

Does Vagueness Exclude Knowledge?

David Barnett

University of Colorado at Boulder

ABSTRACT

On two 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 truth, and so there is nothing to know, or because vagueness is a type of ignorance, and so even though there is a truth of the matter, nobody can know what that truth is. Vagueness as to whether Harry is bald does entail that nobody *clearly* knows that Harry is bald and that nobody *clearly* knows that Harry is not bald. But it does not entail that nobody *knows* that Harry is bald or that nobody *knows* that Harry is not bald. Hence, the two standard views of vagueness are mistaken.

Where exactly does your chin meet your cheek?

It is *vague* where your chin meets your cheek. On two 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 truth—and so there is nothing for you to know—or because vagueness is a type of ignorance—and so even though there is a truth of the matter, you cannot know what it is. More generally, proponents of vagueness-as-absence-of-truth and vagueness-as-ignorance endorse the following principle:

Excluded Knowledge Vagueness as to whether p entails that nobody knows whether p .¹

In this paper, I challenge vagueness-as-absence-of-truth and vagueness-as-ignorance by challenging *Excluded Knowledge*.^{2,3} First I make two preliminary points (§1). Then I argue that *Excluded*

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 Cian Dorr (2003) and Brian Weatherson (MS) also argue against *Excluded Knowledge*. Dorr’s argument is based on a defense of the Law of Excluded Middle, 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. Crispin Wright (2001) refuses to accept *Excluded Knowledge*, but does not positively reject it. In 2001, I challenged *Excluded Knowledge* with an alleged counterexample during a talk at NYU.

Knowledge is not clearly true—at best there is vagueness as to whether it is true (§2). To prepare for my argument that *Excluded Knowledge* is clearly false, I argue that vagueness does not threaten the Law of Excluded Middle (§3). Finally, I argue that *Excluded Knowledge* is clearly false (§4).

1. Preliminaries

Two preliminary points are in order.

First, I will 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 knowing *p* fails to clearly obtain, then there can be no clear instance of knowing *p*; thus, nobody clearly knows *p*. By the same reasoning, if it is vague whether *p*, then nobody clearly knows not-*p*. Given that this reasoning is clearly sound, *Excluded Clear Knowledge* is clearly true.

³ My argument is consistent with theories on which (i) there is a truth of every vague matter, but (ii) it is unsettled, not merely epistemically, what that truth is. For example, Cargile 1969, Campbell 1974, Burgess 1990, Horwich 1990, McGee and McLaughlin 1995, and Field 2000. In a case of vagueness as to whether *p*, proponents of these theories are free to accept that a thinker could either know *p* or know not-*p* (and thus know whether *p*) even though it is unsettled, not merely epistemically, which proposition the thinker knows, *p* or not-*p*.

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 the clear truth of *Excluded Knowledge* emerges. Consider how the two principles bear on the question of what it takes to *clearly* know that Harry is bald. The clear truth of *Excluded Clear Knowledge* merely requires that Harry be *clearly* bald. The clear truth of *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 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 being clearly clearly male. Thus, an initial reason to doubt the clear truth of *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 more to 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. 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 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 from the armchair 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. 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.

Note a couple of things that I am not claiming. First, I am not claiming that members of Z are infallible when it comes to determining whether a given person is bald*. For all I have said, members of Z have terribly unreliable perceptual faculties and, as a result, are systematically deceived, in the first instance about which basic hair condition members of their community are in and, in the second instance, about whether members of their community are bald*. I have only claimed that, under good conditions, upon considering *from the armchair* whether a given hair condition *would* suffice for being bald*, members of Z would find it intuitive that the condition would suffice, if and only if, it would suffice. And so I am not committing to the strong claim that members of Z can never be wrong about the actual pattern of application of ‘is bald*’. Second, I am not claiming 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. I am only claiming that ‘is bald*’ and other mundane predicates of Zenglish are not like this.

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 hair condition, H, is sufficient for being bald*. In other words, imagine that it is *vague* whether it would seem to the members of Z that H 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 H 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:

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: 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 give 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*. Imagine that our subject, Harry, starts off with a full head of hair and that, one by one, we gradually pluck his hairs, until none is left. Let Sophie be an ideally rational member of Z. For each of Harry's consecutive hair conditions, Sophie carefully considers whether it suffices for being bald*. For simplicity, we shall say that, at each stage of the plucking, Sophie considers whether *Harry* is bald*. But in fact what she considers is whether the corresponding *hair condition* suffices for being bald*. She makes this consideration from the armchair. So the accuracy of her judgments does not in any way depend on her ability to determine what hair condition Harry is in.

Sophie's 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 Harry is bald*.

Moreover, Sophie has no reason, evidence, or justification for thinking that her intuitions are misleading. If it seems to her that Harry is bald*, then she justifiably believes that Harry is bald*. This is not to say that, if it seems to her that Harry is bald*, then she justifiably *infers* that Harry is bald* from her beliefs about how things seem to her. Rather, if it seems to her that Harry is bald*, then this intuition *itself* justifies her in believing that Harry is bald*. That is to say, the intuition provides *non-inferential* justification for her belief that Harry is bald*.

That seemings—both intellectual and perceptual—may provide non-inferential justification is widely accepted. In the case of perceptual seemings, Jim Pryor (2000) argues that one has an immediate justification for believing *p* if one has a perceptual experience as of *p* and one has no reason, evidence, or justification for thinking that one's perceptual experiences are misleading. Michael Huemer (2001) defends the same idea for seemings in general—both intellectual and perceptual. In our example, Sophie has no reason, evidence, or justification for thinking that her intuitions are misleading. So, by Huemer's account, if it seems to Sophie that Harry is bald*, then Sophie is immediately justified in believing that Harry is bald*.

Tyler Burge (2003) provides an *externalist* defense of the idea that our perceptual experiences may provide non-inferential justification. On this account, a perceptual experience immediately justifies a corresponding belief only if the experience is produced by a reliable mechanism. Of course, a reliabilist condition on non-inferential justification is no obstacle to the idea that Sophie's intuitions provide non-inferential justification. For, given the physiological facts about members of her community, together with the other stipulations of the scenario, it is clear that her intuitions are produced by a reliable mechanism.

Those who have explicitly defended the idea that intuitions provide non-inferential justification include George Bealer (1999), Lawrence BonJour (1998), Michael Huemer (2001, 2007), Gideon Rosen

(2001), and Ernest Sosa (1991). For present purposes, we need not distinguish between internalist and externalist versions of the idea, for Sophie's intuitions plausibly satisfy whatever internal or external conditions there might be on non-inferential justification. Indeed, given that her intuitions are *perfectly* and *robustly* accurate indications of whether Harry is bald*, they plausibly provide not only for non-inferential *justification*, but for non-inferential *knowledge*:

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

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

Does Sophie *know* that Harry is bald*? Obviously, Sophie does not *clearly* know that Harry is bald*, for (i) Harry is not clearly bald*, and (ii) Sophie does not clearly believe that Harry is bald*. Might it nevertheless be *vague* whether Sophie knows that Harry is bald*? Given the clear truth of (1), and given that it is vague whether it seems to Sophie that Harry is bald*, it is vague, of a condition that is clearly sufficient for Sophie to know that Harry is bald*, 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 Harry is bald*.

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 Harry is bald* and whether Harry is bald*, then *Excluded Knowledge* is not clearly true. For if it were, vagueness as to whether Harry is bald* would *clearly* exclude Sophie's knowing that Harry is bald*. Of course, if *Excluded Knowledge* is not clearly true, then neither is either of the two standard views of vagueness that clearly entail it: vagueness-as-absence-of-truth and vagueness-as-ignorance.

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' for short. 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.

Before I begin, I need to make a point about the relationship between my case for LEM and my larger aim in this paper, which is to challenge the two standard views of vagueness by challenging *Excluded Knowledge*. I cannot in this limited space offer a decisive case for LEM. Suppose, then, that my case turns out to be unpersuasive. How will that affect my larger aim?

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? To begin with, all versions of vagueness-as-ignorance are committed to LEM. For, on these views, vagueness entails ignorance, and a thinker is *ignorant* of whether p only if either (i) p and the thinker does not know p , or (ii) not- p and the thinker does not know not- p . 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-absence-of-truth that are committed to LEM include a dominant family of views: those that appeal to the method of supervaluations in order to maintain LEM without Bivalence.⁴ So, if my defense of LEM is unpersuasive, this will have no effect on my argument that these versions of vagueness-as-absence-of-truth are clearly false. The only effect it will have is to undermine my argument that versions of vagueness-as-absence-of-truth that are not already committed to LEM are clearly false. These include certain many-valued treatments,⁵ certain intuitionist treatments,⁶ and perhaps a couple of recent psychological treatments.⁷

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 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. But if it is clear that Ning weighs something, and it is clear that she does not weigh anything other than 100 or 120 pounds, then it is clear that she weighs either 100 or 120 pounds⁸—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

⁴ Mehlberg 1958; Przelecki 1969; Lewis 1970; Fine 1975; Kamp 1975.

⁵ Halldén 1949; Körner 1960; Zadeh 1965; Tye 1994.

⁶ Putnam 1983.

⁷ Wright 2001; Field 2003a, 2003b.

⁸ For those readers who object to the idea that, if Ning weighs something, there must be some number, n , such that Ning weighs n pounds, I offer the following substitute: if Ning weighs something, there must be some true sentence of the form 'Ning weighs n pounds'.

clearly true even if none of its 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.

It is perhaps worth emphasizing that what I am offering here is an *intuitive* consideration in favor of the view that vagueness does not threaten LEM. Obviously, those who wish to reject this view must reject some piece of the preceding reasoning. The present point is that the preceding reasoning is intuitive. To reject some piece of it is to reject something that seems right, at least initially. Hence, we have an initial consideration for thinking that vagueness does not threaten LEM.⁹

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 what is required to qualify as rich, that is, it is vague where the cutoff for being rich is. But use of the expression ‘the cutoff for being rich’ seems to presuppose 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—seems to involve 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’ seems to presuppose that there is an answer to this

⁹ For other intuitive examples of clearly true disjunctions that lack clearly true disjuncts, see Fine 1975; McGee and McLaughlin 1995; Edgington 1996.

question. And if there is an answer to this question, then either that answer is that Harry is bald or that answer is that Harry is not bald; thus, either Harry is bald or Harry is not bald. Hence, commitment to what appears to be another basic datum seems to involve 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
- (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). 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, at least ordinarily, 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.¹⁰ 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, ordinarily, 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

¹⁰ I am not denying the ease with which we can imagine cases in which it is vague whether a given person is asserting a given proposition. Rather, I am claiming that, in an actual ordinary conversation, if someone were to *deliberately* try to make it the case that it was vague whether he or she was asserting p , this would require strange behavior that is unlikely to result in successful communication.

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 an argument in favor of weakly rejecting LEM.

Field (2003a, 2003b, 2010) 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.¹¹ 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:

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 suffices for being rich, yet being worth any less does not suffice?

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 doesn't it seem misguided to *speculate*, or to *hope*, that the cutoff is, say, \$1,045,009.01?

B: Initially, it does. But it's not clear that this intuition constitutes a problem for the idea that there is a cutoff, for it's not clear that the intuition is incompatible with the idea.

¹¹ In a similar vein, Smith (2005) suggests that it would be misguided to *guess* where the cutoff lies.

A: What do you mean?

B: For simplicity, let's focus on the case of speculation. Plausibly, the intuition that it is misguided to speculate that the cutoff is, say, \$1,045,009.01 derives from three others:

- (i) that it is metaphysically impossible to *successfully* speculate that the cutoff is, say, \$1,045,009.01;
- (ii) that it is easy to see (i);
- (iii) that, if it is easy to see that it is metaphysically impossible to attain a goal, then it is misguided to set the goal in the first place.

Now, the existence of a cutoff is clearly compatible with (iii). The only reason to suspect that it might be incompatible with (ii) is that it might be incompatible with (i). So let's see whether it's incompatible with (i).

Suppose that there is clearly a cutoff, as I maintain. In this case, is it metaphysically possible to successfully speculate that the cutoff is, say, \$1,045,009.01? To answer this question, we need to identify the conditions for *successful* speculation: clearly, p is necessary and sufficient for a given instance of speculating p to be *successful*. Because there is clearly a cutoff, there is clearly some cutoff claim—some claim of the form 'the cutoff is \$ n '—such that an act of speculating that claim is a success. Thus, it is clearly metaphysically possible to successfully speculate on where the cutoff lies.

A: So the existence of a cutoff *is* incompatible with (i).

B: Not so fast. The fact that it is clearly metaphysically possible to succeed in speculating on *where* the cutoff lies does not entail that there is a cutoff claim, p , such that it is clearly metaphysically possible to succeed in speculating p . Indeed, there is no such p . For it is vague where the cutoff lies. Each of a range of candidate cutoffs is such that there is vagueness as to whether *it* is the cutoff. If p lies in this range, then there is vagueness as to whether an act of speculating p is a success; if p lies outside, then an act of speculating p is clearly not a success. Either way, the act is not a *clear* success. Hence, there is no p such that an act of speculating p is a clear success. And, while the location of the cutoff is a *vague* matter, it is not a *contingent* matter: in a given context, 'rich' expresses a concept that is relativized to a particular distribution of wealth and thus has its cutoff essentially. So, if there is no p such that an act of speculating p is a clear success, then there is no p such that it is clearly metaphysically possible to successfully speculate p .

A: I see. So, given vagueness in the location of the cutoff, it is not *clearly* metaphysically possible to successfully speculate that the cutoff is, say, \$1,045,009.01. And so the existence of a cutoff is not *clearly* incompatible with (i).

B: That's right. And similar considerations show that the existence of a cutoff is not *clearly* incompatible with the initial intuition that it's misguided to speculate that the cutoff is, say, \$1,045,009.01. So it's not clear that the initial intuition constitutes even an *initial* problem for the view that there is a cutoff.

- A: But clearly there *is* a problem. For in the neighborhood of this initial intuition are two further intuitions, both of which *clearly are* incompatible with the existence of a cutoff:
- (iv) that it is *clearly* metaphysically impossible to successfully speculate on *where* the cutoff is;
 - (v) that, for any cutoff claim, *p*, it is *clearly* metaphysically impossible to successfully speculate *p*.

B: Even if (iv) and (v) are initially intuitive, they should be rejected. First consider (iv). It appears to conflict with the following intuition:

- (vi) that it's clearly *vague* where the cutoff is.

For, if it's clearly *vague where* the cutoff is, then—intuitively—there clearly is a cutoff. And if there clearly is a cutoff, then it's clearly metaphysically possible to successfully speculate on where the cutoff is. So the two intuitions appear to be incompatible.

A: Looks like we've got a stalemate.

B: I don't think so. First of all, before doing any theorizing about vagueness, (vi) would seem to be a basic datum. (iv), on the other hand, plausibly derives from the following intuition:

- (iv*) that there's clearly no cutoff.

And (iv*) has a plausible *surrogate*; we can satisfy our initial temptation to accept (iv*) by instead accepting:

- (iv*-s) that no number is clearly the cutoff.

By contrast, (vi) has no plausible surrogate. So we should reject (iv*) in favor of (vi). Once we reject (iv*), there is no good reason to accept (iv). Moreover, just as (iv*) has a surrogate, so does (iv):

- (iv-s) that it's metaphysically impossible for an act of speculating on where the cutoff is to *clearly* succeed.

We can satisfy our initial temptation to accept (iv) by instead accepting (iv-s). So there is no stalemate: we should accept (vi) and reject (iv).

A: And what about (v)?

B: The same strategy can be used to show that we should reject (v). For (v) conflicts with (vi); and (v), but not (vi), admits of a surrogate. The conflict was brought out above, when we saw that, given (vi), there is a candidate cutoff, *p*, such that it is clearly *vague* whether it is metaphysically possible to successfully speculate *p*. Obviously, this conflicts with (v). As for the surrogate for (v), our initial temptation to accept (v) can be satisfied by instead accepting (v-s):

(v-s) that, for any cutoff claim, p , it is metaphysically impossible to *clearly* succeed in speculating p .

Because (iv) has no plausible surrogate, there is no stalemate: we should accept (iv) and reject (v).

A: To summarize, then, according to you, it's not clear that the initial intuition *that it's misguided to speculate that the cutoff is, say, \$1,045,009.01* is incompatible with the idea that there's a cutoff. And, while the two other initial intuitions *are* clearly incompatible with the idea, both conflict with the core intuition that it's clearly vague where the cutoff is, and both admit of surrogates. So there is reason to reject both in favor of the core intuition. Hence, I've not yet given you a convincing reason to doubt the existence of a cutoff.

B: That's right.

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 will now try to give a strong counterexample to *Excluded Knowledge*.

4. *Excluded Knowledge* is clearly false

In this section, I 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* Harry is bald*, even though it is vague (i) whether Harry is bald*; (ii) whether Sophie knows *that* Harry is bald*; and (iii) whether Sophie knows *that* Harry is not bald*.

Recall that Sophie is carefully considering whether Harry is bald*. Sophie's cognitive conditions are ideal, and she fully understands the predicate 'is bald*'. Her population is unsusceptible to inconsistencies and instabilities of the relevant sort. Under present conditions, Sophie's intuitions are perfectly and robustly accurate indications of whether Harry is bald*. Given Sophie's complete mastery of the concept of negation, her intuitions are also perfectly and robustly accurate indications of whether Harry is not bald*. Moreover, Sophie has no reason, evidence, or justification for thinking that her intuitions are misleading. So, under present conditions:

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

- (6) If it seems to Sophie that Harry is not bald*, then Sophie knows that Harry is not bald*.

Given LEM and that it is now vague whether Harry is bald*, clearly, either Harry is bald* or Harry is not bald*, even though it is vague which. Clearly, if Harry is bald*, then it seems to Sophie that Harry is bald*. Clearly, if Harry is not bald*, then it seems to Sophie that Harry is not bald*. Thus, clearly, either it seems to Sophie that Harry is bald* or it seems to Sophie that Harry is not bald*, even though it is vague which. Given the clear truth of (1) and (6), clearly, either Sophie knows that Harry is bald* or she knows that Harry is not bald*, even though it is vague which. Because a thinker knows *whether p* just in case she knows *p* or she knows *not-p*, clearly, Sophie knows *whether* Harry is bald*.

To be sure, it might be conversationally inappropriate for Sophie to *say* that she knows whether Harry is bald*. For, by saying this, Sophie might imply that she is in a conversationally appropriate position to say whether Harry is bald*. And, as we saw in the preceding section, at least in any community of speakers similar to ours, it is typically *not* conversationally appropriate to assert *p* or to assert *not-p* when one knows that it is vague whether *p*. So, even though Sophie knows whether Harry is bald*, it might be conversationally inappropriate for her to publicly declare this knowledge. In a similar vein, it might be conversationally inappropriate for those speakers who know that Sophie has this knowledge to declare that Sophie has this knowledge. For such a declaration might itself suggest that Sophie is in a conversationally appropriate position to say whether Harry is bald*, and she might not be.

These facts might help to explain our initial discomfort with the idea that Sophie knows whether Harry is bald*. Not only does Sophie lack *clear* knowledge of whether Harry is bald*—she does not clearly know that he is, and she does not clearly know that he is not, bald*—it might be conversationally inappropriate for Sophie and others to *say* that Sophie has knowledge of whether Harry is bald*.

I have argued that it is clearly the case that Sophie knows whether Harry is bald* 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: vagueness-as-absence-of-truth and vagueness-as-ignorance.

Let me conclude by considering two worries about my argument.

The first concerns my characterization of how things seem to Sophie under the given conditions. I claimed that, at a certain point in the plucking, it is vague both whether it seems to Sophie that Harry is bald* and whether it seems to Sophie that Harry is not bald*. In such a case, we might say that it *sort of* seems to Sophie that Harry is bald* and that it *sort of* seems to her that Harry is not bald*. The worry is that, on one natural view, vagueness presents itself to our minds, not as vagueness in intuition, but rather as an *absence* of intuition, so that if it clearly seems vague whether p , it clearly does not seem that p and clearly does not seem that $\text{not-}p$. On this view, there is a clear absence of intuitions separating the cases where it *sort of* seems to us that p from the cases where it *sort of* seems to us that $\text{not-}p$. Taken as a view about the nature of vagueness and/or intuition, this view entails that it is impossible for it to *sort of* seem that p and simultaneously *sort of* seem that $\text{not-}p$, and thus that my description of how things seem to Sophie cannot be satisfied. The worry, then, is that this natural view might be correct.

To see that the view is mistaken, first consider an artificial example that plausibly *does* give rise to a clear absence of intuitions as to whether a certain condition suffices for the application of a certain predicate. Following Fine (1975), let us introduce the predicate ‘nice*’ by uttering only two sentences:

(7) n is nice* if $n > 15$;

(8) n is not nice* if $n < 15$ (for natural numbers n).

At least initially, it clearly seems that 16 is nice*; it clearly seems that 14 is not nice*; it clearly does not seem that 15 is nice*; and it clearly does not seem that 15 is not nice*. So, there is a clear absence of intuitions about whether being the number 15 suffices for being nice*. But this example is not an example of *vagueness*. It is not vague whether 15 is nice*. Perhaps it is *indeterminate*, but it is not *vague*.

The phenomenology of vagueness is quite different. In a clear case of perceived vagueness, there is not a clear absence of intuitions about whether p . Rather, there is a distinct sensation of *ambivalence*

about whether p : it *sort of* seems that p , and it *sort of* seems that not- p . In my experience, I have found that it is easiest for philosophers to recognize this ambivalence by considering a case of vagueness, not as to whether a given person is, say, bald, but rather as to whether it *seems* to us that a given person is bald. How does *this* vagueness present itself to ours minds? Imagine yourself in a state of mind that you would feel comfortable reporting by saying that it *sort of* seems to you that Harry is bald. Now ask yourself: does it seem to you that it seems to you that Harry is bald? Your candidate answers are as follows:

- (i) clearly so
- (ii) clearly not
- (iii) sort of.

(iii) is the obvious answer. Vagueness as to whether it seems to you that Harry is bald presents itself, not as an *absence* of intuition about whether it seems to you that Harry is bald, but rather as *vagueness* in intuition. Of course, if it is vague whether it seems to you that Harry is bald, then it is vague whether it does not seem to you that Harry is bald: it *sort of* does not seem to you that Harry is bald. Again, ask yourself: does it seem to you that it does not seem to you that Harry is bald? The obvious answer, once more, is (iii). So, if it is vague whether it seems to you that Harry is bald, then it *sort of* seems to you that it seems to you that Harry is bald and it *sort of* seems to you that it does not seem to you that Harry is bald.

There is no reason to think that this example is unique in this respect. If it seems vague whether Harry is bald, then—in ordinary conditions—it *sort of* seems that Harry is bald and it *sort of* seems that Harry is not bald. More generally, if it seems vague whether p , then—in ordinary conditions—it *sort of* seems that p and it *sort of* seems that not- p . As Stephen Schiffer (2000, 2003) and Mark Sainsbury (1986) have noted, appearances of vagueness as to whether p manifest themselves in terms of a distinctive, felt, *ambivalence* about whether p . I conclude that the view that vagueness presents itself as an absence of intuitions is false.

Why, then, is this view *natural* in the first place? I suspect that its initial appeal derives from our failure, once more, to carefully distinguish between being *F* and being clearly *F*. In the present case, we might misreport the sound intuition, (9), as the unsound intuition, (10):

- (9) If it seems vague whether *p*, then it does not clearly seem that *p* and it does not clearly seem that not-*p*.
- (10) If it seems vague whether *p*, then it does not seem that *p* and it does not seem that not-*p*.

The second worry about my argument is 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-absence-of-truth, what it *is* for it to be *vague* whether *p* is for there to be no truth of the matter whether *p*, for some distinctive reason. On vagueness-as-ignorance, what it *is* for it to be *vague* whether *p* is for it to be epistemically, but not metaphysically, unsettled whether *p*, for some distinctive reason. On both theories, 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 be sure, a proponent of one of the two standard theories of vagueness might retreat from her original position by moving to one of the two corresponding weaker claims, that ordinary cases of vagueness happen also to be cases of missing truth, or that ordinary cases of vagueness happen also to be cases of ignorance. But she cannot maintain her original theory that ordinary cases of vagueness are cases of vagueness *in virtue* of being cases of missing truth or ignorance. What is needed, if my counterexample is sound, is an alternative to the two standard theories of the nature of vagueness. Such a theory should not entail that vagueness by nature excludes knowledge.

On my own view, vagueness is neither a type of missing truth 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 arguments against *Excluded Knowledge*, for nothing in the idea that vagueness is *sui generis* suggests that vagueness as to whether *p* excludes knowing whether *p*.^{12,13}

¹² I develop and motivate this view in my 2009.

¹³ Earlier versions of this paper were presented at University of Missouri-Columbia, New York University, University of Vermont, and Yale University. For helpful comments and discussion, I am grateful to members of

REFERENCES

- Barnett, D. 2009. Is Vagueness Sui Generis? *Australasian Journal of Philosophy* 87: 5-34.
- Bealer, G. 1999. A Theory of the A Priori. *Philosophical Perspectives* 13: *Epistemology*: p. 31.
- BonJour, L. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Burge, T. 2003. Perceptual Entitlement. *Philosophy and Phenomenological Research*: 503-548.
- Burgess, J.A. 1990. The Sorites Paradox and Higher-Order Vagueness. *Synthese* 85: 417-474.
- Campbell, R. 1974. The Sorites Paradox. *Philosophical Studies* 26: 175-191.
- Cargile, J. 1969. The Sorites Paradox. *British Journal for the Philosophy of Science* 20: 193-202.
- 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. *Noûs*: 1-30.
- . 2003a. No Fact of the Matter. *Australasian Journal of Philosophy* 81: 457-80.
- . 2003b. The Semantic Paradoxes and the Paradoxes of Vagueness. In Beal and Glanzberg, eds., *Liars and Heaps*. Oxford: Oxford University Press: 262-311.
- . 2010. This Magic Moment: Horwich on the Boundaries of Vague Terms. In Dietz and Moruzzi, eds., *Cuts and Clouds: Vagueness, its Nature and its Logic*. Oxford: Oxford University Press.
- Fine, K. 1975. Vagueness, Truth and Logic. *Synthese* 30: 265-300.
- Halldén, S. 1949. *The Logic of Nonsense*. Uppsala: Uppsala Universitets Arsskrift.
- Horwich, P. 1990. *Truth*. Oxford: Blackwell.
- Huemer, M. 2001. *Skepticism and the Veil of Perception*. Lanham, Maryland: Roman and Littlefield.
- . 2007. Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research* 74: 30-35.

those audiences and to Yuval Avnir, George Bealer, David Christensen, Sinan Dogramaci, Hartry Field, Kit Fine, Mike Huemer, Adam Pautz, and Stephen Schiffer.

- Kamp, J.A.W. 1975. Two theories about adjectives. In *Formal Semantics of Natural Language*, edited by E. L. Keenan. Cambridge: Cambridge University Press.
- Keefe, R., and P. Smith, eds. 1996. *Vagueness: A Reader*. Cambridge, MA: MIT Press.
- Körner, S. 1960. *The Philosophy of Mathematics*. London: Hutchinson.
- Lewis, D.K. 1970. General semantics. *Synthese* 22: 18-67.
- McGee, V., and B. McLaughlin. 1995. Distinctions without a Difference. *Southern Journal of Philosophy* 33 (Supplement): 203-51.
- Mehlberg, H. 1958. *The Reach of Science*. Toronto: University of Toronto Press.
- Przelecki, M. 1969. *The Logic of Empirical Theories*. London: Routledge and Kegan Paul.
- Pryor, J. 2000. The Skeptic and the Dogmatist. *Noûs*: 517-549.
- Putnam, H. 1983. Vagueness and alternative logic. *Erkenntnis* 19: 297-314.
- Rosen, G. 2001. Nominalism, Naturalism, Epistemic Relativism. *Philosophical Perspectives* 15: 69-91.
- 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.
- . 2003. *The Things We Mean*. Oxford: Oxford University Press.
- Smith, N. 2005. Vagueness as Closeness. *Australasian Journal of Philosophy* 83: 157-183
- Sorensen, R.A. 1988. *Blindspots*. Oxford: Clarendon Press.
- . 2001. *Vagueness and Contradiction*. Oxford: Oxford University Press.
- Sosa, E. 1991. *Knowledge in Perspective*. Cambridge: Cambridge University Press.
- Tye, M. 1994. Sorites Paradoxes and the Semantics of Vagueness. In *Philosophical Perspectives 8: Logic and Language*, edited by J. E. Tomberlin. Atascadero, CA: Ridgeview.
- Williamson, T. 1994. *Vagueness*. New York: Routledge.
- Wright, C. 2001. On Being in a Quandry. *Mind* 110: 45-98.
- Zadeh, L.A. 1965. Fuzzy sets. *Information and Control* 8: 338-353.