent by any direct comparison: they are *perceptually indistinguishable*.”

While the weak version seems to us incontrovertible, it is easy to slip from here into the following, much more contentious, *epistemic* statement of a general nature:

“If the agent’s minimal margin of error is ϵ , then for every possible value x of the observed feature, if x is actual, then the agent does not know that the value is not $x + \epsilon$ or $x - \epsilon$ (or anything in between).”

Our first theme is how this ‘strong version’ of **MoE** fares when observations are made, providing a starting point for exploring the role of observations in perceptual knowledge.⁴

3.2 Single Observations

Once he observes the tree, Mr Magoo gains information: he can now exclude some values. If his margin of error is ϵ inches, then in principle he may be able to narrow his estimate to an interval of radius equal to ϵ : this is a most accurate or *best observation* consistent with his perceptual limitations. One way this can happen is if he gives a ‘best guess’ y , a point estimate of the tree’s height as it appears to him⁵. Knowing his own margin of error ϵ , he can then give $[y - \epsilon, y + \epsilon]$ as his best estimate. To be concrete, suppose his point estimate is $y = 665$ inches, and his margin of error at this distance is $\epsilon = 25$ inches. Then he is sure that the height is in the interval $[640, 690]$, but he cannot rule out any of these values.

At this point, we notice a tension. The above scenario complies with the weak version of **MoE**, but it is incompatible with strong **MoE**. This is because, if $m = 666$ then the general version of **MoE** tells us that Mr Magoo cannot know that the tree is not $m + \epsilon = 666 + 25 = 691$ inches. But he *does* know this, since $691 \notin [640, 690]$. Thus, the epistemic situation after just one observation refutes the strong version of **MoE**.

⁴Note that ‘strong’ **MoE** is the general version of Margin-of-Error that underlies Williamson’s argument.

⁵Bennet calls this point estimate y the “precisification”, while Spector [26] identifies it with the “subjective signal” perceived by the agent. In the next sections of this paper, we will use the term “apparent value” for y , distinguishing it from the *actual value* of the feature under observation.

MoE as a ‘Centering Fallacy’ Strong MoE implies that the margin of error interval for the actual value is contained in any measured interval. Hence, even if we allow observations other than point estimates, [33], any *best* observation is an interval centered at the actual value. Say, an agent knows the margin $\epsilon = 25$, and observation told him that $m \in [640, 690]$. He knows this is a best observation, as its radius equals the margin of error. So, **K-MoE** says the actual value is the center $(640 + 690)/2$. This is obviously unwarranted.

One might call this the ‘Centering Fallacy’: the assumption that actual values always lie at, or near, the center of measured intervals. Clearly, this symmetry assumption is not generally warranted in statistics: conclusions from measurement about the location of the actual value come with probability only, and they also depend on the type of probability distribution assumed. In our qualitative setting, the same is true: as we have seen, the actual value may lie close to the border of the measured interval. In fact, this is one of the crucial features that accounts for the beneficial epistemic effects of repeated observations.

The agent, or a device? The use of measurement theory in the preceding analysis suggests that one can trade perceptual imprecision for measurement imprecision. Margin of Error is not only about deficiencies of epistemic agents: equally well, the agent’s perception may be totally accurate, while some measuring device is to blame.

Example 2, The scale. I step on a digital scale to measure my weight: it shows exactly 82 kilograms. But I know the scale has a certain tolerance, say 0.1 kilograms: if I measure the same weight twice in a row, sometimes the scale gives results that differ by at most 0.1 kg. So, after I measure my weight once, I conclude that the weight is 82 ± 0.1 , i.e. in the range $[81.9, 82.1]$. But the same principles MoE and K-MoE can be stated in this case, such as: “I know (given the scale’s error margin) that, if my weight is $i + 0.1$ kilograms, then using the scale I cannot know that it is not in fact i kilograms.” Formally, this is

$$K(p_{i+0.1} \rightarrow \neg K\neg p_i), \text{ or equivalently, } K(K\neg p_i \rightarrow \neg p_{i+0.1}) \text{ for all } i.$$

Now, for any introspective agent who knows himself to have non-zero weight, again the conclusion follows that he has no weight at all. Thus, the existence of an inaccurate apparatus would seem to imply that no perceptually perfect agent using it can ever be introspective.

But the earlier counter-example reappears in this new setting. Once again, the Centering Fallacy is the culprit: if say, the correct value of my actual weight is in fact 81.9 kg, then the scale may just show 82, as in the story above. Given the scale’s tolerance, I then know that my weight is in the range $[81.9, 82.1]$, and thus I am able to exclude 81.8, despite this being close enough (within the margin 0.1) to the actual weight.

These points about observation and centering are not new. In various forms, they were raised by Mott [21], Halpern [17], Spector [26], and others. Now there is an issue of whether these outcomes stem essentially from our simple view of observational events.

Observations with variable margins In [32, 33], Williamson has emphasized *higher-level imprecision*, which characterizes inexact knowledge (obtained by perception, memory, etc.), in contrast with exact, but imprecise knowledge (obtained by measurements). According to this view, the above analysis only applies to cases of the second type, but not to the first.

One’s knowledge of one’s perceptual limitations is inexact, and itself subject to the Margin of Error principle: analyses assuming that individuals know the interval estimate fall short.

There are many issues of how to phrase these intuitions precisely, but certainly, it makes sense to consider perceptual events with arbitrarily large errors.⁶ Still, without analyzing observations with varying accuracy in detail here, we merely note that **MoE** still fails.

First, in discrete settings, the two approaches are equivalent, and all earlier points apply. In infinite continuous settings, there may indeed be no best observation, but despite the errors, we can approximate the actual value with indefinite accuracy by using better and better observations. The important topic of repeated observations will return below.⁷

Are perceptions radically different from measurements? Although we are sympathetic to the idea that there are various types of knowledge, with different properties depending on their sources, we see no fundamental barrier between inexact informal estimates and imprecise results of measurements. The reality is rather gradual progression from rough estimates by eye (based on mental comparisons) to a more systematic use of *benchmarks*, and then all the way to scales and other measurement devices. So, if the **MoE** principle were valid for inexact human perception, it would also apply to the knowledge obtained from imprecise measurements. The second is just a systematic elaboration of the first.

3.3 Benchmarks

The role of benchmarks is worth high-lighting, since they play an important role in many natural forms of perception. We show this in some variations on the earlier scenario of Mr Magoo seeing a tree, but other cases would do just as well: our points are general, and not exclusively addressed to the dialectics of the earlier epistemological argument.

Example 1, revisited. Suppose the tree is indeed 666 inches tall. Does Mr Magoo know that it is not 639 inches tall, as claimed in our ‘best observation’ analysis? This depends on what heights are available to him as benchmarks. Suppose he also sees another tree, 639 inches tall. He perceptually compares the two, and as the difference is larger than his margin of error, he can distinguish them. Alternatively, there is no other tree around, but Mr Magoo has a clear and distinct memory of a past tree seen from the same distance, which was in fact 639 inches tall. The memory is so vivid that he can easily compare the two perceptions mentally, tell them apart, and conclude that the observed tree is not 639 inches tall.

⁶In [31], such scenarios are described using ‘variable margin frames’. Then P is known in w iff there exists some $\delta > 0$ such that P is true in all worlds within a $\epsilon + \delta$ distance from w . In such a model, a margin ϵ wide is just not enough to ensure the ‘safety’ needed for knowledge, but anything wider will do. So there is no notion of ‘best observation’, but only better and better observations with smaller and smaller δ ’s.

⁷Let $\delta > 0$ be any desired degree of accuracy. In the continuous case, there is an observation accurate enough to narrow down to an interval of radius $\epsilon + \delta$. Let y be the center of this interval. Suppose the actual world x differs from y by more than the quantity δ ; then the modified **MoE** (together with his knowledge of the minimal margin ϵ) tells him that the values $x \pm \epsilon$ should be epistemically possible. But (since $|y - x| \geq \delta$), at least one of these values $x \pm \epsilon$ falls outside the interval estimate $(y - \epsilon - \delta, y + \epsilon + \delta)$, and it is thus epistemically impossible. The agent can conclude that the actual world x must be within a distance δ from y : he has pinpointed the real value with an accuracy better than his margin of error.

Example 1, digitalized and synchronized. Mr Magoo is given a box of 5000 photos of trees of all heights, photographed from the same distance, and arranged in a sequence in such a way that each successive ones differ by an inch or so, and hence they are so close that they are perceptually indistinguishable. (Recall that his accuracy is 25 inches from this distance, so two trees are similar if they differ by at most 25 inches.) He currently looks at the photo of a 666-inch high tree. He then sets it side by side with another photo, far away in the sequence, that happens to be of a tree exactly 639 inches high. He can easily tell them apart: so now he knows that the first (666-inch tall) tree is not 639 inches tall.

It might be objected that the original scenario of seeing the tree involved no explicit benchmarks, and was perhaps just about raw impressions or sense data. But it is hard to understand **MoE**, and the successive steps in the earlier argument, without thinking of comparisons and benchmarks. More generally, all observations tend to be *structured*: there is structure in the space of values, and there is structure to basic acts of observation.

From comparisons to measurements We can now link the earlier-mentioned ‘inexact’ knowledge based on rough perceptual comparisons to exact but imprecise knowledge based on measurements, and show how one leads to the other. In measurement, we have explicitly given benchmarks, arranged regularly in a scale. One looks for some benchmark y such that y is indistinguishable from the actual value x of the feature under observation. Several such matching items y may exist, but we can assume that each observation produces at most one salient match y . This y may be called the appearance or the *apparent value* of the observed feature. After the observation, a world is epistemically possible only if the given feature in that world is some x' indiscernible from the apparent value y .

But perception need not be quantitative. We now turn to single or repeated observations without numerical point or interval estimates. With such qualitative events, our analysis of **MoE** becomes more simple and compelling, without depending on a ‘best estimate’ version.

3.4 Qualitative Perceptual Comparisons

No quantitative notion of distance, intervals, signals or point estimates need to be assumed in a general account of perception. All we need is a notion of ‘closeness’ between worlds, a notion of *epistemic possibility*, and acts of *perceptual comparison*, as the primary observational basis for rough, inexact perceptual judgments. We develop this more general perspective in terms of a few basic general notions, providing a concrete illustration afterwards.⁸

Closeness Perception takes place in spaces with a natural binary relation \approx of *similarity* (or ‘closeness’) between possible worlds, or values of an observed feature. So far, these values were real numbers or objects in some other metric space, where $x \approx y$ was given by $|x - y| \leq \epsilon$ (or maybe $|x - y| < \epsilon$), for some margin $\epsilon > 0$. But we can be more general: any *reflexive and symmetric* relation \approx will do. We think of this similarity relation as a qualitative measure of how ‘close’ two worlds (or values) are to each other, in some objective sense.

⁸The numbers in the qualitative examples to come are essentially just names of possible worlds.

Knowledge and epistemic possibility We also assume a notion of factual perceptual knowledge K . This comes with an *epistemic accessibility* relation \sim (also called epistemic possibility, or uncertainty), understood as follows: world s is epistemically accessible from world w iff, whenever w is actually the case, the agent does not know that s is not the case. This relation must be reflexive (by factivity), but nothing else needs to be assumed about it, making our discussion independent from epistemic assumptions of introspection.

Known margin: ability to compare The qualitative analogue of the assumption that the margin of error is known is that agent has the ability to *compare* two possible states, or values of an observed feature, whenever these values are presented to him, and to decide if they are “similar” or not. In other words, the similarity relation \approx is known to the agent.

Closeness as perceptual indistinguishability In this setting, we can call two states, or values of an observed feature, *perceptually indistinguishable* iff the observer has the same subjective impression when observing w in isolation as when observing w' in isolation, or when the comparison set is just $\{w, w'\}$. Thus, the agent cannot tell the difference by comparing them directly only against each other, without comparison with a third state.

If an agent’s powers of discrimination are limited, closeness implies perceptual indistinguishability. The converse may fail in skeptical scenarios such as Brain-in-a-Vat, or when the actual error is much larger than the minimum margin. But our focus will be on “normal”, non-skeptical scenarios of perception, and then the converse implication holds, too.

Does closeness imply possibility? Prima facie, it might seem reasonable to go further, and also identify the closeness relation \approx with the epistemic accessibility relation \sim , especially in non-skeptical scenarios of high accuracy. Or at least, one might propose a similar implication to the one claimed just now for perceptual indistinguishability, namely:

‘if two states are close, then they are epistemically accessible from each other’.

In fact, this is precisely the semantic content of the general Margin of Error principle, in its qualitative version – for instance, when viewed as an axiom of modal logic. But this seems premature. As we will see, epistemic accessibility is *not* based only on a simple comparison between the two states in isolation, but on a coherent set of comparisons and assessments against the available benchmarks. Assessing whether two states w and w' are epistemically indistinguishable, even with exclusively perceptual knowledge, depends on the full *information state* of the observer. This information state may be determined by many factors: the nature of the current observation, memory of past observations, and the availability of other states w'' that play the role of yardsticks or benchmarks for the current perception.

These notions underlie qualitative scenarios without any essential use of numbers.

Example 3, Color cards. I have cards in colors c_1, \dots, c_{100} , each of a shade close to the next shade, perceptually indistinguishable: if I compare *only* cards c_n and c_{n+1} side by side, I cannot tell the difference. But I can of course tell the difference between c_1 and c_{100} . So there is a number $k \geq 1$ such that any c_n is indistinguishable from c_{n+1}, \dots, c_{n+k} (by direct one-by-one comparisons), but distinguishable from c_{n+k+1} : this k is my *accuracy level*. Say,

$k = 4$. Then I can definitely tell apart cards c_n, c_m when compared side by side if and only if $|m - n| > 4$, while I cannot tell the difference when comparing cards with $|m - n| \leq 4$.

The reader can easily use examples like this to re-check all our earlier points.

3.5 Repeated Observations

Precise knowledge from imprecise observations So far, our scenarios mostly concerned one single observation. But the surprising fact about imprecise observations is that, when *repeated*, they may well pinpoint the actual world with high accuracy. This is one more way of seeing that closeness and epistemic uncertainty can easily come apart.

Example 1, enhanced by a third comparison. Mr Magoo looks at two successive photos in the sequence, one of a 666-inch high tree and the other one of a 665-inch high tree. He obviously cannot tell them apart by a direct comparison: they are so close as to be perceptually indistinguishable. But next, he looks through the sequence of other photos, and he stumbles upon one of a 640-inch high tree. He compares it side by side with the second photo (of the 665-inch tree): they seem to match, the difference of $665 - 640 = 25$ inches is too small for him to tell the difference. He then compares the third photo against the first (of the 666-inch tree): now there is a mis-match, since the difference of $666 - 640 = 26$ inches is large enough to be (barely) discernible. Mr Magoo has broken his own margin of error: he now knows how to tell apart two trees differing only by one inch.

Likewise, in the qualitative scenario of the Color Cards, I can use observations up to my accuracy level to cut down my information state to subintervals of $\{1, \dots, 100\}$.

Example 4, Repeated imprecise observations. Suppose we observe discrete values with a margin of error of 1, and our first observation has yielded the interval $\{3, 4, 5\}$. If we now observe a new interval, say, $\{2, 3, 4\}$, then our new information state is the intersection of these intervals: $\{3, 4\}$. After this, it depends. If we observe, say, what appears to be 3, then the interval for this is $\{2, 3, 4\}$, and no gain no more information than $\{3, 4\}$. But if we observe, say, what appears to be 2, then the interval for this is $\{1, 2, 3\}$, and by intersection with $\{3, 4\}$, we obtain $\{3\}$. We have found the actual world.

What is interesting here is that we could not improve from $\{3, 4\}$ by observing one of these two points themselves. The separation has to come by observing some value further out. As in soccer, scores often come from the wings, not through the center.

Whether we find the actual world here depends on what observations are available. There are interesting further issues here, if we can freely choose observations, or have another person elicit them. Suppose our repertoire is all direct comparisons with benchmarks as discussed before. How many imprecise observations are needed to get to the truth? This is a matter of information theory: we have a total interval of values $\{0, \dots, n\}$, and subintervals of size k matching the precision level of available measurements. The comparison tells us either that the object is similar to the benchmark, in which case we are inside an interval of length k , or that it is dissimilar, and we are in a set of size $n - k$. This gives our repertoire of *questions* that we can ask, and now we can calculate how many questions are needed in given cases.

Example 5, Best questions. As an illustration, suppose that we have obtained a positive match with object 4 at precision level ± 2 , resulting in the interval $\{2, 3, 4, 5, 6\}$. We might now try to match with an object inside, say 5, but in the worst case, this gives a positive match and we have 5 remaining possibilities. From an information-theoretic point of view, it is best to choose observations that cut the current interval into roughly halves – corresponding to the standard \log_2 bit measure of information content. Significantly, and this reflects an earlier point, what we need for this optimal information flow are comparisons with objects *outside* of the interval. If we compare with the outside object 7, then a positive match will leave the subinterval $\{5, 6\}$, while a mismatch will leave $\{2, 3, 4\}$.

Precise results exist for choosing the optimal questions toward the truth for a given interval size and observation quality, but our point here is just that imprecise observation can lead to precise knowledge in ways that can be analyzed completely.

We end by stating some general features behind the preceding points.

Comparisons and benchmarks revisited As already mentioned, perceptual observations are not neutral, context-free interactions. They come in a context of available *benchmarks*, and given an unknown object and a benchmark, we can elicit a *comparison*. Suppose I observe some feature of the world (say, the color of a card), either by direct perception or using a measurement device. The actual value $x = c_n$ of that feature is the *object* of the observation. Now, with one color or a number of colors in mind as benchmarks, I try to identify which of these (if any) corresponds to the card. In the simplest case with only one benchmark y , I merely check whether the actual feature is indistinguishable from the benchmark: we can get either a positive result (a ‘match’), or a negative result (a ‘mismatch’). In the general case, there may be more such steps before we have located the observed feature.

Information state after observation The result of a comparative observation is captured by the agent’s *information state*, which is usually represented as a *set* of possible worlds. In simple scenarios of perceptual knowledge, this is often just a set of possible values of the relevant observed feature. In case of a match with benchmark y , the resulting information state is $\{x' : x' \approx y\}$: this reflects the fact that, after a positive observation, a world is epistemically possible only if the given feature in that world is some x' similar to y . For instance, this set can be an open or closed interval $[y - \varepsilon, y + \varepsilon]$, or even just a discrete set of possible values y_1, \dots, y_k . In case of a mismatch, the result is negative: we learn that the observed state x can be distinguished from y , so the information state is $\{x' : x' \not\approx y\}$.

Having discussed imprecise observations in a semi-formal style, in the following section, we show how the preceding examples can be modeled in a standard formal style. This is not a matter of pedantry: we see this formal turn as a stepping stone toward a more general logical theory of observation and measurement. As a side benefit, it will also allow us to find more tenable versions of **MoE** that survive the counter-examples in this section.

4 Imprecise observation with known margin

The simple models that follow here were implicit in our discussion so far, but we now highlight their formal structure. Our first analysis uses well-known semantic tools: standard epistemic logic plus a dynamic logic *PAL* of public announcement or public observation [23, 12, 5, 3]. In the next section, we enrich the models to circumscribe the available observations.

Epistemic similarity models We use *epistemic perceptual models*, i.e. structures

$$\mathbf{M} = (W, \approx, \sim, V),$$

where W is a set of *possible worlds*, \approx is a reflexive and symmetric relation of *closeness* (or “similarity”), \sim is a reflexive relation of *epistemic accessibility* (representing “possibility” or “uncertainty”), and V is a valuation map, associating some set of worlds $V(p) \subseteq W$ to each atomic sentence $p \in P$ from a given set P . Here the atomic sentences denote ontic (i.e., non-epistemic) facts, so the valuation gives us a description of each world’s ontic features.

In using these models, we ignore deceptive perception. This excludes Brain-in-a-Vat scenarios, Gettier cases such as fake trees having the same appearance as a real 666-inch tree, or extremely far-off colors outside the visible spectrum. Our emphasis in analyzing observation is on ordinary settings where the closeness relation $w \approx w'$ matches perceptual indistinguishability of agents – defined as before: the values x, x' of the perceived feature in the two worlds are indistinguishable by a direct perceptual comparison of these two values, in the absence of any third benchmark. In contrast with closeness \approx , the epistemic possibility relation \sim is determined by the results of prior and posterior observations.⁹

It is convenient to introduce a bimodal logic over these models, with two modalities:

| | | |
|-----------|------------|--|
| knowledge | $K\varphi$ | what is true in all epistemic alternatives |
| closeness | $C\varphi$ | what is true in all similar states |

$C\varphi$ says that φ is true in all worlds that are similar to the actual world, while $K\varphi$ says that is true in all the worlds that are compatible with all the agent’s perceptual observations. It is useful to also introduce the dual ‘possibilistic’ modalities: $\langle C \rangle \varphi$, defined as $\neg C \neg \varphi$, says that the actual world is similar to some φ -world; while $\langle K \rangle \varphi$, defined as $\neg K \neg \varphi$, says that φ is epistemically possible: for all the agent knows, φ might well be the case.

Special case: epistemic metric models A special case close to the way we have discussed our earlier examples occurs when similarity is given in terms of closeness in a metric space. An *epistemic metric model* is a structure $\mathbf{M} = (W, d, \sim, V, \varepsilon)$, where W , \sim and V are as above, while d is a metric¹⁰ on W , and $\varepsilon \geq 0$ is a given real number (the “accuracy level”). This gives rise to an epistemic similarity model, by putting:

$$w \approx w' \text{ iff } d(w, w') \leq \varepsilon.$$

⁹A more general kind of models, that we would prefer eventually since it prejudices still fewer issues, would use *three* independent relations: closeness \approx , epistemic possibility \sim , and perceptual indistinguishability.

¹⁰A symmetric function $d : W \times W \rightarrow [0, \infty)$ s.t.: $d(w, w') = 0$ iff $w = w'$, and $d(w, w') \leq d(w, s) + d(s, w')$.

Mathematical digression: distances in similarity models A rough notion of distance d between worlds can be defined in *any* epistemic similarity model, by putting:

$$d(w, w') := \text{the smallest integer } n \text{ s.t. there is a chain } w = w_0 \approx w_1 \approx \dots \approx w_n = w',$$

if such a (finite) chain exists, and otherwise:

$$d(w, w') := \infty.$$

While d need not be a metric, it is, if infinite distances are excluded. Call an epistemic similarity model *Archimedean* if $d(w, w') < \infty$ for all worlds $w, w' \in W$. Thus, the transitive closure \approx^* of similarity is the universal relation (every two worlds are related by some finite chain of close worlds). Although not all models are Archimedean, our intended models will typically be so, since Euclidean space satisfies Archimedes' Postulate that small non-zero numbers, if repeated a large number of times, can add up to arbitrarily large quantities.

Now we return to an earlier example, and represent it formally.

Example 3, Color cards, revisited: before the observation. A model for this example (with accuracy level $k = 4$) can be obtained by taking atomic sentences $P = \{c_i : 1 \leq i \leq N\}$, where c_i says that “the real color of the new card is c_i ”. We can take any set W of possible worlds, together with a valuation $V : P \rightarrow \mathcal{P}(W)$ that ensures that every color is in principle possible (i.e. $V(c_i) \neq \emptyset$ for all i), but that in any world the newly chosen card has a unique actual color (i.e. $\forall w \in W \exists! c_i \in P \text{ s.t. } w \in V(c_i)$). The second assumption allows us to denote by $c(w)$ the color of the chosen card in world w .

For the moment, we will use systematic ambiguity to simplify the notation, by *identifying the worlds w with the corresponding colors $c(w)$* , and thus taking $W = P = \{c_i : 1 \leq i \leq N\}$ and $V(c_i) = \{c_i\}$.¹¹ We define similarity $\approx \subseteq W \times W$ between worlds/colors, using the assumption that the accuracy level is $k = 4$:

$$c_n \approx c_m \text{ iff } |n - m| \leq 4.$$

Finally, *before* the agent looks at the new card, no observation of the real world has yet been made, so the epistemic possibility relation is the universal relation at this stage:

$$c \sim c' \text{ holds for all } c, c' \in W.$$

MoE semantically Next, in this setting, we can determine the semantic content of the **MoE** principle stated syntactically in the original formulation. It says that for any world, any close world is epistemically possible, or more formally:

$$c \approx c' \Rightarrow c \sim c'$$

which is easily seen to be equivalent to

$$c \rightarrow \neg K \neg c', \text{ for all } c \approx c'.$$

¹¹We later drop this identification when keeping track of a world's history through product update.

This can be made precise by standard modal frame correspondence on the form of **MoE**, but the equivalence should be clear by itself. In our model of the situation before the observation, this implication holds indeed for all colors, for the trivial reason that the epistemic relation is the universal one. However, soon we will be questioning this implication.

The current setting supports new notions beyond the ones we employed so far.

Absolute indistinguishability As we saw, perceptual indistinguishability is only relative: one *may* be able to tell apart similar states by observing them against a third state as a benchmark. But in principle, there can be states that cannot be told apart in this way. The relation of *absolute indistinguishability* can be defined by putting, for all worlds $x, y \in W$:

$$x \cong y \text{ iff } \forall z \in W (x \approx z \leftrightarrow y \approx z).$$

In all our examples so far, this relation \cong is trivial: only identical states were absolutely indistinguishable, as is the case in all metric models based on Euclidean spaces \mathbb{R}^n . But there exist epistemic similarity models with distinct states that are absolutely indistinguishable.

Unlike the dubious semantic condition $x \approx y \rightarrow x \sim y$ underling the **MoE** principle, its analogue in terms of our new absolute indistinguishability is a natural requirement (though not one that is valid in all models): it is consistent to add the semantic condition

$$x \cong y \rightarrow x \sim y$$

to our definition of similarity models, without trivializing knowledge.

Next we turn from the statics to the dynamics of observation.

Events with preconditions When perceptual events take place, the epistemic models must be *updated* to represent the new state of information, in line with the standard semantic view of information growth. So, we must identify the relevant perceptual events of observation for our scenario. To keep things concrete, we continue with our cards example.

As a first stab, we choose to model acts of observing the new card by comparing its color with some randomly chosen card from the deck. There are two possibilities: either we find a match (i.e., we cannot tell the difference between the two), or a mismatch (we can definitely tell them apart). These correspond to two different observing events:

| | | |
|-----------------|-------------|--|
| positive result | $See : i$ | the observed color matches c_i |
| negative result | $See^- : i$ | the observed color is distinguishable from c_i |

We can model these observational events in *PAL* style as *updates* (also called “public announcements”, though there is nothing ‘public’ to their use here.) Furthermore, these are not updates $!c_i, !(\neg c_i)$ involving just the real values, as one might think at first sight: we need to incorporate the margin for error. So the precondition of $See : i$ is actually

$$pre(See : i) = \langle C \rangle c_i = \bigvee \{c \in W : c \approx c_i\} = c_{i-4} \vee c_{i-3} \vee \dots \vee c_i \vee \dots \vee c_{i+4},$$

and similarly, the precondition for mismatches runs like this:

$$pre(See^- : i) = C_{\neg c_i} = \bigvee \{c \in W : c \not\approx c_i\} = c_0 \vee \dots c_{i-5} \vee c_{i+5} \vee \dots$$

In terms of the formal language that we introduced, we thus identify the event $See : i$ with an update $!\langle C \rangle c_i$, and the event $See^- : i$ with an update $!C_{\neg c_i}$.

Updating models by observation To represent the updates induced by observations, we choose a *link-cutting* approach (rather than the more common worlds-deletion of *PAL*), since it is convenient to keep around worlds even after they become epistemically impossible from the perspective of the actual world.¹² A link-cutting update $!\varphi$ changes a current model \mathbf{M} simply by *disconnecting all epistemic accessibility relations* between φ -worlds and non- φ -worlds. If $\equiv_{\mathbf{M}}^{\varphi}$ stands for “agreement on the truth-value of φ in \mathbf{M} ”, then in the updated $\mathbf{M}|\varphi$, everything stays the same except for the new epistemic accessibility relation:

$$c \sim^{\varphi} c' \quad \text{iff} \quad c \sim c' \text{ and } c \equiv_{\mathbf{M}}^{\varphi} c'$$

Remark Note that, in this sense, the updated model $\mathbf{M}|See : i$ is exactly the same as the updated model $\mathbf{M}|See^- : i$, we just disconnect all the epistemic relations between worlds $c \approx c_i$ and worlds $c \not\approx c_i$. So in a sense, the two events are ‘update-equivalent’. The only difference comes from which of the two zones contains the real world. But informationally speaking, as in our earlier discussion of repeated observations, this makes a big difference: $See : i$ narrows down the epistemic range (of worlds accessible from the actual one) to no more than 9 possible worlds. In contrast, $See^- : i$ leaves open no less than $100-9=91$ epistemic possibilities! This reflects the fact that $See : i$ is a maximally accurate observation: its error margin is as small as the agent’s inherent level of perceptual acuity (hence, the apparent value c_i is really a best guess), while a negative observation $See^- : i$ is less informative.

What becomes of the original problem This semantic modeling makes our analysis in Section 2 precise. We can check concretely what is true and false about relevant epistemic assertions, contrasting the initial model with the new models that result after having made successive observations. Since we tend to start with the universal epistemic relation, all subsequent updated epistemic accessibility relations will be equivalence relations in *S5*-style: so not only the *KK* principle of Positive Introspection is true throughout, but all the laws of the modal logic *S5*, including even Negative Introspection. However, this is not an inevitable consequence. For instance, if we start with a transitive non-symmetric relation \sim , we will only get the modal logic *S4*. Since we have no strong opinion on this issue, the only constraint imposed on our epistemic relation \sim in our epistemic similarity models is reflexivity.

Finally, returning to the original issue at the start of this paper, after *any* single observation of the form $See : i$ or $See^- : i$, some instances of the Margin-of-Error principle **MoE** will fail in some worlds. For example, after $See^- : 1$, $K_{\neg c_5} \rightarrow \neg c_6$ is false at world c_6 , since c_5 has been eliminated although it is perceptually similar to c_6 . This can be seen even more

¹²This feature allows us to reason counterfactually about the perceptual similarities between colors that have been eliminated, as distinguished from the new card.

easily for the above semantic version of **MoE**: any observation disconnects some epistemic links between some perceptually indistinguishable states. So, **MoE** is not a generally valid principle for perception, and imperfect observation is perfectly compatible with assuming epistemic introspection, or even stronger principles such as the negative introspection of *S5*.

5 Richer modeling: observational event models

Until now, observational events were introduced in an ‘external’ manner, so to speak, in the form of specific updates and their corresponding dynamic modalities. But we may want to be more explicit about the agent’s ‘potential knowledge’: that is, the total set of observations which are actually available to her. We can do this by specifying an ‘event model’, in a style that is by now standard in Dynamic Epistemic Logic [2, 12, 5, 3]. This allows us to generalize, or when needed, to restrict the framework introduced in the previous section, and to highlight additional features of perceptual observation.

Event models In addition to the static model **M**, we are now given an *event model* **E** = (E, pre) , consisting of a set E of *observational events*, together with a *precondition map* that associates to each event $e \in E$ some sentence $pre(e)$, describing the condition of possibility of the corresponding event: the observation e can happen in a world $w \in W$ iff the precondition pre_e is true in that world. In standard Dynamic Epistemic Logic, event models come with epistemic possibility relations on events in **E**, similar to the epistemic relations on states in **M**. But in our single-agent, ‘hard information’ framework, this is not necessary: we assume that the observing agent knows what observation is performing. So our dynamic semantics is an obvious generalization of the one in the previous section: any event $e \in E$ will be treated as an update $!pre(e)$, that cuts all epistemic links between worlds which satisfy the precondition $pre(e)$ and worlds which do not satisfy it. Hence, in this simple single-agent setting, $pre(e)$ captures the *information* carried by observation e .¹³

Fewer benchmarks: restricting the event model In terms of observational event models, earlier informal points can now be made precise. For instance, in our Example 2 of the color cards, the event model will consist of all possible updates of the form $See : i$ and $See^- : i$. This incorporates the assumption underlying the card scenario that all the cards in the deck are available as benchmarks. If instead one has fewer available benchmarks (as in our other examples), the available observation events are fewer, and the agent’s potential knowledge is more limited. We can formalize this by restricting the event model **E**, to allow only some of the events $See : i, See^- : i$: that is, only some c_i are available as benchmarks.

More general comparison events Our observational events so far cover only two extreme cases: $See : i$ are the most accurate observations, in which the error matches the agent’s inherent accuracy level, while $See^- : i$ are the most inaccurate ones (except for no-observation actions *skip*). But we can also model observations that are *less accurate* than the inherent

¹³If we want to insist that these observations are *external perceptions* (that observe external ontic features of the world, rather than self-perception by introspective acts), then we may restrict the preconditions $pre(e)$ to non-epistemic formulas (which do not contain knowledge operators K).

accuracy, still centered at some apparent value i . For any integer $n \geq 1$, we can have an update $See_n : i$ with $pre(See_n : i) = \langle C \rangle^n c_i = \bigvee \{c \in W : d(c, c_i) \leq n\}$ the closed ball or interval of radius n centered at i . The negative observation $See_n^- : i$ is defined likewise. In event models \mathbf{E} such observations generalize the metric view of measurement results as closed or open intervals $[y - \epsilon, y + \epsilon]$. Another option are *relative comparison events* not centered at a point estimate, say, acts $See : <_i j$ of seeing that the observed value is closer to a given benchmark i than is another benchmark j . The precondition is $pre(See : <_i j) = \bigvee \{c \in W : d(c, c_i) < d(c_j, c_i)\}$. This can represent perceptions like those suggested in [33], e.g. “I know by eye that the building opposite is less than 200 feet high without estimating its height” becomes $See : <_0 200$.

Indistinguishability with respect to a set of observations Any set $F \subseteq E$ of observations (where E is the set of all observational events of our fixed event model \mathbf{E}) can be thought of as a generalization of the notion of ‘scale’ (or set of benchmarks): it specifies an “observational context”, in which the available observations are restricted to the events in E . With each such context F , comes a notion of *observational indistinguishability* \cong_F , defined as inseparability of the two states by observations in F :

$$x \cong_F y \text{ iff } \forall e \in F (x \in pre(e) \leftrightarrow y \in pre(e)).$$

A dynamic-epistemic take on absolute and relative indistinguishability We can now understand perceptual indistinguishability \approx as a dynamic-epistemic notion, namely, a kind of observational indistinguishability against a *variable context* (given by observing either or both of the two states, but no others). More precisely, two states x, y are perceptually indistinguishable iff the agent might make observation $See : x$ in the state y (and vice-versa):

$$x \approx y \text{ iff } x \cong_{\{See:x, See:y\}} y.$$

Similarly, *absolute indistinguishability* \cong can be understood as observational indistinguishability $\cong_{See(W)}$ wrt the set $See(W) = \{See : x | x \in W\}$ of *all* possible observations against all possible benchmarks. In this way, we find that absolute indistinguishability, that we introduced as a static notion, also functions as a dynamic-epistemic notion: it sets the absolute limits of what can be in principle discerned by perceptual observations.

A natural next theme is logical calculi for reasoning about the preceding models, and what further insight these offer. So far, we had an epistemic similarity logic with a knowledge modality K and a similarity modality C . In the following section, we go further.

6 Reasoning with knowledge and similarity

In this section, we explore some uses and properties of logics over our epistemic similarity models. These logics have a double function. They allow us to explore what is valuable about Margin of Error intuitions and our earlier special scenarios, but they also yield a platform for linking up with other mathematical accounts of observation and measurement.

6.1 Logics of imprecise observations

There are various logical languages for reasoning about knowledge, similarity and observations. For the basic language introduced above, a complete axiomatization is easily available.

A complete logic The complete dynamic-epistemic over arbitrary epistemic similarity models and arbitrary event models consists of the following:

- the axioms and rules of the standard KT proof system for the epistemic modality K , saying that knowledge is closed under logical consequence, all validities are themselves known, and knowledge is factive,
- the axioms and rules of the standard KTB proof system for the closeness modality C , reflecting the fact that perceptual similarity is reflexive and symmetric,
- with atoms¹⁴ of the form c_i (for various values of a given variable, e.g. ‘color’, ‘weight’, etc), and a conventional accuracy level $k = 1$, we also have special axioms

$$c_i \rightarrow \langle C \rangle c_j, \quad \text{for } |i - j| \leq 1,$$

and this generalizes straightforwardly to other atoms and accuracy levels,

- all standard instances of the usual recursion axioms of link-cutting public announcement for all observational events $e \in E$, treated as updates $!pre(e)$.

Of course, since we have shown that both Positive and Negative Introspection principles are consistent with this setting, one can restrict our class of models to the ones in which \sim is transitive, or even an equivalence relation, depending on one’s philosophical preferences, and then get richer axioms for knowledge (e.g., the modal systems $S4$ or $S5$).

Enriched epistemic similarity logics We can also extend the modal C, K language to reflect further constraints on models. For instance, the preceding logic has no meaningful interaction between C and K : $K\varphi$ does not imply $C\varphi$ or vice versa, as we rejected both implications in our analysis of observation. However, if our models are Archimedean in the sense discussed above, this does impose a connection between the relations \approx and \sim : namely, that \sim must be included in the transitive closure \approx^* . To capture this, we enrich the earlier syntax with *iteration* $C^*\varphi$ (the modality for the transitive closure of \approx), in which case the Archimedean condition gives us the following validity in the logic:

$$C^*\varphi \rightarrow K\varphi.$$

It is also natural to add to the syntax *dynamic modalities* $[See : i]\varphi$, $[See^- : i]\varphi$ for our basic positive/negative observations, or more generally formulas $[e]\varphi$, for observational events $e \in E$, saying that “after observation e , the statement φ will be true”. The dual possibilistic

¹⁴In terms of modal syntax, the variables c_i, c_j, x, y serve as so-called “nominals” here.

modality $\langle e \rangle \varphi$, defined as $\neg[e]\neg\varphi$, then says that “it is possible to perform observation e in the current state, in such a way that afterwards, φ becomes true”.¹⁵

Using dynamic modalities, our dynamic-epistemic take on perceptual indiscernability can now be reflected in the valid equivalences

$$\langle C \rangle c_i \leftrightarrow \langle See : i \rangle true, \quad C\neg i \leftrightarrow \langle See^- : i \rangle true \leftrightarrow [See : i] false,$$

which show that the static closeness modalities can be given a dynamic definition.

6.2 Principles of observation: Margin of Error reconsidered

Next, we use our logical framework to investigate whether we can make sense of the prima facie appeal of the Margin of Error principle that makes the initial puzzle of the Epistemic Sorites so intriguing. There are several ways of doing this, and we present a few. However, our main intention here is not ‘looking back’. An equally informative way of reading the remainder of this Section is as a further exploration of the logic of observation.

First, we look at some validities in our extended language that combine operators K and C with the observational modalities $[See : y]\varphi$, $[See^- y]\varphi$.

A provable weak version of K-MoE For start, here is a weak version of **K-MoE**, which is a trivial consequence of the assumptions behind this setting:

I know that, if two colors c_n and c_m are perceptually *distinguishable*, then they *cannot* be immediate successors in the sequence; i.e., there exists an intermediary color c_k , with k between n and m , distinct from both of them). Formally:

$$K(c_n \not\approx c_m \Rightarrow n \neq m + 1) \text{ holds for all } n, m.$$

This is just a statement about colors in general, but using the closeness modality C , we can restate it as a statement about the (real color in the) actual world:

$$K(C\neg c_m \Rightarrow \neg c_{m+1}) \text{ holds for all } m. \quad (!\mathbf{K-MoE})$$

This says that I know that, if I could perceptually distinguish (by direct comparison) the color of the actual card from c_m , then that color is not c_{m+1} .

Contrast this with the original principle

$$K(K\neg c_m \Rightarrow \neg c_{m+1}) \text{ holds for all } m. \quad (\mathbf{K-MoE})$$

It is easy to see that our weak version **!K-MoE** is *provable* in the above-mentioned axiomatic system. But, to go from **!K-MoE** to **K-MoE**, we need the additional assumption

$$K\phi \Rightarrow C\phi$$

¹⁵In particular, $\langle e \rangle \top$ means that “observation e can be performed” in the actual state (thereby being equivalent to the sentence $pre(e)$), while $[e]\perp$ means that e cannot be performed.

which in semantical terms corresponds to the condition that

$$w \approx w' \Rightarrow w \sim w' \text{ holds for all } w, w' \in W.$$

We already identified this in Section 2 as an unwarranted assumption underlying the failure of **MoE** and **K-MoE**. Our present logical analysis confirms this diagnosis, while adding the valid principle **!K-MoE** that may provide a partial explanation for the appeal of **K-MoE**.

Precision limits of single observations Here is another legitimate role for **MoE**-type intuitions in our setting. Given our findings so far, their full impact cannot be directly on the truth of static epistemic assertions, but it can well be in the *dynamics*: namely, on the quality of available single observations. We could put this informally as follows:

MoE as minimal threshold. The accuracy of an observation should not be better than the perceptual distance between adjacent points.

This seems quite plausible. In line with our discussion in Section 2, we phrase this as:

If a single (positive) observation results in an *apparent value* y , and if before the observation there was an epistemic possibility x similar to y , then x is still an epistemic possibility after the observation.

This feature of imprecise observations can be formalized using event modalities as

$$\langle K \rangle x \rightarrow [See : y] \langle K \rangle x, \text{ for all } x \approx y.$$

Using the closeness modality C and its possibilistic dual $\langle C \rangle$, this can be rendered as

$$\langle K \rangle (x \wedge \langle C \rangle y) \rightarrow [See : y] \langle K \rangle x,$$

or equivalently

$$\langle See : y \rangle K \neg x \rightarrow K(x \rightarrow C \neg y).$$

This validity can in fact be *proved* syntactically in our dynamic-epistemic axiom system, by making essential use of the recursion axioms for events $See : y$ with their preconditions.

Uniform precision limits on observations Intuitively, the minimal threshold restriction does not apply only to events $See : y$, but to all perceptions of a given agent. This would give us a **MoE**-type restriction on event models **E**, namely that *no single observation in E can be more accurate than an observation of the form See : y*. Formally, this requires that all preconditions $pre(e)$ of events $e \in E$ satisfy: for every $x \in pre(e)$, there exists some $y \in pre(e)$ with $x \in \{z : z \approx y\} \subseteq pre(e)$. This corresponds to the validity of

$$pre(e) \rightarrow \langle C \rangle C pre(e).$$

This statement is *not* valid for all event models, but it can be a reasonable constraint on “perceptually feasible” event models. Unlike **MoE**, this is not an explicitly epistemic

assertion, but it has epistemic consequences. For instance, it follows from it that *absolutely indistinguishable states cannot be distinguished by observations*:

$$x \cong y \rightarrow x \cong_E y.$$

Observational precision as a limit to knowledge However, the above dynamic formulations *only limit the potential increase of knowledge* in the future. None of them binds current knowledge: it is consistent to assume some a priori knowledge of distinctions that go beyond observational indistinguishability. But it would be natural to constrain *perceptual* knowledge within the limits of observability. The relevant condition is:

$$x \cong_E y \rightarrow x \sim y,$$

saying that states that are observationally indistinguishable by *all* observations are also epistemically indistinguishable. If we combine this with the above dynamic variant of **MoE** in terms of the “uniform precision” restriction on **E**, we obtain the implication

$$x \cong y \rightarrow x \sim y,$$

saying that *absolutely indiscernible states are epistemically indistinguishable*. This is very close to **MoE**, resembling the invalid implication $x \approx y \rightarrow x \sim y$.

There are yet further valid ways of analyzing perceptual events in which intuitions of the **MoE** variety play a role, but we defer these to an extended version of this paper.¹⁶

7 Comparison with Other Work

There is a large body of literature dealing with the Margin of Error argument, and an even larger body dealing with related paradoxes involving vagueness ([11] is a good overview). But we discuss here only the work that is closest to our setting.

Williamson [30] was the first to suggest modeling vagueness using a modal operator C for *clarity*: this is the special case of our closeness operator C in the case of metric spaces. Williamson also showed that the modal logic KTB is sound and complete for this semantics. But he studied it in isolation, without pairing it with a separate epistemic operator – while later on, in his definitive presentation of the Margin of Error principle [32], he reuses the semantics of the clarity operator to model perceptual knowledge K .

Egré [14] and Arnesen [1] distinguish between perceptual indiscernibility and epistemic uncertainty, introducing two operators similar to our modalities for closeness C and knowledge K , and discussing variants of **MoE** in this setting. But they do not make a systematic use of this distinction, nor argue for it using observational dynamics.

¹⁶For instance, consider the ‘robustness’ or ‘safety’ of knowledge some authors read in **MoE**. At first sight, this may seem to correspond to the bridge law $K\varphi \rightarrow CK\varphi$: ‘if we know something, we also know it in the perceptual neighborhood of the current world’. But this fails in our models, as it implies the invalid $K\varphi \rightarrow C\varphi$. There is indeed a sense in which knowledge is safe in our setting, but it is a different one.

Bonnay and Egré [9] were the first to use a dynamic-epistemic logic to deal with the Margin of Error argument. Their solution uses an elaboration of the bidimensional ‘centered semantics’ introduced in [8]. This builds Introspection into the model, by keeping track of the actual world, as distinct from the world of evaluation, and interpreting K operators by always looking at the neighbors of the actual world (even with nested K ’s), instead of neighbors of the world of evaluation. This blocks the propagation of error inherent in Williamson’s semantics for nested knowledge, and ensures the models are (positively and negatively) introspective. In contrast, our semantics, though compatible with Introspection, is not committed to it. Moreover, the two approaches are conceptually very different. Centered semantics captures the idea that higher-level knowledge obeys different laws than those governing perceptual knowledge. Bonnay and Egré argue that reflective knowledge is not subject to the **MoE** principle – so that the KK Principle cannot be refuted in this way. But they do seem to concede that **MoE** may be valid for perceptual knowledge. In subsequent work [15], these authors combine epistemic logic with probabilities to analyze **MoE**.

As we stated at the outset, our objection against **MoE** was raised previously in different forms by Mott [21], Halpern [17] and Spector [26], following insights by Bennet [4]. They also analyzed the results of observations as intervals centered at the apparent value (cf. Mott’s ‘point estimate’, Halpern’s ‘subjective estimate’, Spector’s ‘signal’).¹⁷ Their formal solutions are different though. Mott’s analysis is mostly syntactic and semi-formal, without a formal semantics. Halpern, following the ‘runs and systems’ approach from [16], uses a bidimensional semantics, representing a world as a pair (s, I) , where s is the ontic state (comprising the non-epistemic features of the world), and I is an information state (representing the agent’s local state, that comprises all his information).¹⁸ He concludes that perceptual indistinguishability is transitive after all, against common-sense intuitions.¹⁹ Spector’s formalism is somewhat related: starting from Halpern’s bidimensional semantics, he built a one-dimensional semantics (with worlds identified with ontic states, as in Williamson’s model) for notions of ‘necessary knowledge’ (defined by quantifying universally over information states) and ‘possible knowledge’ (which quantifies existentially).

It would be interesting to make a detailed comparison of our framework with those of Halpern, Spector, and Bonnay & Egré. For now, we note that, in contrast to their bidimensional approaches (with ‘thick’ worlds having inner structure), we model knowledge and similarity in a standard Kripke semantics with ‘thin’ worlds identified with ontic states. This is despite the fact that we agree with these authors’ analysis of perceptual knowledge

¹⁷Mott also raises another interesting type of objection to **MoE**, based on the propagation of imprecision to knowledge about others and common knowledge.

¹⁸In our case, we could identify the information state I with the *precondition* of the corresponding observation event e , i.e. the range of ontic states consistent with this observation. This resembles the ‘Subset Space Semantics’ of [20], especially in its specialization to an elegant ‘topo-logic’, where evidence sets in \mathcal{O} form a *topology*. This generalizes our earlier use of intervals and balls.

¹⁹It seems to us that this conclusion is due to the lack of dynamics, or some other context-shifting device, to capture the indexical nature of similarity. In our terms, Halpern’s point is that, if one fixes the context (the current set of observations in our formalism), the observational indistinguishability relation *is* transitive. But, as we saw, the intuitive notion of perceptual indistinguishability \approx is an indexical observational relation against a *variable context* (given by observing either or both of the two states, but no others).

as being not fully determined by the ontic features of the world, but needing an additional parameter: the perceived ‘signal’, or apparent value. But thick worlds, though interesting in themselves, are not necessary for capturing this. Kripke semantics can already do it, via its *relational structure*. Indeed, Williamson’s line of recovering epistemic relations from ontic states is unusual in modal logic. The standard approach takes the relations as primitives that cannot be reduced to or recovered from the states. As such, uncertainty relations can directly store an agent’s information. So it is not surprising that in our dynamics new observations only affect the epistemic relations, but leave unchanged the ontic states.

Finally, to us, the dynamic viewpoint on observation is the crux of the matter. Being able to understand perceptual indiscernibility and epistemic uncertainty in dynamic terms, as instances of observational indistinguishability, throws new light on these concepts. It also allows us to rethink Margin-of-Error intuitions in terms of dynamic-epistemic principles limiting perceptual observations. Unlike the above authors, we supplemented the critique of **MoE** with positive proposals for substitute principles. These reformulations are for us a small step towards a better understanding of the logic of inexact observations.

For this dynamic perspective, the relevant work that should be high-lighted is Cohen’s [10], which proposes, like us, a dynamic-epistemic analysis of Williamson’s Margin of Error argument that puts observational events at center place. Moreover, this analysis carefully distinguishes between closeness and epistemic possibility, just as we have done. However, whereas we have used only standard update mechanisms, this work claims that the semantics of inexact observation requires introducing new forms of dynamic-epistemic update. Moreover, its analysis of tenable forms of **MoE** seems quite different from ours, based on intuitions about ‘boosting’ which lead to new dynamic-epistemic principles of reasoning with interesting epistemological content. Again, a further detailed comparison is needed before we can assess the similarities and differences with our dynamic-epistemic approach more clearly.

8 Conclusions and further directions

We have given some simple, though hopefully illuminating, logical accounts of the functioning of imprecise perception. In doing so, we hope to have dispelled some conceptual confusions, and introduced useful distinctions. But we were not after discussing just one philosophical puzzle, or in dissecting putative philosophical intuitions in a protracted debate. Rather, our goal was understanding how imprecise observation works.

Imprecision is ubiquitous in observation, but even so, it can lead to precision. All of science depends on inexact measurements, but in the long run, repeated observations against different background contexts can lead to greater precision and more knowledge. As we saw, the inherent margin of error of a given perceptual ability or measurement device is not by itself an absolute barrier; the observational context, given by the set of available benchmarks, plays an equally important role, and may allow us to indirectly tell apart perceptually indistinguishable states. Our dynamic epistemic framework throws a new light on the limits of perception, by allowing us to separate various *types* of (in)discernibility.

We end by sketching a few directions for further research based on what we presented.

Counterfactual knowledge: keeping track of the history Link-cutting semantics made observations precise, but it did not keep track of the observational history of a world. Related to that, in the updated models, worlds that are no longer epistemically possible (from the standpoint of the actual world after an observation) could be compared as to perceptual similarity, but they were all declared to be epistemically equivalent. This is because we did not keep track of the alternative observations that might have been made at those words. To address this, we can set up things in a more discerning style, by adopting the so-called *product update* semantics, again a standard tool in Dynamic Epistemic Logic [2, 12]. Given any initial state model (M) and any event model \mathbf{E} , the updated model is now represents as a *product model* \mathbf{MxE} , with pair-worlds (w, e) which track the history: w is the old state of world, and e is the intervening observational event.²⁰ Only pairs (w, e) in which the old state w satisfies the precondition of event e are allowed, since only they represent consistent observational histories. For instance $(c_{20}, \text{See} : 16)$ is a consistent pair, while $(c_1, \text{See} : 16)$ is not a consistent such pair (given the accuracy level $k = 4$). Since we assumed the observing agent knows the result of his own observation, the new epistemic relation is given by:

$$(w, e) \sim (s, f) \quad \text{iff} \quad \text{both } w \sim s \text{ and } e = f.$$

In such a product model, we can reason about what the agent *would know* if the state of the world was different (and different observations were made).²¹

Knowability In the setting of event models \mathbf{E} with many potential events, it is also natural to add an “arbitrary observation modality” quantifying over all possible sequences of observations of the actual state. Roughly put, $\Box\varphi$ is equivalent to $\forall e \in E : [e]\varphi$.²² Its possibilistic dual $\Diamond\varphi := \neg\Box\neg\varphi$, is equivalent to $\exists e \in E : \langle e \rangle\varphi$. The combination $\Diamond K\varphi$ expresses “knowability” by observations: φ may come to be known after more observations. Using this, we can formally check our earlier assertions about identifying the actual world from suitable consecutive observations. For instance, in the full *DEL* model of Example 2, where all colored cards are available as benchmarks (i.e. $(\text{See} : i) \in E$ for all i), we have $c_j \rightarrow \Diamond Kc_j$, expressing that the agent can eventually identify the color of the new card, despite his inability to tell apart adjacent colors by a direct comparison.

Plausibility and belief Observations often come with an indication of the trust we place in them. This trust can have to do with what we take to be the reliability of the source: our eyes, glasses, other devices. In this setting, successive observations construct richer epistemic plausibility models that support not just knowledge, but also belief as truth in all

²⁰In a product model, possible worlds have a richer structure: we can no longer identify a world w with its corresponding ontic features, e.g. the color $c(w)$ of the chosen card. This is similar to bidimensional semantics, except that the second component encodes the intervening events, not the information state.

²¹Note that the answer to this question is not uniquely determined by that initial different state of the world, but it also crucially depends on the intervening event.

²²In the Subset Space Semantics approach mentioned in the previous section, this operator corresponds to the “effort modality” \Box , quantifying over the agent’s potential future information states. It should also be noted that modalities quantifying over events are used essentially in Cohen’s analysis of **MoE** in [10].

most plausible worlds – and various update mechanisms are known, cf. [7, 3]. Still, our earlier distinctions apply. Greater plausibility, too, is an epistemic-doxastic notion that is not at all the same as greater perceptual prominence. Instead of updates, observations now trigger doxastic upgrades that change the plausibility on worlds in the view of the new information. For single observations, there is a natural **MoE**-style choice: make states that are closer to the apparent value resulting from the observation to be more plausible than farther worlds, while excluding as before the ones that are ‘too far’, i.e., perceptually distinguishable from the apparent value. This will come with a natural notion of *belief*, according to which the agent believes (without knowing) that the apparent value is the true object-value. This is the plausibilistic analogue of a probabilistic “normal distribution” assumption.²³ It is not clear how to generalize this to repeated observations: some more structure might be needed, so that one could perform some form of averaging over observational results.

Fuzzy boundaries in perception Events so far had sharp boundaries, observed intervals had precise edges. One might also think, more in line with the literature on vagueness, that an observation itself has fuzzy boundaries. Consider intervals $[x - 1, x + 1]$ of length 2 as before, but now assume that x itself is definitely *in* as an epistemic possibility, everything outside of the interval is definitely *out*, but the boundary points $x - 1$ and $x + 1$ are *indeterminate*. This three-valued setting is attractive as it can refine our view of **MoE**. We might read it as saying that there should never be abrupt transitions in a model from being ‘in’ for one world to being ‘out’ for an immediately adjacent world. This might be more compatible with our preceding analysis than the original formulation of **MoE**. Moreover, technically, it is quite feasible to combine our earlier updates with a three-valued epistemic logic, cf. [18], [27].²⁴

Unknown margin What if the degree of precision of a perception is not known to the agent? This higher-order uncertainty can make scenarios hard to fathom. Say, the agent is unsure whether a measuring device has precision $[i - 1, i + 1]$ as in the above, or $[i - 2, i + 2]$, but it must be one of the two. At first sight, a dynamic-epistemic approach can still model this using uncertainty relations on events, but the details need to be worked out.

Non-factive evidence Finally, what if agents just make successive observations with no level of precision specified at all, indeed possibly containing actual errors (i.e. wrong results, rather than imprecise ones)? Then we may have to shift to the abstract evidence dynamics of [6]. In its simplest version, observations give subsets of the domain of worlds, viewed as pieces of evidence. Using this accumulated evidence, *belief formation* becomes the main epistemic process for perceptual agents, not acquisition of knowledge.²⁵ A study of imprecise perception and measurement in this evidence setting remains to be undertaken.

²³Of course, instead of plausibility, one could also work from the start with a subjective probability distribution over worlds, and use a probabilistic update rule to model new expectations induced by observations. The normal distribution is only one possible choice, but it is not generally warranted.

²⁴However, it is not immediately clear how we should interpret the third indeterminate truth value. If it is epistemic, it seems to involve a potential object-meta level confusion: we are modeling uncertainty with formal means, but now we say that we are uncertain about the precise extension of the uncertainty relation \sim . And if indeterminacy is something else, then what is it?

²⁵In this setting, not all observations need to be reliable, that is, the associated evidence set need not contain the actual world. This is clear for instance, if we allow conflicting disjoint evidence sets.

Dynamic logic of observation and measurement theory A more ambitious next task is linking up between the logics of observation proposed here and Measurement Theory, the best available mathematical account of error and imprecise observation. Such a linkage raises non-trivial questions of how to deal with statistical averaging and other typical quantitative devices in measurement that have no direct qualitative logical counterpart. This junction is indeed our eventual goal, but we leave a serious attempt to future work.

Acknowledgments

We thank Michael Cohen, Paul Egré, Alessandro Giordani, Wes Holliday, and Jan-Willem Romeijn for their insightful critical comments and feedback on a first draft of this paper, as well as for many helpful pointers to a number of relevant papers in the literature.

References

- [1] Arnesen, I. Core Logic: Additional project, Period C, part 2. Manuscript, ILLC, Amsterdam. 2004.
- [2] Baltag, A., L. Moss and S. Solecki. The Logic of Public Announcements, Common Knowledge and Private Suspicions. Proceedings of TARK'98, 43–56. Morgan Kaufman Publishers. 1998.
- [3] Baltag, A., and B. Renne. Dynamic Epistemic Logic. Entry in *Stanford Encyclopedia of Philosophy*. 2016. <https://plato.stanford.edu/>
- [4] Bennet, B. Modal semantics for knowledge bases dealing with vague concepts. In A. G. Cohn, L. Schubert, & S. Shapiro (eds.), *Principles of knowledge representation and reasoning: Proceedings of the 6th international conference (KR-98)*, 234–24). Morgan Kaufman. 1998.
- [5] van Benthem, J. Logical dynamics of information and interaction. Cambridge University Press. 2011.
- [6] van Benthem, J., and E. Pacuit. Dynamic logics of evidence-based beliefs. *Studia Logica*, 99(1), 61–92. 2011.
- [7] van Benthem, J., and S. Smets. Dynamic Logics of Belief Change. In H. van Ditmarsch, J.Y. Halpern, W. van der Hoek and B. Kooi (Eds.), *Handbook of Logics for Knowledge and Belief*, College Publications, 313–393. 2015.
- [8] Bonnay, D., and P. Egré. Inexact Knowledge with Introspection. *Journal of Philosophical Logic*, 38 (2): 179–228. 2009.
- [9] Bonnay, D., and P. Egré. Knowing One's Limits: an Analysis in Centered Dynamic Epistemic Logic, chapter 6 in P. Girard, O. Roy and M. Marion (eds), *Dynamic Formal Epistemology*. Synthese Library 351, Springer. 2011.
- [10] Cohen, M. Inexact Knowledge, Dynamic Introspection and Modesty. Manuscript. Department of Philosophy, Stanford University. 2018.
- [11] van Deemter, K., *Not Exactly: in praise of vagueness*. Oxford University Press. 2010.
- [12] van Ditmarsch, H., W. van der Hoek and B. Kooi. *Dynamic epistemic logic*. Synthese Library 337, Springer. 2007.

- [13] Dutant, J. Inexact knowledge, margin-for-error and positive introspection. In D. Samet (ed.), *Proceedings of the 11th conference on theoretical aspects of rationality and knowledge* (TARK XI), 118–124. Presses Universitaires de Louvain-la-Neuve. 2007.
- [14] Egré, P. Reliability, Margin for Error and Self-Knowledge. In V. Hendricks and D. Pritchard (eds), *New Waves in Epistemology*. 2008.
- [15] Egré, P., and D. Bonnay. Vagueness, uncertainty and degrees of clarity. *Synthese* 174:47–78. 2010.
- [16] Fagin, R., J.Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about knowledge*. MIT Press. 1995.
- [17] Halpern, J.Y., Intransitivity and vagueness, in D. Dubois, Ch. Welty and M.-A. Williams (eds), *Principles of Knowledge Representation and Reasoning: Proceedings of KR2004*. AAAI Press. 2004.
- [18] Jaspars, J. Calculi for constructive communication. A study of the dynamics of partial states. Dissertation, Institute for Logic, Language and Computation, University of Amsterdam. 1994.
- [19] Lehrer, K., *Theory of Knowledge*. Westview Press. 2000.
- [20] Moss, L.S. and R. Parikh. Topological Reasoning and The Logic of Knowledge. In *Proceedings of TARK 95–105*. 1992.
- [21] Mott, P. Margins for error and the sorites paradox. *The Philosophical Quarterly* 48(193): 494–504. 1998.
- [22] Nozick, R., *Philosophical Explanations*. Harvard University Press. 1981.
- [23] Plaza, J. Logics of Public Communication. *Proceedings 4th International Symposium on Methodologies for Intelligent Systems*, 201–216. 1989.
- [24] Rott, H., Stability, strength and sensitivity: Converting belief into knowledge. *Erkenntnis* 61 (2-3): 469–493. 2004.
- [25] Sosa, E., Timothy Williamson’s *Knowledge and its Limits*, in P. Greenough and D. Pritchard (eds), *Williamson on Knowledge*. Oxford University Press. 2009.
- [26] Spector, D. Margin for error semantics and signal perception. *Synthese* 190(15): 3247–3263. 2013.
- [27] Thijsse, E. Partial logic and knowledge representation. Dissertation Tilburg University. Eburon Publishers, Delft. 1992.
- [28] Weatherson, B. Luminous margins. *Australasian Journal of Philosophy*, 83: 373–383. 2004.
- [29] Williamson, T., Inexact Knowledge, *Mind*, 101: 217–42. 1992.
- [30] Williamson, T., *Vagueness*. London/New York: Routledge. 1994.
- [31] Williamson, T. Rational failures of the KK principle, in C. Bicchieri, R. Jeffrey and B. Skyrms (eds), *The Logic of Strategy*. Oxford University Press. 1999.
- [32] Williamson, T., *Knowledge and its Limits*. Oxford University Press. 2000.
- [33] Williamson, T., Margins for error: A Reply. *The Philosophical Quarterly* 50(198): 76 –81. 2003.