



The omniscient speaker puzzle

Aleksander Domosławski^{1,2} 

Received: 5 October 2022 / Accepted: 6 December 2023
© The Author(s) 2024

Abstract

The epistemicist theory aims to explain ignorance due to vagueness by semantic plasticity: the shiftiness of intensions across close possible worlds resulting from shiftiness in usage. This explanation is challenged by the Omniscient Speaker Puzzle (Sennet in *Philos Stud* 161(2):273–285, 2012). Suppose that an omniscient speaker, Barney, who knows all the facts about usage and how these facts determine the intensions of expressions, cooks up a scheme to stabilise the intension of a normally semantically plastic term like ‘rich’. It seems that ‘rich’ would display all the phenomena associated with vagueness without being semantically plastic, thus making the epistemicist explanation of ignorance due to vagueness insufficient. In this paper, I present a few choice points for epistemicism that arise as a result of the puzzle.

Keywords Vagueness · Epistemicism · Semantic plasticity · Safety · Normality

1 Introduction

The epistemicist theory of vagueness associates the phenomenon of a borderline case with a particular kind of ignorance.¹ Specifically, the epistemicist explains the ignorance due to vagueness by appealing to the semantic plasticity² of vague expressions. A natural interpretation of the epistemicist theory is that semantic plasticity is a necessary condition for vagueness: if ϕ is vague, then ϕ is semantically plastic. However, an interesting challenge may be posed for the epistemicist: there seem to be cases where a term exhibits all the phenomena associated with vagueness and yet is not semantically

¹ For the purposes of this paper, I use the term ‘epistemicism’ to refer to the version of the theory presented by Timothy Williamson (1994).

² See Hawthorne (2006).

✉ Aleksander Domosławski
aleksander.domoslawski@gmail.com; aleksander.domoslawski@amu.edu.pl

¹ The School of Philosophy, Religion and History of Science, University of Leeds, Leeds LS2 9JT, UK

² Department of Philosophy, Adam Mickiewicz University, ul. Szamarzewskiego 89C, 60-568 Poznań, Poland

plastic (Sennet, 2012). This threatens the epistemicist explanation of ignorance due to vagueness: if people whose expressions are not semantically plastic suffer from a relevantly similar kind of ignorance, then it seems that whatever explains their ignorance also explains our ignorance in borderline cases.

In this paper, I will sketch some routes the epistemicist may take to meet such challenges. In Sect. 2, I briefly outline Williamsonian epistemicism. In Sect. 3, I present the challenge to epistemicism coming from the Omniscient Speaker Puzzle (Sennet, 2012). In Sect. 4, I show that in the light of the puzzle, there are several choice points for the epistemicist. Ultimately, the puzzle does not turn out to be a devastating objection to epistemicism, but forces the epistemicist to come up with explanations that push the theory forward.

2 Epistemicism

The prince is clearly rich and the pauper is clearly not, but there are borderline cases where we are not in a position to know whether someone counts as rich or not. Suppose that Dan has \$40,000 and is a borderline case of 'rich'. According to the epistemicist, there is a boundary between the rich and the non-rich: there is a smallest amount of dollars one needs to have to count as rich. The key question for the epistemicist theory is to explain why borderline cases produce ignorance, e.g. why we are not in a position to know whether Dan is rich. The epistemicist answer is Semantic Plasticity:

(Semantic Plasticity) An expression ϕ is semantically plastic if and only if ϕ 's intension (as a result of shifts in usage) could easily be different and we are insensitive to the ways in which the shifts in usage produce the shifts in intension.³

The epistemicist is in a position to explain ignorance in borderline cases using Semantic Plasticity and the safety condition for knowledge (Williamson, 2000, p. 147):

(Safety) If one knows, one could not easily have been wrong in a similar case.

The epistemicist takes safety to be a necessary condition for knowledge.⁴ The combination of Safety and Semantic Plasticity allows the epistemicist to explain ignorance in borderline cases. Suppose that Dan just falls on the side of being rich. The reason why a belief that Dan is rich does not constitute knowledge is that it is not safe. Since 'rich' is semantically plastic, it could easily denote a slightly different property, richness*, which is slightly more demanding so that Dan does not count as being rich*. Therefore, it could easily be the case that we would say something false by expressing

³ This is not the only way to define semantic plasticity. Some authors, e.g. Hawthorne (2006, p. 290) built in the epistemic condition, i.e. the claim that we are insensitive to the ways in which the shifts in usage produce shifts in intension, into the definition of semantic plasticity. Others, e.g. Sennet (2012, p. 276), do not. I opt for the former definition to exclude problems with variations in intension that result from the 'wrong kind' of variation in usage. For instance, suppose that it could easily be the case that the parents of Venus Williams would name her 'Serena'; then the name 'Serena' could easily have a different intension, but it is not the right kind of plasticity that the epistemicist is interested in.

⁴ See Sosa (1999), Pritchard (2009), Williamson (2000, 2009a) for defences of safety-based approaches to knowledge.

our belief that Dan is rich (say, by uttering ‘Dan is rich’). So ignorance due to vagueness is explained by the beliefs in borderline cases failing to be safe due to semantic plasticity.

Thus, there is an interesting interplay between knowledge, Semantic Plasticity, and Safety. A belief being safe is a necessary condition for knowledge. So failure of Safety is sufficient for ignorance. However, it also seems that the epistemicist wants to explain ignorance in borderline cases by semantic plasticity. Ignorance in borderline cases is explained by the failure of safety due to semantic plasticity. A natural interpretation of epistemicism commits the theory to the following Central Thesis:

(Central Thesis) If ϕ is vague, then ϕ is semantically plastic.

In the next section I will present a puzzle due to Adam Sennet (2012) that challenges the Central Thesis.

3 The omniscient speaker puzzle

Consider a linguistic community where the cut-off point for ‘rich’ is \$40,000. Additionally, imagine that one of the speakers in the community, Barney, is an Omniscient Speaker. Barney is able to detect all uses of the term ‘rich’ within the linguistic community. He also knows how usage of ‘rich’ determines its intension for the community. Barney is dedicated to keeping the cut-off point for ‘rich’ fixed at \$40,000. Whenever someone makes an utterance that would shift the intension of ‘rich’, Barney makes sure that a counter-utterance is made to perfectly offset the shift in the intension and keep it fixed at \$40,000. Consider a member of the linguistic community, Jessica, who believes that (#) is true:

(#) The cut-off point for ‘rich’ is \$40,000.

Sennet (2012) argues that there are a few propositions that we can make about the case.

- (1) *‘Rich’ is not semantically plastic in the community.* If Barney is committed to keeping the intension of ‘rich’ fixed at \$40,000, then it does not seem like the intension could easily shift. The intension of ‘rich’ is stable across all the nearby worlds.
- (2) *‘Rich’ is vague in the community.* Sennet (2012, pp. 278–279) argues that ‘rich’ exhibits all the phenomena that vagueness is associated with. Since Jessica and other members of the community (other than Barney) don’t know about Barney’s scheme to keep the intension of ‘rich’ stable, they will treat ‘rich’ as any other vague term, e.g. they will exhibit the same kind of uncertainty in borderline cases for ‘rich’ as in the case of any other vague term etc.
- (3) *Jessica does not know that the cut-off point for ‘rich’ is \$40,000.* Even if Jessica has a true belief about the cut-off point for ‘rich’, she is not in a position to know the location of the cut-off point.

This constitutes a puzzle for the epistemicist. Firstly, (1) and (2) are jointly inconsistent with the Central Thesis: if epistemicism relies on the Central Thesis, this is a

problem for epistemicism. Secondly, Jessica is not in a position to know that (#) is true. However, ‘rich’ in Jessica’s community is not semantically plastic, so it cannot be the semantic plasticity that explains Jessica’s ignorance. Since epistemicism postulates that vague expressions have sharp boundaries, it seems that the burden is on the epistemicist to explain Jessica’s ignorance of the location of sharp cut-off for ‘rich’. The worry of course has a wider scope: if epistemicism cannot explain Jessica’s ignorance by semantic plasticity, then perhaps it cannot explain our (seemingly relevantly similar) ignorance as well.

4 Choice points for epistemicism

Since (1) and (2) are jointly inconsistent with the Central Thesis, there are three strategies that the epistemicist might adopt to try to solve the puzzle: deny (1), deny (2), or deny that epistemicism is committed to the Central Thesis.⁵ In the next subsection, I will consider the epistemicist option of denying (1). In Sects. 4.2 and 4.3, I consider ways for epistemicists to weaken the Central Thesis to avoid the puzzle. Lastly, I outline an epistemicist strategy of denying (2).

4.1 The insisting-on-plasticity route

Let’s consider the question whether ‘rich’ is semantically plastic in Barney’s and Jessica’s community. ‘Rich’ is clearly not semantically plastic for Barney since he is an omniscient speaker: the intension of ‘rich’ could not be easily different unbeknownst to him. The question is whether it is semantically plastic for Jessica and other members of the community. ‘Rich’ is semantically plastic for them only if there is a close possible world where, unbeknownst to them, the intension of ‘rich’ is slightly different. Clearly, there are worlds where the intension of ‘rich’ is slightly different; for instance, Barney could decide to stop his stabilising scheme or decide to keep it fixed at some different value. The question is whether such worlds would count as *relevantly close*.

One can imagine an epistemicist strategy along the following lines. The epistemicist could resist the problem simply by claiming that worlds where Barney decides to fix the reference of ‘rich’ at a different value (or where he abandons his scheme) should count as relevantly close. Such worlds are definitely epistemically possible for Jessica: what she knows does not rule them out as she’s unaware of Barney’s scheme and she is not an omniscient speaker herself. Williamson (2009b, p. 305) writes: ‘on

⁵ One might think that the epistemicist has an easy way out by arguing that they are not strictly speaking committed to the Central Thesis. The epistemicist claims that borderlineness is explained by semantic plasticity, but they are not committed to the more specific thesis that the borderlineness of ϕ is explained by the semantic plasticity of ϕ (and not some other sentence). For instance, consider the sentence “‘Michael is tall’ is true”. Suppose that it is borderline. It is entirely possible that the borderline status of “‘Michael is tall’ is true” is explained by the semantic plasticity of another sentence: ‘Michael is tall’. So the epistemicist is not committed to the Central Thesis in full generality. However, solving the problem is not that simple. Whereas in the case of “‘Michael is tall’ is true” it is not the semantic plasticity of the sentence itself, but rather the semantic plasticity of another sentence, that explains the borderlineness, in case of (#) there does not seem to be any other sentence that could explain the (potential) borderline status of (#). So pursuing this strategy of solving the puzzle does little for the epistemicist.

my view, one knows in a given case only if one avoids error in cases similar to that one, but we cannot specify in a non-circular way how similar a case must be to fall under that condition, or what relative weights should be given to different respects of similarity'. The specification of which worlds count as relevantly close cannot be made independently of our understanding of which worlds are relevant for knowledge. So it's not clear that 'rich' is not semantically plastic after all.

However, the dialectical effectiveness of such a strategy is limited, since we are not able to clearly state all the relevant conditions that would tell us which worlds count as relevantly close. This is a point raised by Sennet (2012, p. 279): circular explanations are not very explanatory. Williamson (2000, 2009b) disagrees: the point of his safety-based explanation is to demonstrate some structural mechanisms at play (e.g. the failure of the KK principle); circular explanations may be informative when it comes to demonstrating these structural features. However, Williamson (2009b, p. 305) agrees that if one were to use the account to decide some particular hard cases, one would likely get things wrong. Nevertheless, if it's not clear which worlds count as close, what the epistemicist achieves at best is a position in which it's not clear whether the epistemicist has a problem. At worst, this insisting that 'rich' is semantically plastic despite the stabilising scheme may seem like grasping at straws. The epistemicist can do better.

4.2 The non-ambitious route

How serious a problem this puzzle poses for epistemicism will depend on the exact version of epistemicism that one endorses. The epistemicist associates ignorance in borderline cases with semantic plasticity; however, it is up for debate how strong this association between semantic plasticity and ignorance is. One natural candidate is the approach that claims that epistemicism understands ignorance in borderline cases as ignorance due to semantic plasticity. This approach is Ambitious: it requires that the epistemicist explains what it means for our ignorance to be due to a particular factor. The Ambitious approach seems to be the standard way of interpreting epistemicism.⁶ However, it is not the only possible way.

The epistemicist may take the Non-Ambitious route. The Non-Ambitious epistemicist may claim that semantic plasticity plays only a part in the explanation of our ignorance in borderline cases. The source of our ignorance is our insensitivity to language: we are insensitive to all the usage facts (say, facts about which utterances members of the linguistic community make or are disposed to make) and we are insensitive to the ways in which these usage facts determine the intensions of expressions in our language. The epistemicist explanation of ignorance in borderline cases arises straightforwardly from the account of inexact knowledge presented by Williamson (1992). To see how semantic plasticity can play only a part in the explanation of our ignorance, let's consider a typical case of Williamson's (1992) inexact knowledge model.

⁶ For instance Caie (2012), Magidor (2018), Litland and Yli-Vakkuri (2016) and Yli-Vakkuri (2016) interpret epistemicism 'ambitiously'.

Suppose that there are 10,000 leaves on a tree that I happen to look at. I know that there are more than 100 and less than 1 million leaves on the tree, but I don't know their exact number. Furthermore, suppose that I luckily estimate and come to believe correctly that there are exactly 10,000 leaves on the tree. On a safety conception of knowledge that the epistemicist embraces one can show that this belief does not constitute knowledge by showing that there is a nearby world where the belief is false. However, there are two standard ways of showing this. The first (call it the Plasticity Approach) is to show that the facts with which the belief is concerned could themselves be different (which would make the belief false). The second (call it the Methods Approach) is to show that the method used to obtain the luckily true belief could easily produce a false one in a nearby world.

In attempting to show that my true belief that there are 10,000 leaves on the tree is unsafe, we could employ the Plasticity Approach. The number of leaves on a tree is modally plastic: it could easily be different. Therefore, we might look at a nearby world where the number of leaves is slightly different, but where I still believe that there are 10,000 leaves on the tree. We can demonstrate that the belief is unsafe by pointing to this world. We could also employ the Methods Approach and show that my belief is unsafe, by looking at nearby worlds where I reach a slightly different conclusion about the number of leaves on the tree on the basis of the same method that I actually use. Since my method of estimating the number of leaves (looking at the tree without any aids) is not very exact, there is a nearby world where I, using the same method, estimate the number of leaves to be something else than it is in that world, e.g. I could easily estimate that the number of leaves is 10,001, when the number of leaves would still be 10,000.⁷

Our knowledge of the meanings of vague expressions in our language is an example of inexact knowledge, e.g. we are in a position to know that 'tall' picks out people who are over 200 cm in height and does not pick out those who are below 160 cm in height, but we are unable to say exactly where the boundary for 'tall' is. Therefore, in Jessica's case, we may attempt to employ the standard ways of explaining ignorance in cases of inexact knowledge. Since the boundary for 'rich' is kept fixed by the Omniscient Speaker in the community, we assume that there are no nearby worlds where the boundary for 'rich' is different than in the actual world. However, for Jessica's belief about the boundary of 'rich' to be safe, she would need to have a reliable method for arriving at a true belief. If she merely correctly guesses where the boundary for 'rich' is, her belief would not be safe as there is a nearby world where she guesses incorrectly.⁸ Therefore, the Non-Ambitious Epistemicist can explain Jessica's ignorance in the puzzling case.

⁷ The first way is used more often to demonstrate the unsafe status of beliefs. However, the second way is also employed, especially to deal with problematic cases like explaining our ignorance of necessary truths. Why cannot I know whether Goldbach's Conjecture is true? Answer: there is no proof of the conjecture and my method of arriving at any belief regarding the conjecture (guessing) could easily lead me to make an error even if I actually guess correctly. See Williamson (1994, p. 210).

⁸ It's unclear whether the Non-Ambitious Route constitutes a departure from the original Williamsonian version of epistemicism. For instance, Williamson (2016b, p. 849) appeals to the Methods Approach when explaining our ignorance of certain borderline sentences. However, he also seems to require that semantic plasticity play a part in the explanation of ignorance in every borderline case (Williamson 2016b, p. 850).

Note that the Non-Ambitious Epistemicist does not appeal to multiple sources of our ignorance to explain ignorance due to vagueness. The source of our ignorance is our insensitivity to language: we cannot discriminate between cases when the intensions of terms are non-identical but similar. Since our powers of discrimination are imperfect, we cannot ‘track’ the relevant truths, e.g. our judgement about where the boundaries for vague terms lie do not perfectly correlate with where they in fact lie. The reason why shiftiness in intensions across nearby worlds produces ignorance is precisely this imperfect discrimination. If Jessica’s judgement is imperfectly correlated with the facts regarding the intension of ‘rich’, then she will fail to update her belief regarding the intension as the facts change (e.g. she will believe that the boundary for ‘rich’ is $\$N$ even in a world where the boundary is slightly different, say, $\$(N+1)$); she will also be prone to form false beliefs: even if she actually truly believes that the cut-off for ‘rich’ is $\$N$ she might have (using the same method) easily formed the false belief that $\$(N+1)$ is the cut-off for ‘rich’. Normally, we can demonstrate our ignorance using the Plasticity and the Methods Approaches; however, in unusual circumstances (such as the setup of the Omniscient Speaker Puzzle), one of the Approaches may be blocked. Nevertheless, the Non-Ambitious Epistemicist has an explanation for Jessica’s ignorance of the truth of (#).⁹ Therefore, the Omniscient Speaker Puzzle is not a problem for the Non-Ambitious Epistemicist.

However, as noted above, much work currently done on epistemicism takes the Ambitious route. One reason for this is that it makes epistemicism more attractive due to the elegance of always being able to appeal to the Plasticity Approach in explaining ignorance due to vagueness.¹⁰ Therefore, we should also consider whether the Ambitious epistemicist can deal with the Omniscient Speaker Puzzle. As I see it, there are two options for the Ambitious epistemicist to try to deal with this problem. First, the epistemicist might try to tweak the account of ignorance due to semantic plasticity by adopting a normality-based explanation of ignorance (instead of the standard safety-based explanation). Second, they could take the error-theoretic route of trying to explain away the intuition that ‘rich’ is vague despite not being semantically plastic within the community in question. The next two subsections are devoted to these two strategies.

⁹ Of course, the epistemicist may also adopt a more confrontational strategy and insist that there is nothing to explain. After all, ignorance, and not knowledge, is our natural state. Unless the critic presents an argument for why we should expect Jessica to possess the relevant knowledge, there is nothing for the epistemicist to explain. Safety is merely a necessary condition for knowledge: even if it turns out that a belief is safe it does not mean it constitutes knowledge. However, this is a less dialectically effective strategy for the epistemicist, because the critics expect epistemicists to explain why the sharp cut-offs for vague expressions postulated by epistemicists are unknowable.

¹⁰ For instance, taking the Ambitious route would make defining the epistemicist definiteness operator more difficult. The recent attempts by Litland and Yli-Vakkuri (2016), Yli-Vakkuri (2016), Magidor (2018) and Domosławski (2023) to define the definiteness operator for epistemicism all take the Ambitious route. Whereas the Ambitious route requires only incorporating the Plasticity Approach into the semantics for ‘definitely’ (so that, in short, ‘Definitely ϕ ’ comes out as false if ϕ expresses a false proposition in a nearby world), taking the Non-Ambitious route would require the additional incorporation of the Methods Approach into the semantics (e.g. when evaluating ‘Definitely ϕ ’, we would look not only at the propositions expressed by ϕ at nearby worlds, but also at propositions that could easily be believed based on the same method). So there are some (non-fatal) costs to pursuing the Non-Ambitious route.

4.3 The normality route

The epistemicist explanation of ignorance is crucially based on the safety condition for knowledge. Roughly, a belief is safe if and only if there is no nearby possible world where the belief is false. The safety condition for knowledge is a well-established principle in epistemology. However, the recent literature¹¹ has seen a shift from safety-based explanations of ignorance to normality-based ones. Whereas safety-based explanations aim to evaluate beliefs at nearby (similar) worlds, e.g. in accounting for Gettier cases, normality-based explanations look at relevantly close normal worlds instead. For instance, the reason why I don't have knowledge that it's 12 o'clock while looking at a broken clock that shows 12 o'clock is that the process of belief formation is not normal (Loets, 2022, pp. 159–160). Normality-based explanations of ignorance could be potentially used to show that even though (#) is safely true (true in all the nearby cases), it is not true in all the relevantly close normal cases.

Suppose that we change slightly the definition of semantic plasticity. Instead of defining a term as semantically plastic if and only if its intension shifts across close nearby worlds (in ways unbeknownst to us), we can define it as a term whose intension is *disposed* to shift: it shifts across the relevantly close normal worlds. Call the property to be disposed to shift in this way 'dispositional semantic plasticity'. The Omniscient Speaker Puzzle does not pose an objection to the claim that 'rich' is dispositionally semantically plastic: I think it's pretty clear that the term 'rich' is dispositionally semantically plastic in Barney's community. Barney's reference-fixing scheme is a fink: even though the intension of 'rich' is disposed to shift, through Barney's strange scheme the intension is kept stable. This case is quite similar to the standard examples of finkish dispositions discussed in the literature.¹² In a standard case, if x has a finkish disposition to ϕ in case ψ , then whereas in normal circumstances x would ϕ if it were the case that ψ , in the strange local circumstances x would not ϕ if it were the case that ψ . For instance, imagine that there is a porcelain vase that is fragile, but that would not break in any nearby counterfactual case because the local deity has decided to protect the vase at all costs. If a rock were thrown at the vase the deity would solidify the air around the vase; if the vase fell, the deity would make the floor soft etc. Similarly, in Barney's case, normally the intension of 'rich' has the disposition to shift with usage; however, in the strange local circumstances of Barney's community, the intension of 'rich' would not easily shift, because of Barney's stabilising scheme. Of course the strange finkish circumstances don't mean that we shouldn't identify fragility with propensity to break or associate ignorance in borderline cases with the effects of dispositional semantic plasticity.

One of the promises of the normality-based approaches to knowledge is their ability to deal with cases like these. Consider Pritchard's (2012) case of Temp, who looks at a broken thermometer, which he believes to be working. Temp forms a true belief about the temperature, because a benevolent demon is controlling the thermostat to match the readings of the broken thermometer. The proponents of the normality approach

¹¹ For instance: Ball (2013), Beddor and Pavese (2020), Dutant (2016), Goodman (2013), Goodman and Salow (2018, 2023), Greco (2014, 2016), Leplin (2007), Littlejohn and Dutant (2020), Peet and Pitcovski (2018), Smith (2010, 2016, 2018a, 2018b), Stalnaker (2005, 2015).

¹² See Lewis (1997).

claim that they can explain Temp's ignorance more easily than the proponents of the safety approach (Beddor & Pavese, 2020, pp. 70–71; Goodman & Salow, 2023, p. 115). Even if Temp's belief is true in every nearby world (thereby making his belief safe), his belief is not true in every relevantly close normal world. After all, there are worlds that are evidentially possible (not excluded by Temp's evidence) that are at least as normal as actuality, in which no one is sneakily controlling the thermostat (Goodman & Salow, 2023, p. 115). The proponent of the normality approach could give an analogous solution to our puzzle case. Jessica doesn't know the boundary of 'rich' at w , because there are worlds that are evidentially possible (for Jessica) and at least as normal as w , in which there is no reference-fixing scheme and where the boundary of 'rich' is different.

Sennet (2012, pp. 280–281) argues that the analysis of Barney's case as an instance of a finkish disposition does not help the epistemicist. According to Sennet, we can adjust the case so that it involves nothing finkish. Suppose that no one is running a scheme to stabilise the intension of 'rich'. However, the intension of 'rich' is stable nevertheless, because whenever someone makes an utterance that would shift the intension, a counteracting utterance is made by another member of the community, thus keeping the intension stable.¹³

The proponent of the normality approach could give a similar response as in the case of Barney's scheme. Even if at w the intension of 'rich' is kept fixed, there are worlds, that are evidentially possible (to Jessica) and that are at least as normal as w , in which the intension of 'rich' is different: either because the community does not have a disposition to keep the boundary of 'rich' stable or because they have a disposition to fix a different boundary of 'rich' (e.g. a world where the community is disposed to fix the boundary of 'rich' at \$40,001 is at least as normal as the world where the community is disposed to fix the boundary of rich at \$40,000).¹⁴ To reraise the problem for the normality approach, one would have to argue that the *only* worlds that are at least as normal as w are the ones where the community keeps the boundary fixed at \$40,000. However, it's difficult to find the motivation for such a view. After all, the tacit assumption in analysing such cases is that the speakers in the community are epistemically limited agents like us and that their normal disposition is to use language sloppily and without being particularly attuned to the slight changes in the communal

¹³ The case would be analogous to a modified case involving Temp: suppose that at w there is no benevolent demon controlling the thermostat; it's just that that the temperature in the house adjusts itself to the readings of the broken thermometer. We can give an analogous reply as in the demon case: the worlds in which the temperature in the house does not adjust itself to the readings of the broken thermometer (or adjusts itself to a different temperature than shown by the thermometer) are at least as normal as w . So Temp's beliefs based on the readings of a broken thermometer are false in some relevantly close worlds that are at least as normal as w .

¹⁴ One could argue that worlds where the community chooses a different intension to keep fixed is not evidentially possible; but that would mean that speakers are able to discriminate between utterances that would shift the reference of 'rich' from the ones that would not. This goes against the epistemicist orthodoxy: Williamson's model relies on the claim that speakers are not able to make these kinds of discriminations (Williamson 1994, pp. 234–237). If speakers are able to make these kinds of discriminations, then it seems that they would be in a position to know the boundary of 'rich'. See Williamson (2016a) for a discussion of the possibility of such bizarre knowledge of mathematical axioms. Williamson (2016b, p. 850) also accepts that there are cases of such bizarre knowledge when it comes to instances of what to us look like borderline cases.

usage. If the members of the community are careful about their own usage in case of some vague term like ‘rich’, then presumably for all they know they use the language sloppily because this is how they usually use vague terms (if they know they are being careful, then the term is not vague for them). So cases where they use ‘rich’ carefully are abnormal; consequently, cases where they use it sloppily are at least as normal.

The overall picture of this Dispositional Epistemicism is this. Vagueness is associated with dispositional semantic plasticity: ϕ is vague if and only if it is dispositionally semantically plastic. If ϕ is dispositionally semantically plastic, then the intension of ϕ will shift across the normal cases. The reason why dispositional semantic plasticity produces ignorance is that we combine it with a normality-based explanation of ignorance (replacing the safety-based one).¹⁵ Roughly, on such a conception, a belief constitutes knowledge only if it is true in all the relevantly close normal cases. Dispositional semantic plasticity will produce errors in normal cases in the same way as semantic plasticity produces errors in nearby cases, which means that if we will be able to explain ignorance using dispositional semantic plasticity in the standard way.

On the standard treatment of definiteness,¹⁶ roughly, ϕ is definitely true if ϕ is true in the actual world relative to interpretations¹⁷ assigned to ϕ across every nearby world.¹⁸ There are several ways to adjust that account to the normality-based explanation of ignorance. I will not be able to map out all the possibilities here, but there are a few constraints that the proponent of the normality version of epistemicism should be mindful of. Firstly, ‘definitely ϕ ’ is factive: if something is definitely true then it’s true. Factivity is easily achieved by the safety-based approaches. On a safety-based approach, a world counts as relevantly close if it’s sufficiently similar to the actual world: no matter how strict we are about similarity, the actual world always counts as sufficiently similar to itself. On the other hand, though factivity is achieved by many normality-based approaches to knowledge,¹⁹ it’s not an automatic feature of the approach as in the case of a safety-based approach. For instance, if we are evaluating X’s belief that p in all the relevantly close normal worlds to determine whether it constitutes knowledge, there is no guarantee that the actual world counts as a normal world (Loets, 2022, p. 173). One way of guaranteeing the factivity of knowledge is to evaluate beliefs at relevantly close worlds that are *at least as normal as actuality* (Goodman & Salow, 2023, p. 93).

Secondly, the friend of the Normality Route needs to be careful as not to undermine one of the main advantages of epistemicism, i.e. the elegance with which it deals with higher-order vagueness. On the standard epistemicist view, the 4-axiom for definiteness ($\Delta\phi \rightarrow \Delta\Delta\phi$) fails and the epistemicist has a good explanation for this: ϕ may be safely true (true in all the nearby worlds) without being safely safely true (true in all the worlds that are nearby to all the nearby worlds). However, some (but not

¹⁵ Taking the Normality Route would clearly indicate a departure from Williamson’s (1994) version of epistemicism.

¹⁶ See Litland and Yli-Vakkuri (2016), Yli-Vakkuri (2016), Williamson (2016b), Domosławski (2023).

¹⁷ Interpretations are assignments of intensions to expressions in the language.

¹⁸ I am simplifying here of course: some important complications appear once we enrich the language to include indexicals like ‘actually’ or metalinguistic vocabulary like ‘true’. These complications can be handled, but they are irrelevant for our purposes here.

¹⁹ See, for instance, Goodman and Salow (2023, p. 92).

all²⁰) normality-based approaches validate the 4-axiom for knowledge: the infamous KK principle (Greco, 2014). Therefore, the epistemicist sympathetic to the normality-based account needs to be careful to avoid a version of the normality view that would undermine one of the main strengths of epistemicism: its treatment of higher-order vagueness.

Lots more work would be required to make the Normality Route a genuine alternative to the standard epistemicist account that could merit and allow for assessment. I am not convinced by the normality account; I prefer the road more travelled by. However, I believe that there are good reasons to think that a normality-based version of epistemicism could handle problems like the Omniscient Speaker Puzzle. Furthermore, I think the above sketch of the normality-based account is a good starting point for future research on non-standard versions of epistemicism.

Can the standard safety-based (Ambitious) version of epistemicism be saved? In the next subsection, I will outline the error-theoretic route to solving the puzzle that gives hope to the traditional epistemicist. My aim will be to deny that ‘rich’, as is used by the Omniscient Speaker’s community, is vague despite appearances. I will also attempt to explain why it nevertheless may seem to be vague without being so.

4.4 The error-theoretic route

Sennet (2012) argues that in the puzzle case, ‘rich’ exhibits all the phenomena associated with vagueness. That is inconsistent with the standard epistemicist picture: semantic plasticity is one of the manifestations of vagueness; if Sennet is correct that ‘rich’ is not semantically plastic in Jessica’s and Barney’s community, then a key component of vagueness is missing by the epistemicist standards. However, we can assume that all members of the community (apart from Barney) would treat anyone with \$40,000 as a borderline case of ‘rich’; since any term that has borderline cases is vague, ‘rich’ would be vague to them. How to explain their treatment of ‘rich’ as vague?

The first thing to note is that the inference from ‘ ϕ is not borderline’ to ‘definitely, ϕ is not borderline’ is not valid. If it were the case, then the members of the community could not be mistaken about there being borderline cases of ‘rich’. However, this inference is not valid. There are general reasons to think that the inference from ‘ ϕ is not borderline’ to ‘definitely, ϕ is not borderline’, i.e. an inference from $(\neg\nabla\phi)$ to $(\Delta\neg\nabla\phi)$, is not valid. For if it were valid, there would be no higher order vagueness.²¹ Similarly, the inference from $\nabla\phi$ to $\Delta\nabla\phi$ is invalid for the same reasons. We could be mistaken about what’s borderline and what’s definitely true.

However, the fact that we could be mistaken about which cases are definite and which are borderline does not yet show that we could be mistaken about a term being vague. I could be wrong about whether someone is a borderline case of ‘tall’ (even if I knew their exact height): I could be wrong where the borderline region for ‘tall’

²⁰ See Loets (2022) for an outline of the different choices regarding the logic of normality that the normality theorist may opt for.

²¹ The statement $(\neg\nabla\phi) \rightarrow (\Delta\neg\nabla\phi)$ entails that there is no higher order vagueness, because it entails the (4) axiom: $\Delta\phi \rightarrow \Delta\Delta\phi$ (if we assume that Δ obeys the standard *K* and *T* axioms for modal logic).

lies. However, it is one thing to be wrong about the location of the borderline region and another to be wrong about its existence. I am fairly certain that there is such a borderline region for ‘tall’ and that ‘tall’ is vague. Could we be wrong about a term being vague?

One reason to think so is Hart’s (1992) argument against the vagueness of ‘heap’. Hart (1992) argues that the term ‘heap of sand’, probably a term most often used in the literature as a paradigm case of vagueness, is not in fact vague (and hence has no borderline cases): the smallest stable heap structure consists of 4 grains of sand (three grains on the bottom and one on the top), so ‘heap’ has a knowable sharp boundary at 4 grains. It is often assumed that terms like ‘bald’ or ‘rich’ don’t pick out any hidden natural boundaries (in Lewisian (1983) terms, all the candidates for reference of such terms are equally ‘eligible’). However, it may be the case that we are not in a position to know whether, for every term we take to be vague, there is no such natural boundary. Suppose that Hart (1992) is right and there is a hidden boundary we pick out when using the term ‘heap of sand’. If Hart had not discovered this fact, we would still treat ‘heap of sand’ as a vague term, because we would not know that such a hidden natural boundary is there to be discovered.²² Of course we don’t have to accept Hart’s arguments regarding ‘heap of sand’. Nevertheless, his position demonstrates that there may be instances where a term is not vague while appearing to be so.

The second thing to note is that people in the Omniscient Speaker’s community (other than the Omniscient Speaker themselves) are unaware of the reference-stabilising scheme. They treat the term ‘rich’ as a vague term. If the scheme were made publicly known, then there would be no puzzle. Suppose that the Omniscient Speaker announces his scheme to stabilise the boundary for ‘rich’ at \$40,000. Then, other members of the community would be in a position to know that (#) is true, because learning what the reference of ‘rich’ is from the Omniscient Speaker who stabilises its meaning is a good (safe) method of obtaining knowledge about its reference. ‘Rich’ would no longer be treated as vague, as the boundary for ‘rich’ would be known. However, suppose that the Omniscient Speaker announces that he is running the stabilising scheme for ‘rich’, but does not specify *which* boundary they decided to keep stable. In such a case, ‘rich’ would presumably be no longer considered to be vague by the members of the community, but would rather be treated analogously to a case where a term picks out a hidden natural boundary (with us being ignorant of which boundary that is). So the epistemicist might claim that the non-omniscient members of the Omniscient Speaker’s community treat the term ‘rich’ as vague because they are unaware of some relevant facts; if they knew about the scheme, they would not treat the term as vague.

²² Another example of unclarity as to whether a term is vague comes from the literature on moral vagueness (vagueness in terms like ‘wrong’ or ‘permissible’). It seems that there are cases where a term like ‘permissible’ is vague, e.g. at which millisecond after conception abortion becomes impermissible. However, there is a discussion about whether moral terms are in fact vague (and whether vagueness can be blamed for our ignorance of the boundaries of moral terms) or whether moral terms are in fact precise and our ignorance is explained by our insufficient knowledge of morality. For example, Miriam Schoenfield (2016, p. 264) claims that on some moral theories (e.g. hedonistic utilitarianism or scalar utilitarianism (Norcross, 2006)), can take the position that there is no moral vagueness. So there are ‘real life’ cases where a term may turn out not to be vague despite appearances.

One might worry that even if the term ‘rich’ is not vague in the Omniscient Speaker’s community, the non-omniscient members in the community still don’t know the boundary of ‘rich’. Their ignorance seems relevantly similar to the ignorance in borderline cases, so the former kind of ignorance should be given the same explanation as the latter; but the former kind of ignorance cannot be explained by semantic plasticity, so the latter cannot either. However, the epistemicist may simply reply that the cases are not relevantly similar, so they don’t require the same explanation. Suppose that I have two red balls and I decide to name the one on the left ‘Phil’. You may not know which ball ‘Phil’ refers to (even if I do), but your ignorance is not explained by vagueness: ‘Phil’ definitely refers to the ball on the left. This case is similar to the Omniscient Speaker case. In the former case, ignorance of the intension of ‘Phil’ is generated by the ignorance of the facts of Phil’s christening; in the later case, the ignorance of the intension of ‘rich’ is generated by our ignorance of which boundary is favoured by the Omniscient Speaker.²³ The epistemicist does not have to treat the cases of ignorance due to vagueness and ignorance in the Omniscient Speaker case as requiring the same explanation.

Nevertheless, it seems that there is still a task for the epistemicist to complete. The epistemicist has to explain why members of the community would think that the term is vague despite the fact that its reference is stable (so should not count as vague by epistemicist standards). I think that despite the fact that we are not perfectly reliable when it comes to assessing whether a term is vague, our intuitions about what’s vague are right most of the time. That is to say the inference from ‘ ϕ seems vague’ to ‘ ϕ is vague’ is a pretty reliable heuristic.²⁴ Even if Hart (1992) is right that ‘heap’ is not in fact vague, which would show that this heuristic is not perfectly reliable, there are not many such examples: we are correct about what’s vague most of the time. Therefore, we should expect the community to treat ‘rich’ as vague, even though (by epistemicist terms) it is not: they are simply relying on a pretty reliable heuristic.²⁵ Of course we could also imagine a case, where Barney becomes more ambitious and stabilises the intensions of all the expressions in the language. By epistemicist terms, there would be no vagueness in such a language. How could then the heuristic be deemed reliable, if there are no terms in the language in that world for which it works? The general answer we can give is that such a case is analogous to other demon-scenarios in epistemology: Barney acts as an ‘evil demon’ subverting the standardly reliable heuristic inference from ‘ ϕ seems vague’ to ‘ ϕ is vague’. If we want to explain why someone forms an erroneous belief in a bad (demon-infested) case, we should look at the way they would operate in a good case that the bad case mimics.²⁶ The reason why inhabitants

²³ For instance, the epistemicist may explain the ignorance in such cases by employing the Methods Approach discussed in Sect. 4.2.

²⁴ See Williamson (2020) on recent philosophical applications of reliable heuristics.

²⁵ One might worry that vagueness would turn out to be an illusory phenomenon on this view: perhaps there are no vague terms and the heuristic systematically misleads us into thinking there are. However, it seems crucial for the development of heuristics or ‘rules of thumb’ that they be very often right: otherwise there would be no reason for us to adopt them. Cf. Williamson’s (2020, pp. 63–67) analysis of tolerance principles that generate the Sorites Paradox.

²⁶ See Williamson (2000, Chapter 8) for the analysis of the distinction between good and bad cases.

of Barney's world (bad case) believe that the terms in their language are vague is that in a good case the inference from ' ϕ seems vague' to ' ϕ is vague' is pretty reliable.

Note that this explanation does not only work for the puzzle at hand, but also in cases where we are mistaken about a term being vague for some other reason. For instance, if Hart's (1992) argument against the vagueness of 'heap of sand' is correct, then until Hart's discovery we mistakenly took 'heap of sand' to be vague. So there may be cases where we are wrong about a term being vague, because there is a previously uncovered joint-carving (natural) referent for some term we thought to be vague. We need the heuristic explanation to account for the appearance of vagueness. Similar mistakes could occur in case when (unbeknownst to the majority of a population) a term is introduced by a precise stipulation. Suppose that w is a world like ours apart from the fact that the term 'bald' was first introduced as a term referring to those with 1000 hairs or less. However, the linguistic community at large is not aware of this initial stipulation: they use 'bald' in the same way as we do. Nevertheless, arguably, 'bald' as used at w has a stable sharp boundary at 1000 hairs, unbeknownst to the members of the community at w .²⁷ We need to explain why to the community at w the term 'bald' would seem vague: the explanation provided above can serve that purpose. Therefore, we need an explanation for why we would be mistaken about a term being vague that is independent of the puzzle.

5 Conclusion

The Omniscient Speaker Puzzle can elicit a variety of responses from the epistemicist. There are a few choice points for the epistemicist: I outlined four possible routes that the epistemicist might take. Firstly, the epistemicist might insist that the term 'rich' is semantically plastic after all (despite the stabilising scheme), because possibilities where the intension of 'rich' is different count as relevantly close; however, this strategy seems to be dialectically weak. Secondly, the epistemicist has to decide whether to take the Ambitious or the Non-Ambitious route. The puzzle does not present a problem for the Non-Ambitious Epistemicist, but their Ambitious cousin does have some explaining to do. Thirdly, the epistemicist has to decide whether to stick to the standard account based on the safety explanation of ignorance or whether to pursue a new path of normality-based explanations. The normality-based approach offers a non-standard path for the epistemicist; however, we don't know where such a path leads yet: more work needs to be done to properly assess such proposals. Lastly, if the Ambitious epistemicist sticks to the safety explanation, they should take the error-theoretic route. They can argue that 'rich' is not vague, because of the reference stabilising scheme, and explain why it seems to be vague for the members of the community. Such an explanation is available: it seems that the community relies on an inference from ' ϕ seems vague' to ' ϕ is vague'; though this inference is not always truth preserving, it is pretty reliable. The explanation of why we can be mistaken about what's vague is

²⁷ Of course semantic drift is possible; after a while the intension of 'bald' may start to vary with use, but plausibly at the initial stage of usage, it will stick to the original stipulation.

needed for independent reasons: it could be that a term is not vague despite appearances because of some unknown reference magnet or uncovered initial stipulation.

One interesting question that remains is: which of the outlined routes is compatible with the original Williamsonian version of epistemicism? The Normality Route is definitely incompatible as it requires abandoning the safety principle championed by Williamson. The Insisting-on-plasticity Route and the Error-theoretic Route don't seem to depart from the original Williamsonian epistemicism. Things are less clear with the Non-Ambitious Approach.

Overall, the Omniscient Speaker Puzzle forces the epistemicist to make some important decisions about the development of the theory and embark on routes that move the theory forward.

Acknowledgements My deepest gratitude goes to Daniel Elstein, Ed Elliott, Michał Maśny, Miriam Bowen, Ofra Magidor, Robbie Williams and Tomasz Zyglewicz for their insightful comments on the paper. Special thanks go also to three anonymous referees at *Synthese*.

Funding This work has been supported financially by the following Grants: Arts & Humanities Research Council through the White Rose College of the Arts & Humanities (Award Number: AH/L503848/1), The Royal Institute of Philosophy (Award: Jacobsen Studentship), School of Philosophy, Religion and History of Science, University of Leeds (Award: Doctoral Scholarship), National Science Centre (Narodowe Centrum Nauki) in Poland (Award Number: 2021/41/B/HS1/00187).

Declarations

Conflict of interest The author has no other relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ball, B. (2013). Knowledge is normal belief. *Analysis*, 73(1), 69–76.
- Beddor, B., & Pavese, C. (2020). Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101(1), 61–79.
- Caie, M. (2012). Vagueness and semantic indiscriminability. *Philosophical Studies*, 160(3), 365–377.
- Domosławski, A. (2023). Defining definiteness. *Ergo*, 10, 20. <https://doi.org/10.3998/ergo.4650>
- Dutant, J. (2016). How to be an infallibilist. *Philosophical Issues*, 26, 148–171.
- Goodman, J. (2013). Inexact knowledge without improbable knowing. *Inquiry*, 56(1), 30–53.
- Goodman, J., & Salow, B. (2018). Taking a chance on KK. *Philosophical Studies*, 175, 183–196.
- Goodman, J., & Salow, B. (2023). Epistemology normalized. *Philosophical Review*, 123(1), 89–145.
- Greco, D. (2014). Could KK be OK? *Journal of Philosophy*, 111(4), 169–197.
- Greco, D. (2016). Safety, explanation, iteration. *Philosophical Issues*, 26, 187–208.
- Hart, W. D. (1992). Hat-tricks and heaps. *Philosophical Studies*, 33, 1–24.
- Hawthorne, J. (2006). Epistemicism and semantic plasticity. *Oxford Studies in Metaphysics*, 2, 289–322.
- Leplin, J. (2007). In defense of reliabilism. *Philosophical Studies*, 134, 31–42.
- Lewis, D. (1983). New work for a theory of universals. *Australasian Journal of Philosophy*, 61(4), 343–377.

- Lewis, D. (1997). Finkish dispositions. *The Philosophical Quarterly*, 47(187), 143–158.
- Litland, J., & Yli-Vakkuri, J. (2016). Vagueness & modality—An ecumenical approach. *Philosophical Perspectives*, 30(1), 229–269.
- Littlejohn, C., & Dutant, J. (2020). Justification, knowledge, and normality. *Philosophical Studies*, 177, 1593–1609.
- Loets, A. (2022). Choice points for a theory of normality. *Mind*, 121(521), 159–191.
- Magidor, O. (2018). Epistemicism, distribution, and the argument from vagueness. *Noûs*, 52(1), 144–170.
- Norcross, A. (2006). The scalar approach to utilitarianism. In H. West (Ed.), *The Blackwell guide to Mill's utilitarianism* (pp. 217–232). Wiley-Blackwell.
- Peet, A., & Pitcovski, E. (2018). Normal knowledge: Toward an explanation-based theory of knowledge. *Journal of Philosophy*, 115(3), 141–157.
- Pritchard, D. (2009). Safety-based epistemology: Whither now? *Journal of Philosophical Research*, 34, 33–45.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3), 247–279.
- Schoenfield, M. (2016). Moral vagueness is ontic vagueness. *Ethics*, 126(2), 257–282.
- Sennet, A. (2012). Semantic plasticity and epistemicism. *Philosophical Studies*, 161(2), 273–285.
- Smith, M. (2010). What else justification could be. *Noûs*, 44(1), 10–31.
- Smith, M. (2016). *Between probability and certainty: What justifies belief*. Oxford University Press.
- Smith, M. (2018a). Coin trials. *Canadian Journal of Philosophy*, 48(5), 726–741.
- Smith, M. (2018b). The logic of epistemic justification. *Synthese*, 195, 3857–3875.
- Sosa, E. (1999). How must knowledge be modally related to what is known? *Philosophical Topics*, 26(1&2), 373–384.
- Stalnaker, R. (2005). On logics of knowledge and belief. *Philosophical Studies*, 128, 169–199.
- Stalnaker, R. (2015). Luminosity and the KK thesis. In S. C. Goldberg (Ed.), *Externalism, self-knowledge, and skepticism: New essays*. Cambridge University Press.
- Williamson, T. (1992). Inexact knowledge. *Mind*, 101, 217–242.
- Williamson, T. (1994). *Vagueness*. Routledge.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford University Press.
- Williamson, T. (2009a). Probability and danger. *The Amherst Lecture in Philosophy*, 4.
- Williamson, T. (2009b). Reply to Goldman. In P. Greenough & D. Pritchard (Eds.), *Williamson on knowledge* (pp. 305–312). Oxford University Press.
- Williamson, T. (2016a). Absolute provability and safe knowledge of axioms. In L. Horsten & P. Welch (Eds.), *Gödel's disjunction* (pp. 243–253). Oxford University Press.
- Williamson, T. (2016b). Reply to Yli-Vakkuri. *Canadian Journal of Philosophy*, 46(4–5), 839–851.
- Williamson, T. (2020). *Suppose and tell: The semantics and heuristics of conditionals*. Oxford University Press.
- Yli-Vakkuri, J. (2016). Epistemicism and modality. *Canadian Journal of Philosophy*, 46(4–5), 803–835.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.