

Does Vagueness Exclude Knowledge?

David Barnett

University of Colorado at Boulder

ABSTRACT

On standard views of vagueness, vagueness as to whether Harry is bald entails that nobody knows whether Harry is bald—either because vagueness is a type of missing fact, and so there is nothing to know, or because vagueness is a type of ignorance, and so even though there is a fact of the matter, nobody can know what that fact is. I argue that vagueness as to whether Harry is bald entails (i) that nobody *clearly* knows that Harry is bald or that Harry is not bald, but not (ii) that nobody *knows* that Harry is bald or that Harry is not bald.

Does Vagueness Exclude Knowledge?

Where exactly does your chin meet your cheek?

It is *vague* where your chin meets your cheek. On standard views of vagueness it is therefore impossible for you to know where your chin meets your cheek, either because vagueness is a type of missing fact—and so there is nothing for you to know—or because vagueness is a type of ignorance—and so even though there is a fact of the matter, you cannot know what it is. On standard views, then, the following principle holds:

Excluded Knowledge Vagueness as to whether p entails that nobody knows whether p .^{1,2}

In this paper, I challenge the standard views by challenging *Excluded Knowledge*.^{3,4} First I argue that *Excluded Knowledge* is not clearly true—at best there is vagueness as to whether it is true. Then I argue that *Excluded Knowledge* is clearly false.

¹ On some versions of vagueness-as-ignorance, a question might be vague for one person but not for another. Proponents of these views would say that vagueness as to whether p entails that nobody “for whom it is vague whether p ” knows whether p .

² *Excluded Knowledge* does not require that vagueness at *any* order with respect to the question of whether p entails that nobody knows whether p . It does not for instance require that vagueness as to *whether it is vague whether p* entails that nobody knows whether p .

³ Cian Dorr (2003) and Brian Weatherson (MS) also argue against *Excluded Knowledge*. Dorr’s argument is based on a defense of bivalence together with intuitive considerations against an epistemic treatment of vagueness. Weatherson’s is based on the view that demonstratives can be substituted salva veritate within the scope of knowledge attributions. In 2001, I challenged *Excluded Knowledge* with an alleged counterexample during a talk at NYU. Wright (2001) refuses to accept *Excluded Knowledge*, but does not positively reject it.

1. Preliminaries

Two preliminary points are in order.

First, throughout the paper I use the expression ‘clearly’ in the following neutral manner, to signify the absence of vagueness:

a is *clearly F* iff (i) *a* is *F* and (ii) it is not vague whether *a* is *F*.

If I say for instance that a certain freckle is clearly part of my cheek, I simply mean (i) that the freckle is part of my cheek and (ii) that there is no vagueness as to whether it is part of my cheek.

The second preliminary point concerns the distinction between *Excluded Knowledge* and a closely related principle:

Excluded Clear Knowledge Vagueness as to whether *p* entails that nobody *clearly* knows *p* and nobody *clearly* knows not-*p*.

If it is vague whether *p*, then *p* is not clearly true, and so a condition that is clearly necessary for knowing *p* fails to clearly obtain. If a condition that is clearly necessary for *x* to be *F* fails to clearly obtain, then *x* cannot be *clearly F*. Hence, *Excluded Clear Knowledge* is clearly true.

Excluded Knowledge might derive some initial appeal from our failure to distinguish it from *Excluded Clear Knowledge*. For this reason, it is important, when evaluating my arguments, to keep in mind that I am arguing against *Excluded Knowledge* and not against *Excluded Clear Knowledge*.

Once *Excluded Clear Knowledge* is carefully distinguished from *Excluded Knowledge*, an initial reason to reject *Excluded Knowledge* emerges. Consider how the two principles bear on the question of what it takes to *clearly* know that Harry is bald. *Excluded Clear Knowledge* merely requires that Harry be *clearly* bald. *Excluded Knowledge*, by contrast, requires that Harry be *clearly clearly* bald. At least initially, this seems wrong. Naturally, to clearly know that Harry is bald, Harry must be clearly bald. To

⁴ I propose an alternative to the standard views in Author #1.

clearly clearly know that Harry is bald, Harry must be clearly clearly bald. And so on. No more does clearly knowing that Harry is bald require Harry to be clearly clearly bald than clearly being a bachelor requires clearly clearly being male. Thus, an initial reason to doubt *Excluded Knowledge* is that it leads to an apparent mismatch in the strength of the requirements on knowing p , clearly knowing p , clearly clearly knowing p , and so on.

2. *Excluded Knowledge* is not clearly true

In this section I attempt to give a *weak* counterexample to *Excluded Knowledge*—one that shows that *Excluded Knowledge* is not clearly true.

To prepare for the example, I begin by saying a word about intuition, or intellectual seeming. Like belief, intuition comes in degrees. My intuition toward the proposition that suffering is intrinsically bad is stronger than my intuition toward the proposition that death is intrinsically bad. This is not to say that I *believe* the former proposition more firmly than I *believe* the latter, or even that I *believe* either proposition. I might for theoretical reasons reject them both. We sometimes have reason to believe that things are not how they seem. We might for instance have reason to believe that we do not fully understand a given predicate. If so, then we have reason to doubt whether things regarding the application conditions of the predicate really *are* as they *seem*. If, on the other hand, we have reason to believe that we fully understand the predicate, then we have reason to take our intuitions about its conditions of application at face value.

For a certain family of predicates, which I will call ‘mundane’, we have very good reason to take our intuitions at face value. Predicates like ‘is bald’ and ‘is rich’ are mundane: their meanings seem shallow, transparent, and without environmental content. There appears to be nothing over and above their meanings than would be reflected by our community-wide pattern of intuitions—elicited in good epistemic conditions, after careful consideration—regarding their conditions of application. This is not because intuition has a mysterious capacity for grasping meanings, but rather because what our community-wide pattern of intuitions would be plays a constitutive role in determining what the

meanings of our mundane predicates are in the first place.⁵ Had our entire community been disposed to have different intuitions about the application of the predicate ‘is bald’, the predicate would have expressed a different concept.

Ideally, this feature of intuition would endow each of us with a disposition to have perfectly accurate intuitions (under good conditions) about the application conditions of our mundane predicates; and, moreover, to *have had* perfectly accurate intuitions, were these predicates to have had slightly different meanings. That is, ideally, it would guarantee not only that our intuitions would be accurate (under good conditions), but that they would be *robustly accurate*: accurate and counterfactually sensitive to minor differences in the meanings of our mundane predicates.

As it happens, there is no guarantee. For even under good conditions there would be slight inconsistencies both in our community-wide pattern of intuitions and in any given member’s pattern. And even if such inconsistencies would not arise at a given time, they would likely arise over any decent stretch of time, for our intuitions are known to be somewhat unstable. Inconsistencies of this sort complicate a theory of meaning by precluding a straightforward story as to how our community-wide disposition toward having intuitions determines the meanings of our mundane predicates. We could not for instance simply read these meanings directly from our pattern of dispositions, with the idea that the correct set of meanings would be whichever one satisfied the pattern, for no set would satisfy the pattern. And so, at least as things stand in our community, the constitutive role of intuition in determining the meanings of mundane predicates is not straightforward enough to endow each of us with perfectly accurate—let alone perfectly and *robustly accurate*—intuitions about the application of our mundane predicates.

But suppose things had been different. Suppose that we had been equipped with better cognitive faculties and a better system of information exchange. Then our tendency toward the preceding sorts of inconsistencies might not have existed. It is easy to imagine such a population, one in all superficial

⁵ A similar, though perhaps more controversial, point can be made about our mundane *concepts*: our collective disposition toward having intuitions about the application conditions of our mundane concepts plays a constitutive role in determining *which* mundane concepts the dispositions are about in the first place.

respects like our own, though unsusceptible to inconsistencies of the preceding sort. Imagine for instance hypothetical population *Z*, whose members are all highly rational speakers of Zenglish, a language that is in most respects just like English and might even have grown out of it. *Z*'s members are not individually or collectively susceptible to inconsistencies of the preceding sort. Their intuitions are consistent, not merely out of luck, but as a result of physiological make-up together with the laws of nature. Indeed, how things would seem to any one of them, upon considering the application conditions of a given mundane predicate, is qualitatively identical to how things would seem to any other. And how things would have seemed to any one of them, if the predicate had had a slightly different meaning, is qualitatively identical to how things would have seemed to any other. As a result, it is not unreasonable to expect that, in community *Z*, the collective disposition toward having intuitions about mundane predicates plays a *straightforward* constitutive role in determining the meanings of those predicates; and that, under good epistemic conditions, the application conditions of *Z*'s mundane predicates are thus guaranteed to *be* just as they *seem*. For instance, upon considering whether hair condition *H* would suffice for falling under the predicate 'is bald*'—a predicate used in almost just the way we use 'is bald'—members of *Z* would under good conditions find it intuitive that *H* would suffice if and only if *H* would suffice.⁶ Moreover, intuitions of this sort would be counterfactually sensitive to differences in meaning; were a mundane predicate of Zenglish to have had different application conditions, every member of *Z* would have had a disposition to have correspondingly different intuitions.

Now, might it be vague what the application conditions of 'is bald*' are? Yes. For consider a state of seeming that you would feel comfortable reporting by saying, "It *sort of* seems to me that *p*," or "It is vague whether it seems to me that *p*." Imagine that this state is qualitatively identical to the state of seeming that all members of *Z* are disposed to go into upon considering whether a given condition, *C*, is sufficient for being bald*. In other words, imagine that it is *vague* whether it would seem to the members

⁶ This is not to suggest that, for *any* predicate of Zenglish, consistency and stability in *Z*'s community-wide pattern of intuitions about the predicate's application conditions would suffice for the truth of these intuitions. There may for instance be a predicate whose associated concept is incoherent, and which thus applies under no condition, even though members of *Z* have not yet noticed this incoherence and find it intuitive that the predicate would apply under various conditions. But 'bald*' and other mundane predicates of Zenglish are not like this.

of Z that C is sufficient for being bald*. Given that the dispositions of the members of Z toward having intuitions about their mundane predicates plays a straightforward constitutive role in determining the meanings of those predicates, we may conclude that it is *vague* whether C is a sufficient condition for the application of the predicate ‘is bald*’. Hence, there could be vagueness in the absence of inconsistency and instability in the linguistic dispositions of a community.

Timothy Williamson disagrees. He says:

A slight shift along one axis of measurement in all our dispositions to use ‘thin’ would slightly shift the meaning and extension of ‘thin’. On the epistemic view, the boundary of ‘thin’ is sharp but unstable. [...] The point is not confined to public language. Even idiolects are vague. You may have no settled disposition to assent to or dissent from ‘TW is thin’. If you were forced to go one way or the other, which way you went would depend on your circumstances and mood. You have no way of making each part of your use perfectly sensitive to the whole, for you have no way of surveying the whole. To imagine away this sprawling quality of your use is to imagine away its vagueness. (1994 pp. 230-32)

Our dispositions to use words *are* unstable. But what is the significance of this instability? Is it the key to understanding vagueness, as Williamson maintains, or does it merely gives rise to more vagueness than would have existed without it?

From a theoretically neutral perspective, it would seem easy enough to imagine vagueness in the absence of such instability. We begin by imagining a population of speakers whose linguistic and conceptual practices are *robustly consistent*—consistent not out of luck, but as a result of physiological make-up together with the laws of nature. Imagining this much is easy. We then pose to ourselves the following question: might it be *vague* how this population uses its expressions? Might it be that its members use a given predicate *consistently* and *robustly*, even though it is vague how they use it? From a theoretically neutral perspective, we found above that the answer is *yes*. We found nothing problematic about the idea of vagueness without inconsistency or instability.

We are now positioned to give a weak counterexample to *Excluded Knowledge*. Let Sophie be an ideally rational member of Z. As she observes the gradual plucking of Harry’s full head of hair, Sophie

continuously and carefully considers whether ‘is bald*’ applies to Harry. She will conclude that ‘is bald*’ applies to Harry just in case it seems to her that ‘is bald*’ applies to Harry. Her perceptual and cognitive conditions are ideal, and she fully understands the predicate ‘is bald*’. Her population is unsusceptible to inconsistencies and instabilities of the preceding sort. Under present conditions, Sophie’s intuitions are perfectly and robustly accurate indications of whether ‘is bald*’ applies to Harry. They are therefore a reliable source of knowledge for Sophie as to whether ‘is bald*’ applies to Harry:

- (1) If it seems to Sophie that ‘is bald*’ applies to Harry, then Sophie knows that ‘is bald*’ applies to Harry.

Now, there is no single plucking prior to which Sophie clearly does not find it intuitive that ‘is bald*’ applies to Harry, and immediately after which Sophie clearly does find it intuitive that ‘is bald*’ applies to Harry. Rather, for a range of pluckings it is vague whether it seems to Sophie that ‘is bald*’ applies to Harry. Suppose that the plucking process is now in this range. If asked at this point whether it seems that ‘is bald*’ applies to Harry, Sophie might naturally say that it *sort of* seems to her that ‘is bald*’ applies to Harry. If asked whether ‘is bald*’ applies to Harry, Sophie might naturally say that ‘is bald*’ *sort of* applies to Harry. And if asked whether Sophie thinks that ‘is bald*’ applies to Harry, Sophie might naturally say that she *sort of* thinks that ‘is bald*’ applies to Harry. In other words, according to Sophie, it is vague whether it seems to Sophie that ‘is bald*’ applies to Harry; it is vague whether ‘is bald*’ applies to Harry; and it is vague whether she believes that ‘is bald*’ applies to Harry.

Does Sophie *know* that ‘is bald*’ applies to Harry? Obviously, Sophie does not *clearly* know that ‘is bald*’ applies to Harry, for (i) ‘is bald*’ does not clearly apply to Harry, and (ii) Sophie does not clearly believe that ‘is bald*’ applies to Harry. Might it nevertheless be *vague* whether Sophie knows that ‘is bald*’ applies to Harry? Given the clear truth of (1), and given that it is vague whether it seems to Sophie that ‘is bald*’ applies to Harry, it is vague, of a condition that it is clearly sufficient for Sophie to know that ‘is bald*’ applies to Harry, whether that condition obtains. Because no other sufficient

condition for Sophie's knowing this proposition clearly obtains, it must be vague whether Sophie knows that 'is bald*' applies to Harry.

This conclusion might at first seem radical. But we have known all along that the notion of *knowing a proposition* is vague: it admits of a thinker and a proposition such that there is vagueness as to whether the thinker knows the proposition. In the preceding scenario, there just happens to be vagueness both as to whether a certain thinker knows p and as to whether p .

If I am right to conclude that it is vague both whether Sophie knows that 'is bald*' applies to Harry and whether 'is bald*' applies to Harry, then *Excluded Knowledge* is not clearly true. For if it were, vagueness as to whether 'is bald*' applies to Harry would *clearly* exclude Sophie's knowing that 'is bald*' applies to Harry. Of course, if *Excluded Knowledge* is not clearly true, then neither is either of the two standard views of vagueness that entail it.

3. Vagueness does not threaten the Law of Excluded Middle

In Section 4 I will argue that *Excluded Knowledge* is clearly false. My argument will depend on the controversial view that vagueness does not threaten the Law of Excluded Middle (LEM). Given LEM, if it is vague whether Harry is bald, then Harry is either bald or not bald, even though it is vague *which*. In this section I will give some initial considerations in favor of LEM, and I will try to explain away some of the initial appeal of rejecting LEM. However, I cannot in this limited space offer a decisive defense of LEM. Suppose, then, that my defense turns out to be unpersuasive. How will that affect the larger aim of the paper, which is to challenge the two standard views of vagueness, vagueness-as-indeterminacy and vagueness-as-ignorance, by challenging *Excluded Knowledge*?

First, it will have absolutely no effect on the argument from the preceding section, which aimed to establish that neither of the two standard views is *clearly* correct. Second, with respect to those versions of the two standard views that are already committed to LEM, it will have no effect on my argument in Section 4 that the two standard views are *clearly false*. Which versions are these? Obviously, all versions of vagueness-as-ignorance are committed to LEM. So, if my forthcoming defense of LEM is

unpersuasive, this will have no effect on my argument that vagueness-as-ignorance is clearly false.

Versions of vagueness-as-indeterminacy that are committed to LEM include what is perhaps the most dominant family of views on vagueness today, namely, those that appeal to the method of supervaluations.⁷ It also includes certain versions of vagueness-as-indeterminacy that do not appeal to this method.⁸ So, if my defense of LEM is unpersuasive, this will have no effect on my argument that these versions of vagueness-as-indeterminacy are clearly false. The only effect it will have is to undermine my argument that versions of vagueness-as-indeterminacy which are not already committed to LEM are clearly false. These include most many-valued treatments,⁹ intuitionist treatments,¹⁰ and a couple of recent psychological treatments.¹¹

Here, then, is my defense of LEM. The defense consists of three initial considerations in favor of LEM, together with an attempt to explain away some of the initial appeal of rejecting LEM.

The first consideration centers on the following hypothetical scenario. As a result of an unfortunate accident, Ning is now in the midst of an operation to reattach one of her legs to her body. At the present stage of the procedure it is vague whether Ning's leg is attached to her body. The surgeon asks an attending student, "What is Ning's current weight?" Well prepared, the student knows that Ning's body weighs precisely 100 pounds without the leg, and that the leg weighs precisely 20 pounds. She offers the intuitive response: "It is vague what Ning weighs. Still, the only live candidates are 100 and 120 pounds. For we can *clearly* rule out all other candidates: clearly, Ning does not weigh 0 pounds, 1 pound, 2 pounds, or any other number of pounds besides 100 and 120. Moreover, although there is no weight such that Ning *clearly* weighs *it*, it is clear that Ning weighs *something*. After all, she is clearly not weightless. Hence, Ning must weigh either 100 or 120 pounds, even though it is vague which." It is vague whether Ning weighs 100 pounds; it is vague whether she weighs 120 pounds; yet, intuitively, it is clearly the case that she weighs one or the other. Intuitively, a disjunction might be clearly true even if none of its

⁷ Mehlberg 1958; Przelecki 1969; Lewis 1970; Fine 1975; Kamp 1975.

⁸ Campbell 1974; Burgess 1990; Horwich 1990, 2005; McGee and McLaughlin 1995; Field 2000.

⁹ Halldén 1949; Körner 1960; Zadeh 1965; Tye 1994.

¹⁰ Putnam 1983.

¹¹ Wright 2001; Field 2003, forthcoming (a).

disjuncts is. More specifically, an instance of LEM might be clearly true even if neither disjunct is: clearly, Ning either does or does not weigh 100 pounds, even though it is vague which.¹²

The claim that the clear truth of a disjunction does not require the clear truth of any of its disjuncts gains further support from the logically equivalent consideration of whether the clear truth of an existential generalization requires the clear truth of any of its substitution instances. Intuitively, it is clearly the case that Ning weighs *something*. Of course, there is no weight such that Ning clearly weighs *it*, for there is vagueness as to what Ning weighs. Hence, intuitively, the clear truth of an existential generalization does not require the clear truth of any of its substitution instances.

Here is the second consideration in favor of LEM. At least prior to theorizing about vagueness, it would seem to be a basic datum that it is vague where the cutoff for being rich is. But use of the expression ‘the cutoff for being rich’ presupposes that there is a cutoff for being rich. And if there is a cutoff for being rich, then, for every amount of worth, either that amount is or is not the cutoff for being rich. Hence, commitment to what would appear to be a basic datum—that it is vague where the cutoff for being rich is—involves a commitment to LEM.

In a similar vein, it would seem to be a basic datum that, if it is vague whether Harry is bald, then it is vague what the answer to the question of whether Harry is bald is. But use of the expression ‘the answer to the question of whether Harry is bald’ presupposes that there is an answer to this question. And if there is an answer to this question, then either Harry is bald or Harry is not bald. Hence, commitment to what appears to be another basic datum involves a commitment to LEM.

I have offered three initial considerations in favor of LEM. Now I want to try to explain away some of the initial appeal of rejecting LEM.

Suppose that Harry is a paradigmatic borderline case of baldness. Upon considering whether Harry is bald, there is an initial temptation to reject both (2) and (3):

(2) that Harry is bald

¹² For other intuitive examples of clearly true disjunctions that lack clearly true disjuncts, see Fine 1975; McGee and McLaughlin 1995; Edgington 1996.

(3) that Harry is not bald.

The rejection of (2) and (3) leads naturally to the rejection of their disjunction, and thus to the rejection of LEM. The rejection should not be interpreted too strongly, as the assertion of the negation of what is rejected, for the negation of (2) contradicts that of (3). Nor should it be interpreted too weakly, as a mere agnosticism, for agnosticism toward both (2) and (3) would not warrant agnosticism toward their disjunction. Just how the rejection should be interpreted is up for debate.

My own suspicion is that our truly initial temptation is to *strongly* reject (2) and (3), by asserting their negations. After all, those who are new to the debate are more than willing to assert, of a borderline bald person, that he is neither bald nor not bald. Sensing contradiction, philosophers resist this temptation and label as ‘weak rejection’ our resultant state of ambivalence toward (2) and (3). If this is right, some of the initial appeal of rejecting LEM might be explained away as follows, in two stages.

First, we locate the source of our temptation to strongly reject (2) and (3). I suspect the temptation arises out of a failure to distinguish *being bald* from *being clearly bald*. On this diagnosis, our temptation to strongly reject (2) and (3) is simply an ill-manifestation of the sound belief that Harry is neither *clearly* bald nor *clearly* not bald. Once the distinction between *being bald* and *being clearly bald* is salient, our urge to strongly reject (2) and (3)—and thus their disjunction—might be satisfied by instead strongly rejecting (4) and (5):

(4) that Harry is clearly bald

(5) that Harry is clearly not bald.

The second stage is to explain away our lingering temptation to weakly reject the disjunction of (2) and (3), by refusing to assert it. I suspect this temptation arises out of a failure to recognize that there is a perfectly adequate explanation of the fact that we should not *assert* (2) or *assert* (3) which is consistent with LEM. Here is the explanation. Consider the following rule:

(M) Assert p only if p .

Ordinarily, we aim to clearly satisfy (M); that is, we prefer a situation in which (M) is clearly satisfied to one in which either (M) is clearly not satisfied or there is vagueness as to whether (M) is satisfied. Now suppose that it is vague whether p . Then one who *clearly* asserts p has not *clearly* satisfied (M) (for, given a clear assertion of p , (M) is clearly satisfied only if p is clearly the case). So one should not clearly assert p . Two options remain: one can clearly refrain from asserting p , or one can try to make it the case that it is vague whether one is asserting p . In any ordinary context, the latter would be inappropriate, for it would require strange behavior that is unlikely to result in successful communication (at least in any community similar to ours). Hence, the former is to be preferred. Of course, clearly refraining from asserting p entails refraining from asserting p . Thus, vagueness as to whether p is—without any further analysis—sufficient to explain why we should not assert p . Applied to our example, vagueness as to whether Harry is bald is sufficient to explain why we should not assert that Harry is bald. Likewise, vagueness as to whether Harry is not bald is sufficient to explain why we should not assert that Harry is not bald. So, without appealing to any particular theory of vagueness, we have an adequate explanation of why we should not assert (2) or assert (3). Clearly, this explanation is consistent with LEM. My suggestion, then, is that our initial failure to recognize the availability of this explanation is responsible for our initial temptation to weakly reject the disjunction of (2) and (3), and, more generally, to weakly reject LEM.

I have attempted to explain away some of the initial appeal of weakly rejecting LEM. Next I want to try to rebut a recent argument for weakly rejecting LEM.

Field (2003 and Forthcoming (b)) argues that once the existence of a cutoff and the relevant instances of LEM are admitted, it is impossible to explain a certain datum: that it seems misguided to *speculate* or to *hope* that the cutoff lies in some exact place.¹³ But perhaps vagueness in the location of the cutoff is—without further analysis, and consistent with LEM—sufficient to explain the datum. To see what I have in mind, consider the following exchange:

¹³ In a similar vein, Smith (2005) suggests that it would be misguided to *guess* where the cutoff lies.

- A: Where is the cutoff for being rich?
- B: It's *vague* where the cutoff is.
- A: You mean there's no cutoff?
- B: No. I mean that it's vague *where* the cutoff is. If there were no cutoff, it wouldn't be vague *where* the cutoff was.
- A: Are you suggesting that there's an amount of money such that being worth that amount fails to suffice for being rich, whereas being worth any greater amount suffices for being rich?
- B: Yes.
- A: That's crazy! You're suggesting that there is a sharp cutoff for being rich!
- B: Nope. That's not what I'm suggesting. The idea of a *sharp* cutoff is crazy. There is a cutoff, but it is not *sharp*: there is a tremendous amount of vagueness as to *where* the cutoff lies. Metaphorically, we say that boundary between what suffices and what fails to suffice for being rich is *blurry*: it has no *clear* location.
- A: You say that there is a cutoff. But wouldn't it be misguided to *speculate* or to *hope* that the cutoff is, say, \$1,045,009.01? And wouldn't it likewise be misguided to *bet* on where the cutoff is?
- B: Typically, it would be misguided. For at best there would be vagueness as to whether such speculation was accurate, whether such a hope was realized, and whether one has won the bet. Typically, when we speculate, hope, and bet, we aim for more than mere vagueness as to whether we succeed: we aim for *clear* success. Of course, if you're in a philosophical mood, and you're willing to settle for vagueness as to whether you've succeeded, then knock yourself out: speculate, hope, and bet away!

I have provided three initial considerations for thinking that vagueness does not threaten LEM, and I have tried to explain away some of the initial appeal of rejecting LEM. Returning to my main line of argument, I now try to give a strong counterexample to *Excluded Knowledge*.

4. *Excluded Knowledge* is clearly false

I now try to develop our example from Section 2 into a *strong* counterexample—one that shows that *Excluded Knowledge* is clearly false. My strategy is to argue that Sophie clearly knows *whether* 'is bald*'

applies to Harry even though it is vague (i) whether ‘is bald*’ applies to Harry; (ii) whether Sophie knows *that* ‘is bald*’ applies to Harry; and (iii) whether Sophie knows *that* ‘is bald*’ does not apply to Harry.

In Section 2 I concluded that it is vague whether Sophie knows that ‘is bald*’ applies to Harry. Let me now argue that it is also vague whether she knows that ‘is bald*’ does not apply to Harry. Recall that Sophie is continuously and carefully considering whether ‘is bald*’ applies to Harry. She knows with confidence that she fully understands the predicate ‘is bald*’, that her epistemic conditions are ideal, and that members of her community are not susceptible to inconsistencies and instabilities of the relevant sort. She thereby knows that, under present conditions, ‘is bald*’ applies to Harry just in case she finds it intuitive that ‘is bald*’ applies to Harry. Thus, if she does not find it intuitive that ‘is bald*’ applies to Harry, she will correctly conclude that ‘is bald*’ does not apply to Harry. Under present conditions, then, Sophie’s *absence* of intuition is a reliable source of knowledge for her that ‘is bald*’ does *not* apply to Harry:

- (6) If it does not seem to Sophie that ‘is bald*’ applies to Harry, then Sophie knows that ‘is bald*’ does not apply to Harry.

Given that it is now vague whether it seems to Sophie that ‘is bald*’ applies to Harry, we may conclude that it is also vague whether it does not seem to Sophie that ‘is bald*’ applies to Harry. So, given the clear truth of (6), it is vague, of a condition that it is clearly sufficient for Sophie to know that ‘is bald*’ does not apply to Harry, whether that condition obtains. Because no other sufficient condition for Sophie’s knowing this proposition clearly obtains, it must be vague whether Sophie knows that ‘is bald*’ does not apply to Harry.

Given LEM, it clearly either does or does not seem to Sophie that ‘is bald*’ applies to Harry, even though it is vague which. If it does seem to her that ‘is bald*’ applies to Harry, then she knows that ‘is bald*’ applies to Harry. If it does not, then she knows that ‘is bald*’ does not apply to Harry. Clearly, then, either Sophie knows that ‘is bald*’ applies to Harry or she knows that ‘is bald*’ does not apply to

Harry, even though it is vague which. Because a thinker knows *whether* p just in case she knows p or she knows $\text{not-}p$, it is clearly the case that Sophie knows *whether* ‘is bald*’ applies to Harry.

To be sure, in a typical context, it would be conversationally inappropriate for Sophie to tell us that she knows whether ‘is bald*’ applies to Harry. For, in a typical context, it is appropriate to tell someone that you know whether p only if you know yourself to be in a position to say whether p . But Sophie knows that she is in no position to say whether ‘is bald*’ applies to Harry. Here is why. Sophie aims to clearly satisfy the following rule of assertion:

(K) Assert p only if you know p .

In other words, Sophie prefers a situation in which (K) is clearly satisfied to one in which either (K) is clearly not satisfied or it is vague whether (K) is satisfied. If Sophie were to *clearly* assert that ‘is bald*’ applies to Harry, then (K) would not be clearly satisfied. For it is vague whether Sophie knows that ‘is bald*’ applies to Harry. Two options remain: either Sophie can clearly refrain from asserting that ‘is bald*’ applies to Harry, or she can attempt to make it the case that it is vague whether ‘is bald*’ applies to Harry. The latter option would require strange behavior that is unlikely to result in successful communication (at least in any community like our own). Thus, as a sincere speaker, Sophie should clearly refrain from asserting that ‘is bald*’ applies to Harry. Similar considerations show that Sophie should clearly refrain from asserting that ‘is bald*’ does not apply to Harry. Sophie, an ideal thinker, knows that she should clearly refrain from asserting these things, and thus that she is in no position to say whether ‘is bald*’ applies to Harry. Thus, in a typical context, it would be conversationally inappropriate for Sophie to tell us that she *knows* whether ‘is bald*’ applies to Harry.

In a similar vein, it typically would be conversationally inappropriate for *anyone* to say, of Sophie, that she knows whether ‘is bald*’ applies to Harry. For typically this would suggest that Sophie is in a position to say whether ‘is bald*’ applies to Harry. Yet she is not. This fact might help to explain our initial discomfort with the idea that Sophie knows whether ‘is bald*’ applies to Harry. Not only does Sophie lack *clear* knowledge of whether ‘is bald*’ applies to Harry, we typically should not even *say* that

Sophie has *knowledge* of whether ‘is bald*’ applies to Harry—for Sophie is in no position to say whether ‘is bald*’ applies to Harry.

I have argued that it is clearly the case that Sophie knows whether ‘is bald*’ applies to Harry despite the fact that it is vague whether he is bald. If my argument is sound, then *Excluded Knowledge* is clearly false. And if *Excluded Knowledge* is clearly false, then so are the two standard views of vagueness that recommend it.

One might worry that I have overstated the significance of my counterexamples to *Excluded Knowledge*. The examples depend crucially on some very special properties of the hypothetical Zenglish-speaking community. These properties are lacking, not only in our actual community, but in any possible community similar to ours. This feature of my example might lead to the following worry: The source of vagueness in the Zenglish-speaking community is very different from the source of vagueness in communities similar to our own. Hence, the example at most establishes that the two standard theories of vagueness cannot accommodate one very special source of vagueness. The example does not establish that the two standard theories cannot accommodate the source of vagueness in communities similar to ours. Hence, the example does not establish that the two theories are incorrect as applied to ordinary cases of vagueness.

But the two standard theories of vagueness are theories of the *nature* of vagueness. On vagueness-as-indeterminacy, what it *is* for it to be *vague* whether *p* is for it to be metaphysically unsettled—for some distinctive reason—whether *p*. On vagueness-as-ignorance, what it *is* for it to be *vague* whether *p* is for it to be epistemically, but not metaphysically, unsettled—for some distinctive reason—whether *p*. On either theory, vagueness *by nature* excludes knowledge. So if there could be an instance of vagueness that does not exclude knowledge, then—no matter how fantastically different from actual the world would need to be to accommodate the possibility—both of the standard theories are false.

To emphasize the point, consider an analogy from ethical theory. Here is a theory of the nature of moral rightness: what it *is* for an agent to do what is *right* is for him to do unto others as he would have them do unto him. Here is a counterexample: imagine that one and only one person wanted, more than

anything in the world, others to torture and mutilate him. Intuitively, it would not be right for that person to torture and mutilate others. Hence, the given theory is false. In light of this counterexample, one might object that the example involves features so foreign to ordinary agents that it only establishes that the given theory does not apply to one very special sort of case, and not that the theory fails to apply to ordinary cases. But this would be a mistake, since the given theory aims to give the *nature* of moral rightness. It does not aim to say what features typical right actions *happen* to have, but rather what features all right actions must have *in virtue* of which they are right.

To be sure, a proponent of one of the two standard theories of vagueness might retreat from his original position by moving to one of the two corresponding weaker claims, that ordinary cases of vagueness happen also to be cases of ignorance, or that ordinary cases of vagueness happen also to be cases of indeterminacy. But he cannot maintain his original theory that ordinary cases of vagueness are cases of vagueness *in virtue* of being cases of certain sorts of ignorance or indeterminacy. What is needed, if my counterexample is sound, is an alternative to the two standard theories of the nature of vagueness—a theory that tells us what vagueness *is*. Such a theory should not entail that vagueness by nature excludes knowledge.

On my own view, vagueness is neither a type of indeterminacy nor an epistemic phenomenon, but is rather *sui generis*. On this view, there is a close-knit family of concepts, none of which admits of analysis in terms of concepts outside the family, and at least some of which are essential to an adequate treatment of vagueness-related phenomena. One might hold a similar non-reductionism about intentional, normative, or modal phenomena. For instance, one might hold that the notions of physical necessity, metaphysical necessity, and causation, together with that of a law of nature, admit of analysis only in terms of one another, if at all. To get a feel for the family of vagueness-related notions, suppose again that it is vague whether Harry is bald. Depending on the source of this vagueness, we might characterize Harry as a *borderline case* of baldness. And we might say that Harry is neither *clearly* bald nor *clearly* not bald. We might say that he is *sort of* bald and *sort of* not bald; and that he both *roughly* qualifies and *roughly* fails to qualify as bald. This, we might say, is due in part to the fact that baldness is a *rough*

concept; it bears only a *rough* relation to more basic concepts; and so there is *vagueness* as to what its basic application-conditions are. According to *vagueness-as-sui-generis*, it is impossible to break out of this family of concepts by way of reduction. This view is consistent with my argument against *Excluded Knowledge*, for nothing in the idea that vagueness is *sui generis* suggests that vagueness as to whether *p* excludes knowing whether *p*.^{14,15}

REFERENCES

Author #1

Author #2

Bealer, G. 1999. A Theory of the A Priori. *Philosophical Perspectives 13: Epistemology*: p. 31.

Dorr, C. 2003. Vagueness without Ignorance. *Philosophical Perspectives 17: Language and Philosophical Linguistics*: 83-113.

Edgington, D. 1996. Vagueness by Degrees. In Keefe and Smith 1996.

Field, H. 2000. Indeterminacy, Degree of belief, and Excluded Middle. *Nous*:1-30.

Fine, K. 1975. Vagueness, Truth and Logic. *Synthese* 30:265-300.

Keefe, R., and P. Smith, eds. 1996. *Vagueness: A Reader*. Cambridge, MA: MIT Press.

McGee, V., and B. McLaughlin. 1995. Distinctions without a Difference. *Southern Journal of Philosophy* 33 (Supplement):203-51.

¹⁴ I develop and motivate this view in Author #1.

¹⁵ For helpful comments and discussion, I am grateful to...

Sainsbury, R.M. 1986. Degrees of Belief and Degrees of Assent. *Philosophical Papers* 15:97-106.

Schiffer, S. 2000. Vagueness and Partial Belief. In *Philosophical Issues 10: Skepticism*, edited by E.

Sosa and E. Villanueva. Boston: Blackwell.

Sorensen, R.A. 1988. *Blindspots*. Oxford: Clarendon Press.

———. 2001. *Vagueness and Contradiction*. Oxford: Oxford University Press.

Williamson, T. 1994. *Vagueness*. New York: Routledge.

Wright, C. 2001. On Being in a Quandry. *Mind* 110: 45-98.