Vagueness, conditionals, and context-sensitivity

By Tom Beevers

Penultimate version

Abstract: I argue that practically all vague language is context-sensitive in a covert and unfamiliar way. I first outline a novel puzzle concerning the interaction of conditionals and vagueness. I then argue that the best way of resolving the puzzle is through positing context-sensitive penumbral connections between sundry parts of language. I argue that these penumbral connections shift through a distinct form of Lewisian accommodation. The upshot is that meaning is a far shiftier thing than has typically been thought.

1. Introduction

STRANGER. Sarah has just arrived at her first day at work and has just met her new colleague, Pete. He has complained about being a little peckish and seems a bit drowsy. She doesn't think much of it, however. Whilst he's not clearly wide-awake or full, he's not clearly sleepy or hungry either. After speaking to him for a while, Sarah recognises him as being borderline hungry and borderline sleepy.

Now consider the following statements:

A. ?? Pete is hungry and full.
B. ? Pete is hungry and sleepy.
C. ? Pete is hungry and not sleepy.
D. ? Pete is not hungry and sleepy.
E. ? Pete is not hungry and not sleepy.

Saying each of (A)-(E) would sound bad, but (A) would sound considerably worse. Whilst (B)-(E) seem borderline, (A) seems clearly false.

Classical theories of vagueness (i.e., epistemicism, supervaluationism, and other related theories) accommodate these judgements by appealing to penumbral connections.[1] Penumbral connections can be roughly understood as constraints on the possible precise things linguistic expressions could mean.[2] If 'this apple is red' is clearly true, then a penumbral connection would be that it's not possible for 'red' to mean something which would make 'this apple is red' false. A different kind of penumbral connection coordinates the precise possible contents of *different* linguistic expressions. One such

penumbral connection requires that no precise content of 'hungry' and 'full' can jointly apply to the same object. This ensures that (A) is not true at any admissible precisification and is thus determinately false. However, there is no such coordination between the possible meanings of 'hungry' and 'sleepy'. They are penumbrally disconnected - any precise possible meaning of one can be combined with another. This ensures that (B)-(E) are borderline since the indeterminacy of the conjuncts transfers to the conjunction.

Classical theories of vagueness thus have no problem in explaining why the following seems true given the above case:

- Statement 1. It's not determinately false that (Pete is hungry and not sleepy)

This is just to say (C) is not determinately false. However, we can put pressure on our judgement here by considering a different situation unfolding at the same time as Stranger.

FRIEND. Fred is Pete's best friend. He knows Pete has a hormonal condition where whenever he gets hungry, he becomes sleepy. The hungrier he gets, the sleepier he gets. Fred doesn't work with Pete and doesn't know how he's doing right now, but knows enough to say:

F. If Pete's hungry, then he's sleepy.

This assertion seems clearly true. Thus, it seems we should accept:

- Statement 2. Determinately (if Pete's hungry, then he's sleepy).

But now we can derive a contradiction. Given the plausible principle that 'if A, C' entails the corresponding material conditional 'it's not the case that A and not-C' and the principle that determinacy is closed under entailment, we can infer that (C) is determinately false. This contradicts Statement 1.

How do we resolve the contradiction? Two options present themselves:

Option 1. Hold that Statement 1 is true and Statement 2 is false.

Option 2. Hold Statement 1 is false and Statement 2 is true.

The first option implies that we were wrong to judge (F) to be determinately true, leading to a revision of our judgement regarding this conditional. The second implies that we were wrong to judge (C) as being borderline, leading to a revision of the judgements that makes classical theories of vagueness attractive. Neither option seems *prima facie* appealing, and I will argue that both are ultimately unviable.

Fortunately, there is another. We can sometimes reconcile two seemingly contradictory claims through appealing to context-sensitivity. For example,

although the sentence 'I am not in Belize' is the linguistic negation of 'I am in Belize', there are clearly consistent assertions of both. Given 'I' picks out whoever asserts it within the context of utterance, two assertions of 'I am in Belize' and 'I am not in Belize' can both be true provided the first is asserted by someone in Belize and the second asserted by someone who is not. Both statements will be true relative to their respective contexts of assertion. We might similarly avail ourselves of context sensitivity to avoid the problematic consequences of Options 1 and 2:

> Option 3. Contextualism. Hold that Statement 1 is true in Stranger but false in Friend and hold that Statement 2 is false in Stranger but true in Friend.

This will allow an assertion of Statement 1 to be true in Stranger *and* an assertion of Statement 2 to be true in Friend. It will be the task of this paper to persuade you that this third option is correct.

In the proceeding discussion, I will use $'>'$ to stand for a generic conditional operator which can be read either as an indicative or counterfactual conditional (the difference between the two will be immaterial here).[3] I will rely only on the following principles:

1. Factivity: $\mathrm{Determinately(A)} \rightarrow \mathrm{A}$.
2. Normality: Determinacy obeys a normal modal logic.
3. Material Conditional Entailment (MCE): 'if A then C' implies 'it's not the case that A and not C'.

Factivity says that what is determinate is true. Normality tells us that modus ponens for the material conditional is valid, that determinacy distributes over the material conditional (so $\mathrm{Determinately}(A \rightarrow C) \rightarrow (\mathrm{Determinately}(A) \rightarrow \mathrm{Determinately}(C))$ is valid), that the classical propositional tautologies are valid, and that the logical necessities are determinate at all orders. Normality implies that determinacy is closed under entailment. MCE tells us that from $A > C$ we can infer the material conditional $A \rightarrow C$ (or equivalently, 'not-A or C' or 'not-(A and not-C)'). I will not justify these principles except to say almost all classical theories of vagueness[4] and almost all theories of conditionals[5] have adopted them. I will take them for granted here.

To argue for Option 3, I will first argue against Option 1 and Option 2. In §2 I show that Option 1 is unacceptable since it implies it is implausibly difficult for conditionals to be determinately true. In §3 I show that Option 2 is unacceptable since it implies it is implausibly difficult to judge whether a complex sentence is determinately true or not. In §4 I outline and defend Option 3. I argue that Statement 1 is true in Stranger but false in Friend, and the reverse is true for Statement 2. I argue that the shift in contextual parameters in (C) and (F) can be attributed to a form of accommodation

(Lewis, 1979). I conclude in §5 by showing that we can generalise my arguments to show that practically all vague language is context-sensitive in a covert and unfamiliar way. This makes meaning a far shiftier thing than has been typically thought.

2. Option 1.

Option 1 tells us we should flatly accept Statement 1 and reject Statement 2. One reason for why this might seem attractive is that it does not require any revisionary implications regarding (A)-(E). We can preserve the intuitions which make classical theories of vagueness attractive. It does, however, lead to the unexpected judgement that the conditional (F) is not determinately true. Let's carefully examine why this.

Let's say that A and C are *backwardly distributive* when borderlineness in A and borderlineness in C leads to borderlineness in $A \rightarrow C$. In Stranger we saw that when 'Pete is sleepy' and 'Pete is hungry' were borderline, (C) also seemed borderline. Something is borderline iff its negation is borderline, and indeed, in Stranger the negation of (C) seemed borderline as well. Thus, in the context of Stranger, it seems that 'Pete is sleepy' and 'Pete is hungry' are backwardly distributive. This is why it seemed intuitively that Statement 1 is true. Statement 1 implies that Statement 2 is false given MCE and Normality since it follows from these principles that if $A \rightarrow C$ is not determinately true, then $A > C$ is not determinately true as well. Thus, if we are to accept Statement 1, we are immediately committed to Statement 2 being false.

We can thus see that there is a tension between our judgements of backwards distributivity and our use of conditionals. This is a consequence of the following fact which we can derive from our assumptions:

> Fact 1: If A and C are backwardly distributive, then $A > C$ is determinately true only if A and B aren't both borderline.[6]

Given Fact 1 and the fact that Peter is borderline hungry and borderline sleepy, the backwardly distributivity of 'Peter is hungry' and 'Peter is sleepy' implies Statement 2 is false.

At first glance, this is an unwelcome result. But perhaps this is as it should be. Fred's assertion of (F) did not seem indeterminate when he asserted it. But one might think the fact that (as it turned out) Pete was borderline hungry and borderline sleepy showed this judgement to be mistaken. In support of this, one might point to the fact that asserting (F) would seem infelicitous if Fred knew that Pete was borderline hungry and borderline sleepy. One might take this as evidence that (F) is not determinately true and that Statement 2 is thus false. Perhaps the first option isn't so bad, after all.

However, on reflection, this explanation does not work. Fred had no idea how hungry or sleepy Pete was when he asserted (F) in Friend. It was thus entirely consistent with what he knew that Pete was borderline hungry and borderline sleepy, yet this still didn't seem to affect his grounds for asserting (F). In fact, we could imagine Fred engaging in the following exchange:

> Interjector. What if Pete is borderline hungry and borderline sleepy?
>
> Fred. He might be for all I know. But clearly, if he is hungry, he is sleepy.

There is nothing wrong with what Fred says here. Yet this would be puzzling were Fred's assertion not determinately true due to having a borderline antecedent and borderline consequent. Fred would be simultaneously claiming that (F) is determinately true but simultaneously asserting something which implies immediately that (F) might be not determinately true. His assertion would be unassertible and a kind of Moorean statement, but it isn't.[7]

Moreover, the fact that asserting (F) would be infelicitous in Stranger does not show that it is not-determinately true when asserted by Fred. To see this, suppose Barry is borderline bald. It would be somewhat strange for anyone who knows this to say, 'Barry might be bald'. However, if someone had previously asserted the same sentence when their evidence did not rule out that Barry was clearly bald, their assertion would seem clearly true. Analogously, they may have said 'Barry might be borderline bald for all I know, but it's still true that he might be bald'. Likewise, the fact that (F) would sound infelicitous when Fred knows Pete is borderline hungry and sleepy, does not mean that his assertion is not determinately true when asserted in a context where he didn't know how hungry and sleepy Pete was. Indeed, if Fred found out that Pete was borderline hungry and sleepy after asserting (F), he would not feel under obligation to retract or apologise for what he said.[8] Thus, it seems the linguistic evidence supports Fred's assertion being determinately true.

This option also makes it implausibly difficult to know whether a conditional is determinately true. If this option were correct, Fred must be able to rule out that Pete is jointly borderline hungry and borderline sleepy to know whether (F) is determinately true. But this sets an implausibly high bar for knowledge. Fred was not in such a position when he asserted (F), yet it seems he could know this to be determinately true.

This problem becomes even more acute when we consider analogous cases involving contrary conditionals (conditionals of the form $A > C$ and $\neg A > \neg C$).[9] If we flatly accept that 'Pete is hungry' and 'Pete is sleepy' are borderline distributive, we will be committed to doing likewise in other cases.

For instance, if the only thing we know about Hailey is that she is borderline happy and we also know that England is borderline doing well in the World Cup, we would judge 'England is doing well in the World Cup and Hailey is unhappy' to be borderline. Thus, it seems 'England is doing well in the World Cup' and 'Hailey is happy' are backwardly distributive. The same is true of any given person, S. If we only know of S that 'S is happy' is borderline and we know that 'England is doing well in the World Cup' is borderline, then these statements will seem backwardly distributive. But if this is true across the board, then one could only know 'if England is doing well in the next football World Cup, S will be happy and if England isn't doing well, S will be unhappy' to be determinately true if one can rule out both that England won't borderline do well and that S won't be borderline happy. This is a consequence of the following fact:

> Fact 2. If A and C are backwardly distributive, then $A > C$ and $\neg A > \neg C$ are determinately true only if *neither A nor C* are borderline.[10]

Given 'S is happy' and 'England is doing well in the World Cup' are backwardly distributive, Fact 2 implies that to know that 'if England is doing well in the next football World Cup, S will be happy' and 'if England isn't doing well, S will be unhappy' to be determinately true one must be able to rule out that England will borderline do well in the next World Cup. But given no one can rule out that England will borderline do well in the next World Cup (at the time of writing), it seems we cannot know any conditionals of this form to be determinately true.

This seems bizarre – we *can* know these conditionals are determinately true of some people. If we know Hailey is a diehard England football fan, then it will be entirely obvious that if England is doing well in the World Cup, Hailey will be happy and if England isn't doing well in the World Cup, Hailey will be unhappy. Thus, it seems our judgements of backwards distributivity cannot be categorically maintained if we are to maintain a descriptively adequate theory of conditionals.

3. Option 2.

Let's turn now to Option 2. Through taking up this option, we can avoid the counterintuitive consequences of Option 1. We revise our judgements of backwards distrbutivity by holding that 'Pete is hungry and not sleepy' is determinately false when Pete is borderline hungry and borderline sleepy. This commits us to taking Statement 1 to be invariantly false. This allows (F) to be determinately true without immediately implying that Pete cannot be borderline hungry and sleepy. Likewise, it allows 'if England is doing well in the next World Cup, S will be happy' and 'if England isn't doing well, S will

be unhappy' to be determinately true without this implying that England won't do borderline well at the World Cup and S won't be borderline happy.

One unexpected theoretical consequence for classical theories of vagueness is that it leads to a wider range of penumbral connections than might initially have been thought. Whilst we may have first thought that predicates like 'hungry' and 'sleepy' are penumbrally disconnected, this turns out not to be so. I do not think this constitutes a decisive objection – the notion of penumbral connectedness is too theoretical to elicit any firm intuitions. However, I think there is a decisive objection around the corner which does not appeal to any theoretical machinery.

Consider again the statement (C) – 'Pete is hungry and not sleepy'. This seemed borderline in Stranger, but determinately false in Friend. The defender of Option 2 tells us to take the latter judgement to be correct and the former to be incorrect. Given Sara had only just met Pete, she did not know about his hormonal condition and thus whether his hunger and his tiredness were related at all. At the risk of stating the obvious, it's possible for someone to be hungry without this making them sleepy. Thus, if Option 2 is correct, she should have reserved judgement on whether (C) was borderline or not. But this seems bizarre. It does not seem like her judgement should depend on whether Pete's hunger and sleepiness are causally dependent on each other.

Taking this approach generally will mean that whenever we are faced with a logically complex and contingent statement, it will typically be impossible to tell whether it is borderline or not unless we know the relevant causal/information relationship between the parts. For example, suppose England is borderline doing well in the World Cup and Hailey is borderline happy. If we take the same approach as we did with Stranger here, we cannot make any judgement about whether 'England is doing well in the World Cup and Hailey is happy' or 'England is doing well in the World Cup and Hailey is unhappy' are borderline or not until we know whether 'if England is doing well in the World Cup, Hailey is happy' and 'if England is not doing well in the World Cup, Hailey is unhappy' are determinately true. Perhaps England doing well in the World Cup would make Hailey happy, perhaps it would make her unhappy, perhaps she doesn't care. Until we know which of these obtains we must reserve judgement on these logically complex statements. This option thus makes our judgements of determinacy and indeterminacy hostage to the conditional facts. Given we are frequently unable to discern these relationships, we are frequently not in a position to know the determinacy conditions of many complex statements. This makes the determinacy conditions of logically complex statements implausibly difficult to know.

Thus, in sum, both options 1 and 2 lead to unacceptable consequences. The former makes it implausibly difficult for conditionals to be determinately true, and thus difficult to know whether they are determinately true. The latter makes it implausibly difficult to know whether logically complex statements are determinately true or not. Both lead to unappealing error theories of different kinds.

4. Option 3.

If we are to avoid these consequences, we seemingly need to accommodate our judgement that 'Pete is hungry' and Pete is sleepy' are backwardly distributive when the causal/informational facts are out of focus, as well as our judgement that (F) is determinately true when they are salient. We thus need Statement 1 to be true in Stranger and Statement 2 to be true in Friend. But given these are incompatible statements, we will need to appeal to some kind of context-sensitivity if we are to allow this to be so. But what kind of context-sensitivity could this be? In this section, I will suggest that we can attribute this context-sensitivity to a distinct kind of accommodation first proposed by Lewis (1979, p.339-40).[11]

According to Lewis (1979) and many others since, contextual parameters will often shift charitably to ensure an assertion is felicitous. Accommodation refers to this general phenomenon. Accommodation occurs through a variety of more specific linguistic mechanisms. For example, imagine I refuse a stranger's offer of a cigarette by saying that I quit smoking a few years ago. It seems like my assertion may be perfectly acceptable despite presupposing the (from the stranger's perspective) unknown fact that I used to smoke. Lewis (1979, p.339-40) suggests that this is because of the following conversational rule:

> Rule of Presupposition: If at time t something is said that requires presupposition P to be acceptable, then if P is not presupposed just before t, then *ceteris paribus* and within certain limits - presupposition P comes into existence at t.

In accordance with this rule, the presupposition 'I used to smoke' is added to the common ground immediately at the point of assertion to prevent presupposition failure and the corresponding infelicity.

Lewis (1979, p.351-54) proposes a different kind of accommodation involving vagueness. He suggests that under low standards of precision, 'Italy is boot-shaped' can count as 'true enough' meaning we would judge an assertion of this sentence to be acceptable. Moreover, even if the standards are not initially low, asserting this sentence may update the conversational score, relaxing the standards of precision and allowing the assertion to be felicitous.

This case and Lewis' treatment of it is controversial. However, we can make the basic idea less so by finding a less extreme case and substituting the tendentious 'true enough' with the widely used 'determinately true'. Suppose that we are talking to Tim who is borderline thin. Suppose further that I assert 'Tim is thin' in conversation with you. Whilst it's possible that you would judge my assertion to be borderline and thus infelicitous, if you are feeling cooperative, you might let me 'get away with it' and allow the standards for thinness to drop, rendering my assertion determinately true. Henceforth people who are at least as thin as Tim would count as determinately thin. There are thus two different things which might follow from my assertion:

1. Accommodation. The standards for 'thin' decrease, 'Tim is thin' becomes determinately true, and my assertion is acceptable.
2. Refusal. The standards for 'thin' remain the same, 'Tim is thin' remains borderline, and my assertion is unacceptable.[12]

Although Lewis does not suggest a conversational rule to explain how this kind of accommodation can occur, we might formulate one like so:

> Rule of Penumbral Connections: If at time *t* '*A*' is asserted which requires some type of penumbral connections for *A* to be acceptable, then if those penumbral connections didn't exist in the time immediately prior to *t*, then *ceteris paribus* and within certain limits - they come into existence at *t*.[13]

Let's call accommodation which occurs through applications of this rule, *penumbral accommodation*. Penumbral accommodation allows us to explain how the conversational score can shift to ensure my assertion comes out determinately true. Whilst before it was possible that 'thin' could have meant something which would have rendered 'Tim is thin' false, after penumbral accommodation has occurred, this is no longer so. Thus, rendering this sentence determinately true.

Let's now see how this same kind of accommodation can be used to account for the context-sensitivity in Stranger and Friend. Let's first take Sarah's context in Stranger. We will maintain the standard story about why Statement 1 seems true in this context. In Stranger, 'hungry' and 'sleepy' are not penumbrally connected. It will thus be possible to precisify these predicates in a way that makes them both apply, not apply, or only one or the other apply to someone who is borderline hungry and borderline sleepy. Because Pete is borderline hungry and borderline sleepy, this ensures that (B)-(E) are borderline and Statement 1 is true.

Now let's take the context in Friend when Fred asserted the conditional (F). If we maintained that the penumbral connections in Fred's context were the

same as they were in Stranger, then Fred's assertion would be unacceptable. Given Fact 1, his assertion could not be determinately true. We would thus be in no better position than if we had pursued Option 1. But provided everything is equal and within certain limits, The Rule of Penumbral Connections tells us that the penumbral connections must shift to allow Fred's assertion of (F) to be acceptable. Thus, this rule can explain how the context shifts to allow both our judgements in Stranger and Friend to be correct.

How is this? Well, a natural suggestion is that the causal relationship between Pete's hunger and his sleepiness provided the right conditions for penumbral accommodation to occur.[14] Fred's assertion of (F) then raised this relationship to salience, meaning that the precisifications on which 'Pete was hungry and not sleepy' were true were rendered inadmissible, ensuring that 'Pete is hungry and not sleepy' was determinately false and Statement 2 was true. Thus, just as an assertion of 'Tim is thin', can precipitate a shift in the penumbral connections involving 'thin', an assertion of 'if Pete is hungry, he's sleepy' can precipitate a shift in the joint penumbral connections between 'hungry' and 'sleepy'.

These penumbral connections are not quite the same as the ones between 'hungry' and 'full. The penumbral connections between 'hungry' and 'full' are context-invariant in the sense that there is no context in which we can precisify 'hungry' and 'full' to apply to the same object. By contrast, the penumbral connections between 'hungry' and 'sleepy' in Friend must dissipate as soon as the causal/information relationship between Pete's hunger and tiredness loses salience. For if Fred was later informed that Peter is borderline hungry and sleepy and the causal/informational relationship between Pete's hunger and anger was not salient, he should seemingly judge 'Pete is hungry and not sleepy' to be borderline, not determinately false. In this way, Fred's assertion of (F) only instigates a temporary constraint on the meaning of its non-logical parts.

The appeal to these ephemeral penumbral connections validates both our judgement that Statement 1 is true in Stranger and our judgement that Statement 2 is true in Friend, avoiding the problematic consequences of options 1 and 2. Indeed, it seems to me to be the only way of avoiding an error theory whilst maintaining a non-logically revisionary theory of vagueness and conditionals. This, I think, provides a strong argument for us embracing option 3. I will conclude by examining the implications of this proposal for a broader debate in the philosophy of language.

5. Conclusion

Many philosophers and linguists have tried to downplay the role of context-sensitivity in natural language. The most extreme example of this is found in the work of the semantic minimalists (see Cappellen and Leopore, 2008; Borg, 2006) who argue that only a basic set of indexicals like 'that', 'now', and 'I' are context-sensitive. At the other end the spectrum, the radical contextualists hold that the contents of most linguistic expressions are context-relative and that the meaning of context-sensitive expressions will often shift in covert, linguistically unmarked, ways. Most theorists, I venture, would fall somewhere in the middle of these two extremes. Along with the set of indexicals, they would allow a further limited set of gradable adjectives like 'big' and 'red' into the class of context-sensitive expressions, provided they shift in an overt and relatively well-understood way. However, if my proposal is right, practically all vague expressions (including controversial examples like 'know') are context-sensitive in a covert and unfamiliar way. This pushes us firmly towards the radical contextualist end of the spectrum.

Let me unpack the claim that the form of context-sensitivity posited above is of a covert kind. Before I outlined the case for Option 3, one might have naturally thought the following to be true:

1. 'Pete', 'hungry' and 'sleepy' have the same interpretations in Stranger as they do in Friend.
2. 'Pete is hungry' and 'Pete is sleepy' have the same interpretations in Stranger as they do in Friend.
3. 'Pete is hungry and not sleepy' has the same interpretation in Stranger as they do in Friend.[15]

After all, there is no change in standard of comparisons, standards of precision, presuppositions, or any other obvious context-sensitive factors between Stranger and Friend. There are no unarticulated constituents present in one context and absent in the other. On the face of it, 'Pete', 'hungry' and 'sleepy' are used in exactly the same way across both contexts. However, according to my proposal, 1-3 are all false.

Why is this? Well, the set of contents which 'Pete', 'hungry', and 'sleepy' could jointly express (i.e., different precise objects and properties) must be different in Stranger than it is in Friend. The set of contents (i.e., propositions) 'Pete is hungry' and 'Pete is sleepy' could jointly express will shift from c_1 to c_2. And thus, the set of contents which 'Pete is hungry and not sleepy' could express (i.e. propositions) Stranger will be different to what it is in Friend. This difference in interpretation is exactly what was necessary to allow Statement 1 to be true in Stranger and Statement 2 to be true in Friend.

An epistemicist who is uncomfortable with radical contextualism might object that the shift in the set of contents that 'Peter is hungry and not sleepy' could express does not entail that the actual content of the sentence has shifted.[16] For the epistemicist holds that the content of a sentence like 'Pete is hungry and not sleepy' is just the proposition it expresses relative to the actual precisification (and possibly some other designated indices). For the epistemicist, even though it's possible that the proposition expressed by this sentence shifts from Stranger to Friend, it is not guaranteed that it will. Likewise, they may hold that the shift in the set of contents which 'sleepy', 'hungry', and 'Pete' could jointly express between Stranger and Friend, does not necessarily entail that even the set of contents they individually could express has shifted.[17] They may thereby maintain that penumbral accommodation may not involve any shiftiness in content.

This is an unsatisfying response to the radical contextualist. Even if the proposition expressed by 'Peter is hungry and not sleepy' remains the same across Stranger and Friend, the truth-value of sentences within which it is embedded (e.g. 'it's borderline whether Peter is hungry and not sleepy') *must* shift.[18] This undermines anti-contextualist arguments emphasising the need for a stable interpretation of linguistic items for semantic comprehension and communication to be possible. Moreover, it is at any rate clear that the epistemicist *will* be committed to some covert shifts in content due to penumbral accommodation. To see this note that, like other forms of context-sensitivity, penumbral accommodation does not by itself bring about any kind of epistemologically privileged context. It is not as if the context brought about by penumbral accommodation is necessarily more epistemically informed than the context which came before it. For, as we said above, if Fred subsequently found out that Pete was borderline hungry and sleepy and the causal relationship between Pete's hunger and tiredness was not salient, he too would plausibly judge 'Pete is hungry and not sleepy' to be borderline. So, knowing what we know, it must be a genuine possibility on the epistemicist picture that 'Pete is hungry and not sleepy' was actually true in Stranger, even though this sentence is definitely false in Friend. If a sentence is true in one context, and false in another, there must be a shift in the content of some of its subsentential parts. Thus, for the epistemicist, it must be a genuine possibility that the content of at least one of 'Pete', 'hungry' and 'sleepy' has shifted between Stranger and Friend. Therefore, given the ubiquity of penumbral accommodation, on the epistemicist picture, covert shifts in content will be inevitable in some cases of penumbral accommodation.

Let's now turn to my claim that practically *all* vague expressions are context-sensitive. To find out if a statement, *A*, is context-sensitive through penumbral accommodation we need only conduct the following test. We first find a conditional with *A* in the antecedent or consequent where this

conditional would seem determinately true even if its antecedent and consequent were borderline.[19] Secondly, we must ensure that, when viewed independently of the causal/informational facts, the antecedent and consequent of the conditional seem backwardly distributive. If we do this then, we can conclude that *A* must be context-sensitive since shifts in the interpretation of *A* are required to ensure that the conditional is determinately true in the first context, but not in the second. Given it is plausible that practically every vague sentence that is not logically or conceptually true or false will pass this test, we can infer every such sentence is context-sensitive.

Take for example, the statement 'John knows that there is no largest prime'. Imagine John was in a class in which the teacher demonstrated that there was no largest prime. If you don't know whether John was listening to the teacher, it would then seem like the claim 'if John was listening in class, he would know there is no largest prime' would be determinately true even if it turned out that John was borderline listening in class and only borderline knows that there is no largest prime. But seen in isolation of the causal/informationally relevant facts, 'John was listening in class' and 'John knows that there is no largest prime' would seem backwardly distributive. If John was borderline listening in class and borderline knew that there was no largest prime, 'John was listening in class and doesn't know there is no largest prime' would seem borderline. Thus, we need penumbral accommodation to explain why this sentence is determinately false in the first context, but not in the second. We can thus infer 'John knows there is no largest prime' is context-sensitive.

We can also easily extend this test to discover the context-sensitivity of some subsentential expression, *K*. All we need to do is to find a sentence *A* containing *K* where *A* passes the above test (and so is context-sensitive) and where the context-sensitivity must be attributed (at least in part) to *K*. For example, given 'there is no largest prime' is clearly context-insensitive, the only possible loci of context sensitivity in 'John knows there is no largest prime' is 'John' or 'knows'. It is implausible to suppose that all context-sensitivity can be attributed to 'John'. So, it seems 'knows' itself must be context-sensitive, contrary to what many theorists have argued (e.g. Williamson, 2005; Stanley 2004). Although this doesn't establish that 'know' is context-sensitive in the specific ways epistemic contextualists have claimed, it does provide a picture of meaning that is difficult to square with invariantism of various stripes. Much has been made of contextualism's supposed cost in preventing the free transmission of assertions by memory, testimony and so forth in the same linguistic form. These purported costs have been used to undermine contextualist accounts of knowledge (see Williamson, 2005 p.100-102), conditionals (Williamson, 2020, p.82), and

language generally (Cappelen and Leopore, 2008). But if the meaning of all linguistic items (including 'know') is constantly shifting from context to context together with the salient causal/informational facts, it seems this cost cannot be as expensive as has been claimed.

Thus, in sum, if we are to maintain both a non-logically revisionary theory of vagueness and a non-logically revisionary conditionals, mass context sensitivity is something we can and must learn to live with. Meaning must be a fleeting thing, indeed.[20]

Bibliography

Bacon, Andrew. *Vagueness and Thought*. Oxford: Oxford University Press, 2018.

Barnett, David. 'Is Vagueness Sui Generis?' *Australasian Journal of Philosophy* 87, no. 1 (1 March 2009): 5–34. https://doi.org/10.1080/00048400802237376.

Borg, Emma. *Minimal Semantics*. Oxford: Oxford University Press, 2004.

Cappelen, Herman, and Ernest Lepore. *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Oxford: John Wiley & Sons, 2008.

Edgington, Dorothy. 'Vagueness by Degrees'. In *Vagueness: A Reader*, edited by Rosanna Keefe and Peter Smith. Cambridge, Massachusetts: MIT Press, 1997.

Hawthorne, John. *Knowledge and Lotteries*. Clarendon Press, 2004.

Incurvati, Luca, and Julian J. Schlöder. 'Meta-Inferences and Supervaluationism'. *Journal of Philosophical Logic* 51, no. 6 (1 December 2022): 1549–82. https://doi.org/10.1007/s10992-021-09618-4.

Keefe, Rosanna. *Theories of Vagueness*. Cambridge: Cambridge University Press, 2000.

Kolodny, Niko, and John Macfarlane. 'Ifs and Oughts'. *The Journal of Philosophy* 107, no. 3 (2010): 115–43.

Lewis, David. 'Scorekeeping in a Language Game'. *Journal of Philosophical Logic* 8, no. 1 (1979): 339–59. https://doi.org/10.1007/BF00258436.

Litland, Jon Erling, and Juhani Yli-Vakkuri. 'Vagueness & Modality—an Ecumenical Approach'. *Philosophical Perspectives* 30 (2016): 229–69.

McGee, Vann. 'A Counterexample to Modus Ponens'. *The Journal of Philosophy* 82, no. 9 (September 1985): 462. https://doi.org/10.2307/2026276.

McGee, Vann, and Brian McLaughlin. 'Distinctions Without a Difference'. *Southern Journal of Philosophy* 33, no. S1 (1995): 203–51. https://doi.org/10.1111/j.2041-6962.1995.tb00771.x.

Stanley, Jason. 'On the Linguistic Basis for Contextualism'. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 119, no. 1/2 (2004): 119–46.

Williamson, Timothy. 'Reply to Yli-Vakkuri'. *Canadian Journal of Philosophy* 46, no. 4–5 (2016): 839–51. https://doi.org/10.1080/00455091.2016.1205850.

———. *Suppose and Tell: The Semantics and Heuristics of Conditionals*. Oxford: Oxford University Press, 2020.
———. *Vagueness*. New York: Routledge, 2002.
Yli-Vakkuri, Juhani. 'Epistemicism and Modality'. *Canadian Journal of Philosophy* 46, no. 4–5 (2016): 803–35. https://doi.org/10.1080/00455091.2016.1201878.

Faculty of Humanities, Philosophy
Northeastern University London.

---

[1] Fine (1975) and Keefe (2000) provide canonical defences of a more traditional form of supervaluationism. McGee and McLaughlin (1995) defend another kind of supervaluationism. Williamson (1994) provides a canonical defence of epistemicism. For other classical theories of vagueness see Edgington (1997), Barnett (2009), and Bacon (2018) who all appeal to something like penumbral connections.

[2] Admittedly, different classical theories of vagueness will have slightly different conceptions of penumbral connections. I gloss over those differences here.

[3] I mainly focus on indicative conditionals, but the only assumption I make about the behaviour of conditionals is that they entail material conditionals. This assumption is equally plausible for counterfactuals as it is for indicatives.

[4] For a recent exception, see Incurvati and Schlöder (2022).

[5] Some theorists have suggested that MCE fails (see McGee, 1985; Kolomody and MacFarlane, 2010). However, the proposed counterexamples to MCE have always involved conditionals with modalized consequents. Our puzzle affects conditionals with and without modalized consequents. Thus, rejecting MCE for conditionals with modalised consequents will be of limited help here.

[6] Suppose that A and C are borderline and backwardly distributive. It follows immediately that $A \rightarrow C$ is borderline. Because $A > C$ implies $A \rightarrow C$ by MCE, it follows from Normality that $A > C$ cannot be determinate and $A \rightarrow C$ not determinate. So given $A \rightarrow C$ is non-determinate, $A > C$ must be non-determinate as well.

7 Compare Fred's claim here with the claim that John is clearly bald, but he might be borderline bald. This claim would be clearly infelicitous, unlike Fred's claim.

[8] Likewise, after asserting "Barry might be bald" in a context where we have no idea whether Barry is bald or not, we would not feel under any pressure to retract our assertion when later finding out that Barry was borderline bald.

[9] Because contraposition is often held to fail in conditional logics, I do not presume that converse and contrary conditionals are equivalent.

[10] Suppose the antecedent of Fact 2. Given Normality and MCE, we know that determinacy in $A > C$ and $\neg A > \neg C$ implies determinacy in $A \leftrightarrow C$ (the material conditional is contraposable, even if ">" is not). But given A and C are backwardly distributive, $A \leftrightarrow C$ can only be determinately true if neither A nor C is borderline. To see this, suppose A is borderline (we can make a parallel argument for C). Either C is borderline or it is not. Either way we derive a contradiction. If C is borderline, then we know from the fact that A and C are backwardly distributive that $A \rightarrow C$ must be borderline as well. But if $A \rightarrow C$ is borderline then $A \leftrightarrow C$ cannot be determinately true (given Normality). So C must not be borderline. But for the same reason that a biconditional $A \leftrightarrow C$ can't be necessary in normal modal logic whilst one part is contingent and the other is not, $A \leftrightarrow C$ cannot be determinately true whilst one part is borderline and the other is not. So A cannot be borderline.

[11] I am very grateful to Timothy Williamson for suggesting (without endorsing) accommodation as a possible mechanism to explain the linguistic data presented here.

[12] Note that the principle that it is impermissible to assert a borderline case is unviolated under both scenarios.

[13] This is an instance of the general rule of accommodation given by Lewis (1979, p.347).

[14] As with other forms of accommodation, there must be some significant limits on when it is allowed to occur. To see this suppose I say, "if the traffic is bad tomorrow, the sea will be choppy". Given there is no causal or informational link between the traffic being bad and the sea being choppy we wouldn't want penumbral accommodation to ensure that my assertion is automatically determinately true when it turns out that the traffic is borderline bad and the sea is borderline choppy. Outlining exactly what these limits are will be a delicate task, and is not one that I will attempt here.

[15] I use the "interpretation" of a linguistic item here to refer to whatever plays the item's role in determining the truth-value of a sentence of which the expression is part.

[16] A similar defence may be available to forms of supervaluationism where vagueness is interpreted as semantic undetermination since, for these forms of supervaluationists, meaning is a precisification-relative notion.

[17] By analogy suppose that I have consecutively pulled three cards at random out of a full pack. I have hidden their faces from you. Suppose I inform you that it's not the case that all the following are true: the first card was a jack, the second was a four, and the third was black. Although you have now ruled out certain combinations of cards, from your perspective, each individual card drawn by me could still be any card in the pack.

[18] On the epistemicist picture, the truth-value of a sentence is not determined purely by the semantic value of its parts (see Litland and Yli-Vakurri, 2016; Yli-Vakkuri, 2016 and Williamson, 2016 for discussion).

[19] Or rather, two contrary or converse conditionals with *A* in the antecedent or consequent which would be determinately true even when both antecedents and consequents were borderline. Note that there is no requirement that these conditionals be something that we would actually say in a natural conversation. As long as the conditionals meet these conditions, they could be as obscure and complex as one likes.

[20] I am extremely grateful to Julien Dutant, Eliot Michaelson, Robbie Williams, Timothy Williamson, Brian Ball, Daniel Rothschild, and an anonymous reviewer for *Pacific Philosophical Quarterly* for providing very helpful feedback on some version of this material.