ISSN 2965-1557
http://doi.org/10.1590/2965-1557.037.e202532322

Scientific article

Revista de Filosofia Aurora
periodicos.pucpr.br/aurora

PUCPRESS

Improbable knowing and ambiguity: a case against rational epistemic akrasia

Conhecimento improvável e ambiguidade: um caso contra a akrasia epistêmica racional

Veronica de Souza Campos [a]
Belo Horizonte, MG, Brasil
[a] Faculdade Jesuíta de Filosofia e Teologia

Como citar: CAMPOS, Veronica de Souza. Improbable knowing and ambiguity: a case against rational epistemic akrasia. *Revista de Filosofia Aurora*, Curitiba: Editora PUCPRESS, v. 37, e202532322, 2025. DOI: https://doi.org/10.1590/2965-1557.037.e202532322

Abstract

This paper challenges interpretations of Timothy Williamson's (2014) cases of improbable knowing as offering support to the idea of rational epistemic akrasia, i.e., the view that it is sometimes rational for an agent to undertake an inconsistent combination of doxastic attitudes, of the form "p, but I shouldn't believe that p". Williamson's scenarios of improbable knowing have been considered by many to be run-of-the-mill examples of situations whereby epistemic akrasia would be not only possible but also rational. I argue that those cases support the idea of rational epistemic akrasia only if the agent's reasoning, in the scenario, is interpreted in a way that presupposes ambiguity. Once the ambiguity is removed, they can no longer be read as such.

Keywords: *Epistemic akrasia. Rational epistemic akrasia. Improbable knowing. Improbable rationality.*

[a] Doctor of Philosophy from the Federal University of Minas Gerais, e-mail: 182vkai@gmail.com

Resumo

Este artigo desafia interpretações dos casos de conhecimento improvável de Timothy Williamson (2014) como oferecendo suporte à ideia de akrasia epistêmica racional, ou seja, a visão de que às vezes é racional para um agente adotar uma combinação inconsistente de atitudes doxásticas, da forma "p, mas eu não deveria acreditar que p". Os cenários de conhecimento improvável de Williamson vem sendo considerados por muitos como exemplos comuns de situações em que a akrasia epistêmica seria não apenas possível, mas também racional. Eu argumento que esses casos dão suporte à ideia de akrasia epistêmica racional apenas se o raciocínio do agente, no cenário, for interpretado de uma forma que pressupõe ambiguidade. Uma vez que ambiguidade sendo removida, eles não mais podem ser lidos como tal.

Palavras-chave: Akrasia epistêmica. Akrasia epistêmica racional. Conhecimento improvável. Racionalidade improvável.

Preliminary remarks

In two recent papers that came about as late developments of his book *Knowledge and its Limits* (2000), Timothy Williamson has argued for the idea that there are cases in which one knows that **p** in spite of its being overwhelmingly unlikely, on one's evidence at the time, that one knows that **p**1. Those are the so-called cases of "improbable knowing". Williamson relies on formal epistemic models of our imprecise knowledge of values of perceptible and measurable magnitudes to advance his basic argument, which, in short, aims at undermining the so-called "KK Principle" (the principle according to which, if you know p, you know that you know p). If there indeed are cases in which a subject acknowledges its being highly unlikely that he knows that **p** but he knows that **p** nonetheless (cases of improbable knowing), then it follows that KK is not a rational constraint, this being Williamson's argument main consequence.

If Williamson's regard of cases of improbable knowing is correct, it also follows that there seems to be at least some circumstances under which it is rational for one to believe that **p** in spite of its being overwhelmingly unlikely, on one's evidence at the time, that **p** is the case. That's because, according to Williamson, a generalization of the description underlying cases of improbable knowing yields the idea of improbable rationality – scenarios in which it is rational for one to do or to believe something even though it is almost certain on their evidence at the time that it is not rational for them to so do, or to so believe. In other words, it also follows, as a sort of by-product of Williamson's analysis, that a person's having a high degree of confidence in both "**p**" and "[it is very unlikely that I know that **p**, hence] I shouldn't believe that **p**" can be this person's most rational response to the total evidence, at least some of the times.

This – being highly confident in both "**p**" and "I shouldn't believe that **p**" – is what scholars having been terming "epistemic akrasia"2. The idea that the doxastic state whereby one believes both "**p**" and "I shouldn't believe that **p**" can be a person's best response under some circumstances has been the object of a long-standing debate, that pre-dates the publication of Williamson's works: the debate around the notion of rational epistemic akrasia, what some term "level-splitting"3. In this debate, scholars have taken different positions. Some of them purport that said doxastic state can indeed be rational under some circumstances (Williamson 2011, 2014; Horowitz, 2014; Hawthorne, Isaacs & Lasonen-Aarnio, 2015; Christensen 2024), whereas others purport that it cannot (Feldman 2005; Kolodny 2005; Elga 2007, 2013; Christensen 2007, 2010; and Smithies 2012). Since some of the cases that have been fuelling defences of level-splitting are variations of Williamson's scenarios of improbable knowing and improbable rationality, one thing that is important for this debate is getting clear on the extent to which those scenarios actually support the case for rational epistemic akrasia. Do they actually entail that rational epistemic akrasia is the best response an agent can give, when faced with evidence that doesn't support what he thinks that he knows, or thinks that he should believe? This is precisely what I aim to discuss here.

The view I offer in this paper is that Williamson's cases of improbable knowing only offer support to the rationality of epistemic akrasia if they are interpreted in a very specific way. In reality, they only offer

[1] Those two papers are Williamson (2011) e Williamson (2014).

[2] For an interesting analysis of the very concept of epistemic akrasia, see Owens (2002) and Campos (2020).

[3] Epistemic akrasia sometimes goes by the name of "level-splitting" because it implies a mismatch between one's first and higher-order beliefs, that is, between what one believes and what one believes that he should believe.

support to the rationality of epistemic akrasia if they are given an interpretation that involves ambiguity. Once the ambiguity is removed, those cases can no longer be seen as offering a basis for the rationality of epistemic akrasia. This being said, here is the plan for the paper. In next section I'll present and discuss UNMARKED CLOCK, the most famous scenario construed by Williamson to exemplify the idea of improbable knowing. Next, I'll present and discuss a semantic objection to the interpretation of UNMARKED CLOCK as a case of rational epistemic akrasia. I'll conclude by highlighting some of the implications of this objection.

Unmarked clock

One of the cases presented by Williamson to discuss the possibility of improbable knowing is the following[4]:

> Imagine a disc something like an irritatingly austere modernist clock, where (as so often happens) the designer preferred the appearance of functional efficiency to the reality. It consists of a plain unmarked circular dial with a single pointer, like an hour hand, which can point at any one of n equally spaced, unmarked positions on the perimeter of the face. You must judge the position of the hand by the naked eye from some distance away. (…) Imagine that you are looking at the dial from a fixed point of view equidistant from all points on the perimeter. You can make some discriminations between positions, (…) but the difference between neighbouring positions is well below your threshold of discrimination. (Williamson, 2014, p. 979).

Williamson's original discussion of UNMARKED CLOCK relies on formal epistemic models to represent our imprecise knowledge of values of perceptible and measurable magnitudes across different possible worlds. I won't discuss the case using Williamson's terminology since, for our present purposes, the relevant idea can be conveyed in much simpler terms; as well as because most of what has been written about epistemic akrasia in contemporary epistemology has been written using natural language. Basically, suppose a rational epistemic agent attentively looks at the unmarked clock from a fixed point of view equidistant from all points on the perimeter of the dial aiming to know what time it is; and then finds that it is 4 o'clock (Williamson, 2014, p. 980). The agent is aware that he is able to discriminate between the unmarked positions, and that this ability is good, but not perfect. It has a margin of error of, say, plus or minus 5 minutes, so that whenever it appears, on one's visual inspection of the dial, that it is 4 o'clock, it is as likely that it is any time between 3.55 and 4.055.

In Williamson's (2014) presentation of the case (pp. 979-983), the agent knows the margin of error of his own perceptual ability. Because the agent knows such margin of error, he knows that the actual time could be any time between 3.55 and 4.05, and he also knows that his credence (about what time it is) should be equally distributed over the eleven hypothetical positions of the pointer (i.e., between 3.55, 3.56, 3.57…

[4] The so-called "puzzle of the unmarked clock" was first presented by Williamson in Knowledge and its Limits (2000, p. 229), with the initial aim of building a case against the "KK Principle". The case has been re-worked out by Williamson himself (2011, 2014), and it has also been separately adapted by David Christensen (2010) as well as by Adam Elga (2013), with slightly different aims. In all four variations, the case was discussed in a highly abstract style. The version I present here is a slightly less abstract, but still faithful to the original case.

[5] Plus or minus 5 minutes is Williamson's original margin of error. If such a margin for error seems implausible, it can be expanded or shrinked accordingly, and that won't hurt the plausibility of the argument.

4.05). In this way, if asked to guess about the actual position of the pointer, for each of those ten positions, the agent should have approximately .1 credence that the pointer is at that position.

So far, no case for the rationality of inconsistent doxastic attitudes of the form "**p**, but I shouldn't believe that **p**" appears to be supported – the agent in UNMARKED CLOCK believes things that are entirely consistent, such as

> (1) "the actual time is between 3.55 and 4.05"
>
> (2) "my perceptual skills have a margin of error of plus or minus 5 minutes"

And those beliefs, in turn, are entirely consistent with what he believes that he should believe, for instance, that

> (3) "my credence should be equally distributed across the positions comprised within the interval 3.55-4.05, so that, for each position, I should have .1 credence that the pointer is at that position".
>
> (4) "I cannot tell the exact time with any more precision".

For a case such as UNMARKED CLOCK to provide an example of a situation it which it is rational for one to believe that "**p** but I shouldn't believe that **p**" (or, in Williamson's terminology, for this to be a case of improbable knowing – knowing that **p** in spite of its being overwhelmingly unlikely, on one's evidence at the time, that **p** is the case), UNMARKED CLOCK needs to be construed in an alternative way – in a way whereby, all else being equal, the agent knew that it is a certain specific time. For that, he would have to form a belief such as

> (5) "it is 4 o'clock sharp".

A situation whereby the agent in unmarked clock believes (5) whilst retaining the other beliefs stressed above is a situation whereby he finds himself in a doxastic state that is inconsistent, because believing it is a specific time is having credence above a given threshold (that it is that time). If you believed, for instance, that (5) and (4), all at once, you would find yourself in the sort of doxastic state that could be expressed by a moorean sentence, such as "it is 4 o'clock, but I don't know whether it is 4 o'clock" (5 & 4). The content of a moorean sentence that expresses the conjunction (5 & 4) has the form "**p**, but [I cannot tell whether **p**, therefore] I shouldn't believe that **p**". That would only happen if the agent actually comes to form the belief that (5). Could an agent in UNMARKED CLOCK come to form the belief that (5) *rationally*?

To Williamson, the answer to this question is yes. A situation in which the agent rationally believes that (5) and that (4) can come by whenever the pointer appears (to the agent) to be at 4 o'clock, and actually is at 4 o'clock. He says:

> We can model each world as an ordered pair <j, k>, where j is the real position of the hand and k is its apparent position. (...) What one believes in <j, k> is true in just those worlds in which the real position

> is close enough to its apparent position in <j, k> and its apparent position is exactly the same as in <j, k>. (…) In this simplified model, whenever the apparent position matches the real position, belief coincides with knowledge (...). Any failure of match between appearance and reality generates some beliefs that fail to constitute knowledge. (Williamson, 2014, pp. 985-986).

So, according to Williamson, if it appears to me that the pointer is at 4 o'clock and it turns out that the pointer actually is at 4 o'clock, I am justified in believing (or, alternatively, I can rationally believe) that it is 4 o'clock sharp. Being justified in believing something that actually is the case equals having knowledge. So if it appears to me that the pointer is at 4 o'clock and the pointer actually is at 4 o'clock, then I know that it is 4 o'clock sharp, even though I also know that it is very unlikely that I actually know that, since I am aware that my perceptual skills are not perfect, they have a margin of error of plus or minus 5 minutes. There you have it: it is rational for me to know that **p** ("it is 4 o'clock sharp") even though it is very unlikely, on my evidence (my visual inspection of the dial), that I know that **p**. This is, if anything is, a case of rational epistemic akrasia.

Williamson himself reassures the reader that it is, in fact, rational epistemic akrasia that we are talking about, even though he does not use the term akrasia. In his own words:

> We have no general difficulty with the idea that a claim may be true even though it is almost certain on someone's evidence that it is false—for example, when the claim states the actual outcome of a long sequence of unobserved coin tosses. What we find harder to accept is the possibility of the same combination when the claim ascribes knowledge to the very person whose evidence is in question at that very time. Their strong evidence that they do not know p seems incompatible with whatever sort of reliance on p is mandated by their knowing p. But the apparent incompatibility is an illusion, which can seriously distort our assessment of particular ascriptions of knowledge and with it our epistemological theorizing. (Williamson 2014, pp. 990-991)

In his view, thus, the apparent incompatibility (between a person knowing that **p** and it being almost certain on their evidence that **p** is false) is an illusion, meaning, the incompatibility is merely apparent. The two doxastic states are actually compatible.

Now, it must be observed that a distinction between appearance and reality has been introduced. The unmarked clock case being a case of rational epistemic akrasia hinges on this distinction. That's because the only type of situation in which the agent would be justified in believing it is a specific time, as in (5), is when the pointer looks like it is pointing at a specific time and it actually is that time. For the unmarked clock case to yield a case of rational epistemic akrasia, the hand has to appear to me as though it is pointing to 4 o'clock, and this appearence has to correspond to reality, i.e., it has to be 4 o'clock. Situations whereby the hand appears to me as though it is pointing to 4 o'clock but it actually is another time, say, 4:02, are situations in which the proposition specified by (5) ("it is 4 o'clock sharp") is false. And when (5) is false, I cannot know (5). Similarly, situations whereby the hand is at 4 o'clock but it looks like (to me) as though it is at another position, say, 3:59, are situations in which it is irrational for me to believe that (5). Those would be Gettier-type situations: situations whereby I would have a justified true belief that (5), that is not knowledge, insofar as there is no rational support relationship connecting the evidence available to me and the belief I end up forming. Williamson discusses this in §6 of the 2014 paper, "Adding Appearances and Constructing Gettier Cases" (Williamson, 2014, pp. 986-987). For there to be rational improbable knowing –

knowing that "(5), even though it is very unlikely that I know that (5)" – the scenario cannot be gettierized. It has to appear to me as though (5).

The semantic objection

In what follows, I'll present a semantic objection to the interpretation of cases of improbable knowing, such as UNMARKED CLOCK, along the lines developed in the previous section, i.e., as cases of rational epistemic akrasia. The objection tackles the semantic dynamics of the assumptions underpinning the case. It suggests that the way the case is put together hides an ambiguity pertaining exactly to the ideas of appearance and reality, and that this is the only reason why it appears to be a case of rational epistemic akrasia. Once the ambiguity is removed, it no longer appears to be acceptable, for the agent, to endorse an inconsistent combination of doxastic attitudes of the form "**p**, but I shouldn't believe that **p**".

There are three assumptions underpinning Williamson's construction of UNMARKED CLOCK as a case in which it is rational, for the agent, to believe both that (5), "it is 4 o'clock sharp", and that he shouldn't believe that (5). Let's begin with the last notion highlighted in the previous section – for an agent to attain not just improbable knowing, but rational improbable knowing that (5), it has to appear to him as though (5). So, one of the implicit presuppositions of the case is

> (6) Apparent specific time$_{4:00}$
> The pointer looks like (to me) it is at 4 o'clock.

If the pointer looks like (to the agent) it is at 4 o'clock, and if the agent knows about the margin of error of his perceptual skills, then the pointer could be at 4 :00, but it could also be at any of the other ten positions within the verge of position 4:00. This claim – that the pointer could be at any of the other positions within the verge of position 4:00 – can be spelled out as follows:

> **(7) Verge$_{4:00 \pm 5}$**
> The actual position of the pointer is one of the positions on the verge of the position corresponding to 4:00, namely, 3:55, 3:56, 3:57, 3:58, 3:59, 4:01, 4:02, 4:03, 4:04 or 4:05 o'clock.

The proposition expressed by **Verge$_{4:00 \pm 5}$** comprises all of the positions within the interval that goes from 3:55 to 4:05 (which reflects the margin of error of the agent's discriminatory skills, which is of plus or minus 5 minutes) *except* the position corresponding to the one it looks like it is at, namely, the position 4:00. In this situation, the agent should have approximately .1 credence that the pointer is at 4:00, .1 credence that the pointer is at each of the individual positions within the verge of 4:00, and .9 credence that the pointer is in the verge of 4:00 *rather than* at 4:00 (i.e., he should be .9 confident that **Verge$_{4:00 \pm 5}$**).

Now – here is where it starts to get tricky – the agent should be .9 confident that (7), **Verge$_{4:00 \pm 5}$**, only if (6), **Apparent specific time$_{4:00}$**, is true. That's because if the pointer looks like it is at, say, 4:02 instead of at 4:00 (and if the margin of error remains of plus or minus 5 minutes6), there will be a new interval corresponding to the positions where the pointer could be – the interval between 3:57 and 4:07. As Williamson puts it:

---

[6] Cases of variable margins of error are discussed by Williamson (2014) afterwards, in pp. 983-985. Those cases do not concern us here, and they make no difference to the argument being put forward.

> Shifting from one world to another (...) always opens up new epistemic possibilities as well as perhaps closing down old ones. (...) This is a plausible feature of real-life examples of inexact knowledge. As we move through logical space, our epistemic horizon moves with us. New epistemic possibilities appear within our epistemic horizon as others disappear. (Williamson, 2014, p. 982)

In other words, in a possible world whereby the pointer looks like it is at 4:02, rather than at 4:00, the agent should have .1 credence that the pointer is at each position of the new interval (the interval comprised between 3:57 and 4:07), and .9 credence that the actual time is any time within this new interval, which could be expressed by the proposition **Verge**$_{4:02\pm5}$ below:

> **Verge**$_{4:02\,\pm\,5}$
> The actual position of the pointer is one of the positions on the verge of the position corresponding to 4:02, namely, 3:57, 3:58, 3:59, 4:00, 4:01, 4:03, 4:04, 4:05, 4:06 and 4:07 o'clock.

The scopes of **Verge**$_{4:02\,\pm\,5}$ and **Verge**$_{4:00\,\pm\,5}$ intersect with each other to a certain extent, but the former contains elements that the latter doesn't, and vice-versa. So, if the pointer looks like it is at 4:02, then the proposition (7), **Verge**$_{4:00\,\pm\,5}$, will deserve a considerable lesser degree of confidence than it would if the pointer looked like it was at 4:00, which means that that the degree of confidence that the agent should have in **Verge**$_{4:00\,\pm\,5}$ should not be .9, but significant lower than .9. In fact, the further the apparent position of the pointer is from 4:00, the lower the credence in **Verge**$_{4:00\,\pm\,5}$ the agent should have.

In the original unmarked clock case, the pointer looks like it is at 4 o'clock sharp, and it is, in fact, 4 o'clock sharp, so the third presupposition from which the case is construed is

> **(8) Actual specific time**$_{4:00}$
> The pointer is at 4 o'clock.

Also, in the original unmarked clock case, the fact that (6) and (8) are true is deemed sufficient for the agent to attain knowledge that (8). So he knows that (8). But he also knows that it is very unlikely that he has that knowledge, since the pointer is much more likely to not be at 4:00 than it is to be at 4:00. That is, he knows that (7), **Verge**$_{4:00\,\pm\,5}$, is much more likely than (8), **Actual specific time**$_{4:00}$; but he believes that (8) nonetheless (assuming, here, with Williamson, that knowing that **p** entails believing that **p**).

Now, let's pause for a moment and think alongside the following lines. How is having this knowledge preferable to not having it? Suppose it is you (rather than a putative agent) at the UNMARKED CLOCK scenario. Would you not reflect on the situation, knowing about the relations described above? That is, would you not think, at first, that your credence should be distributed between **Verge**$_{4:00\,\pm\,5}$ and **Actual specific time**$_{4:00}$, but afterwards, upon reflection, would you not see that the likelihoods of both **Verge**$_{4:00\,\pm\,5}$ and **Actual specific time**$_{4:00}$ hinge on **Apparent specific time**$_{4:00}$? **Apparent specific time**$_{4:00}$ is certain, but the only credences **Apparent specific time**$_{4:00}$ warrants is .9 credence that **Verge**$_{4:00\,\pm\,5}$ and .1 credence that **Actual specific time**$_{4:00}$. So why risk being wrong, and ending up with a false belief, by placing a different degree of confidence on those propositions (different than the one that is warranted by that about which you are certain, i.e., the appearence)? Why not just focus on **Apparent specific time**$_{4:00}$ and accept to have either the credences it warrants, or simply accept that **Apparent specific time**$_{4:00}$ is all you can know for now, and call it a day? If

you thought along those lines, you would soon realize you don't actually have any reason whatsoever to take a step forward from (6), **Apparent specific time**$_{4:00}$, and form a belief that (8), **Actual specific time**$_{4:00}$. For even if it is true that, in taking that step, you would end up with a true belief (or even with knowledge), the whole set up is a russian roulette: there is no way for you to know that you know the specific time. The only rational beliefs to have is .9 credence that **Verge**$_{4:00 \pm 5}$ and .1 credence that **Actual specific time**$_{4:00}$ (or, alternatively, a full blown belief that the pointer is somewhere in the interval that goes from 3.55 to 4:05).

If that makes sense, then how could Williamson arrive at such a different diagnosis? I believe that has to do with a semantic trick in the interpretation of those propositions. One thing that must be observed that there's a prominent asymmetry between (6), **Apparent Specific time**$_{4:00}$, on the one hand, and (8), **Actual specific time**$_{4:00}$, on the other hand. This asymmetry can be conceptualized in terms of the *skeptical load* of those claims. Skeptical load is a notion coined here to represent the extent to which a proposition is to be accepted and used as a basis for reasoning, in a given context. The skeptical load of a claim is a function of the extent to which it is to be taken seriously, by an agent, that the fact expressed by a the claim is uncertain, or unknown. Thus, a claim is said to carry a heavy skeptical load whenever it includes or entails the understanding that what it expresses is not known to be true by the epistemic agent by whom it is entertained, and a light skeptical load whenever it does not carry that understanding – whenever it can be accepted prima facie7.

For instance, the claim "Vitamin C is the l-enantiomer of ascorbic acid" carries a heavy skeptical load whenever it is entertained by me in pretty much any context, because I am no chemist, and I frankly don't know whether Vitamin C is the l-enantiomer of ascorbic acid, although I deem it likely to be the case, because I, say, remember vaguely from my school years having read in a chemistry textbook that Vitamin C is ascorbic acid. That vague remembrance makes it justifiable for me to accept the claim prima facie in at least some contexts, but it also gives me reason to suspend judgment in relation to this claim, in various other contexts – if, for instance, I'm pressed to offer reasons, or if something highly valuable is at stake and it hinges on Vitamin C actually being not just ascorbic acid, but the l-enantiomer of ascorbic acid, and so forth.

Compare it to the claim "chocolate cakes are made of plain flour". This claim does not carry a heavy skeptical load when it is entertained by me, because I take it to be a well established fact that chocolate cakes are made of plain flour. After all, I've been baking chocolate cakes my entire life and I've always used plain flour as the main ingredient. I may of course be wrong, and I might have done chocolate cakes wrongly my entire life, but the point is: I take it to be beyond reasonable doubt, unlike the case regarding vitamin C, which I take to be below reasonable doubt, not beyond. To borrow Williamson's terminology: of all of the possible worlds in which those claims are relevant for whatever reason and true, the number of possible worlds whereby I would be compelled to suspend judgment regarding the claim "chocolate cakes are made of plain flour" is dramatically smaller than the number of possible worlds whereby I would be compelled to suspend judgment regarding the claim "Vitamin C is the l-enantiomer of ascorbic acid".

This being said, we may see now that claim (8), **Actual specific time**$_{4:00}$, carries a heavy skeptical load when it is entertained by a putative agent with imperfect perceptual skills starring at the dial in

[7] This notion of skepitcal load is not altogether dissimilar to Sophie Horowitz "Reasoning Bridge Principle", according to which "It is rational to reason from any proposition, P, just in case one is rationally highly confident in P" (Horowitz, 2014, p. 734). The difference here is that skeptical load admits of degrees, and allows us to compare different propositions across different contexts (as a way of determining the extent to which it is rational to reason from a proposition) whilst Horowitz's bridge principle is an all-or-nothing affair.

unmarked clock cases, whilst claim (6), **Apparent Specific time**$_{4:00}$, in the same circumstances (i.e., for the same agent, in the same context) carries virtually no skeptical load at all, in that one cannot be wrong about what appears to be, only about what is. It is a corollary of analytic philosophy that one cannot reasonably doubt sense datum. If it looks like (to me) **p**, then I can be confident that it looks like **p** (although I might be delusional, and it might not be the case that **p** even though this is what it looks like). In any case, I won't have motivations to suspend judgment in relation to (6) in most contexts. I do have, however, motivation to suspend judgment in relation to (8) in some contexts, since I cannot discriminate finely enough. It could be at 4 o'clock, but it could also be anywhere else in the relevant margin.

To Williamson, however, the agent *knows* (8), and he knows (8) because (8) is true *as well as* because of (6). The problem with the asymmetry described above is that claim (6), **Apparent Specific time**$_{4:00}$, seems to be taken at face value only when it is contrasted with (8), **Actual specific time**$_{4:00}$. The agent in the original UNMARKED CLOCK contrasts (6) with (8) and forms the belief that (8) because he takes (6) to be beyond reasonable doubt (and that belief ends up amounting to knowledge because it turns out that (8) is true). When, however, he contrasts (6) with (7), **Verge**$_{4:00 \pm 5}$, he does not take (6) to be beyond reasonable doubt, but below – for he knows that (7) entails (6), to the extent that the interval specified in (7) is only the relevant interval if (6). But he is not certain that (7), on the contrary, he believes that (8) is true and (7) is false, since (8) and (7) cannot be both true at the same time. And if (7) is false, (6) is no longer certain. In his reasoning, (6) is certain and not certain.

It is clear, thus, that somebody that reasons along the lines just described (which is the way the agent reasons in Williamson's built up of UNMARKED CLOCK) is ascribing (6) two different meanings, depending on the bearing it has on (7) and on (8). Upon entertaining (7), the agent takes (6) as meaning that the pointer must be within the interval yielded by the position it appears to be at plus its relevant verge, whilst in entertaining (8) the agent takes (6) as meaning something entirely else, namely, that provided that the pointer *merely looks like* it is at a certain position, he ultimately cannot know where the pointer is (and that is why he is uncomfortable with his own position, the position of having improbable knowledge). Though these two meanings might be somewhat related and need not be strictly inconsistent, because both (7) and (8) are conditional on (6), switching between the two meanings in moving from an evaluation of (7) to an evaluation of (8) is a form of conceptual gerrymandering: arbitrarily shrinking and expanding the significance of the fact that the exact position of the pointer is unknown, or shrinking and expanding its skeptical load, to use my terminology.

A rational agent in the unmarked clock scenario must "decide" which meaning he is going to ascribe to (6), the more credulous one, or the more skeptically loaded one ("The pointer looks like it is at 4 o'clock, so there is a chance it might be 4 o'clock" and "The pointer *merely looks like* it is at 4 o'clock, so I don't know what the exact time is", respectively). In other words, he must either take (6) as meaning he should have .1 credence that the pointer is at 4:00 (in which case he must also have .9 credence that the pointer is at the relevant verge, that is, in which case he is entitled to (7), an entitlement that won't be disturbed by the fact that the actual position of the pointer is unknown) *or* he should take (6) as meaning that he actually doesn't really know where the pointer is, in which case he should suspend judgment altogether, instead of forming the belief that (8), even if (8) has always been true. A rational epistemic agent should not ascribe both these meanings to (6) all at once, neither should he switch from one to the other arbitrarily.

Final remarks

When I look at a man coming my way from afar, and I utter a sentence such as "It appears as though it is John", this sentence can mean a few of different things. It can mean "I'm not certain whether it is John, I need to check", or it can mean "It merely appears to be John, but it probably ain't him", or it can yet mean "John is here". In the first interpretation, it is implied that I should suspend judgment and engage in inquiry, an activity that will, hopefully, enable me to attain knowledge. In the second interpretation, it is implied that I should not bother, or that it does not matter. In the third interpretation, it is implied that I take myself to already know.

It is a fact about our natural language that it is susceptible to ambiguities of this type. Most of the times, what will resolve the ambiguity is appeal to contexts, and to what is at stake in case I am wrong. When something valuable is at stake and there will be losses in case I assume that reality is identical with appearances and I turn out to be wrong, the fact that it appears at though something is the case will likely be interpreted as expressing the need to suspend judgment and engage in inquiry. In other contexts, it might be alright to take appearances for granted. What is rational for one to know, and to believe, is at least in part a function of the specific demands of their context.

The tricky thing about UNMARKED CLOCK is that we don't know a lot about the specific demands of the context, and the assumptions underpinning the case are implicit. We don't know whether knowing the exact time is important for the agent, we don't know what the costs of asserting an exact time are in case the agent is wrong, or what the costs are of refraining from asserting an exact time. All that we know is that it appears as though the pointer it at 4:00, and that the margin of error is of plus and minus 5 minutes. We don't know what the precise meaning of "It appears as though the pointer is at 4:00" is, apart from the fact that it entails that the pointer must be within the interval that goes from 3:55 to 4:05. The absence of those intricacies facilitates it going unnoticed that the way the case has been construed models an agent that reasons in a way that interprets the implicit assumptions in a way that is not consistent across the board. It seems that the very agent in UNMARKED CLOCK is not aware of the intricacies either. So much so that, in attaining improbable knowledge that it is 4 o'clock sharp, he ascribes two different meanings to "It appears as though the pointer is at 4:00", as I hope I was able to show.

If we forbid a rational agent from ascribing two different meanings to a proposition that is being used as a basis for reasoning, the result is that situations of improbable knowledge don't arise. Either the agent sticks to a more skeptical attitude towards the meaning of the appearance ("the pointer merely looks like it is at 4:00, but I don't know where it is"), or he sticks to a limited knowledge interpretation ("the pointer is within the relevant interval"). In neither case he makes the move from entertaining what appears to be to actually knowing what is. And if he does not make that move, he doesn't find himself in a situation whereby he knows, but thinks he doesn't know.

Data availability statement

The main focus of this article is contributions of a theoretical or methodological nature, without the use of empirical data sets. Therefore, in accordance with the journal's editorial guidelines, the article is exempt from being deposited in SciELO Data.

References

CAMPOS, V. Disclaiming epistemic akrasia: arguments and commentaries. *Principia*, v. 24, n. 2, p. 333–361, 2020.

CHRISTENSEN, D. Epistemic akrasia: No apology required. *Noûs*, v. 58, n. 1, p. 4–76, 2024.

CHRISTENSEN, D. Epistemology of Disagreement: the Good News. *The Philosophical Review*, v. 116, n. 2, p. 187–217, 2007.

CHRISTENSEN, D. Rational Reflection. *Philosophical Perspectives*, v. 24, p. 121–140, 2010.

ELGA, A. Reflection and Disagreement. *Noûs*, v. 41, n. 3, p. 478–502, 2007.

ELGA, A. The puzzle of the unmarked clock and the new rational reflection principle. *Philosophical Studies*, v. 164, n. 1, p. 127–139, 2013.

HAWTHORNE, J.; ISAACS, Y.; LASONEN-AARNIO, M. The rationality of epistemic akrasia. *Philosophical Perspectives*, v. 35, n. 1, p. 206–208, 2015.

HOROWITZ, S. Epistemic Akrasia. *Noûs*, v. 48, n. 4, p. 718–744, 2014.

KOLODNY, N. Why be rational. *Mind*, v. 114, n. 455, p. 509–563, 2005.

OWENS, D. Epistemic Akrasia. *The Monist*, v. 85, n. 1, p. 381–397, 2002.

SMITHIES, D. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research*, v. 85, n. 2, p. 273–300, 2012.

WILLIAMSON, T. *Knowledge and its Limits*. Oxford: Oxford University Press, 2000.

WILLIAMSON, T. Improbable knowing. *In*: DOUGHERTY, T. (org.). *Evidentialism and its Discontents*. Oxford: Oxford University Press, 2011. p. 147–164.

WILLIAMSON, T. Very Improbable Knowing. *Erkenntnis*, v. 79, n. 1, p. 971–999, 2014.

Editores responsáveis: Léo Peruzzo Júnior e Jelson Oliveira.

RECEBIDO: 13/11/2024
APROVADO: 08/11/2025
PUBLICADO: 19/11/2025

RECEIVED: 11/13/2024
APPROVED: 11/08/2025
PUBLISHED: 11/19/2025