

Testimonial Knowledge and Context-Sensitivity: a new diagnosis of the threat

(Alex Davies, alex.stewart.davies@gmail.com)

Forthcoming in *Acta Analytica*

Abstract: Epistemologists typically assume that the acquisition of knowledge from testimony is not threatened at the stage at which audiences interpret what proposition a speaker has asserted. Attention is instead typically paid to the epistemic status of a belief formed on the basis of testimony that it is assumed has the same content as the speaker's assertion. Andrew Peet has pioneered an account of how linguistic context sensitivity can threaten the assumption. His account locates the threat in contexts in which an audience's evidence under-determines which proposition a speaker is asserting. I argue that Peet's epistemic uncertainty account of the threat is mistaken and I propose an alternative. The alternative locates the threat in contexts that provide factors that give audiences a mistaken psychological certainty or confidence that a speaker has asserted a proposition she has not.

1 Epistemic uncertainty vs. misplaced psychological certainty

The fact that a single sentence can be used to assert different propositions in different contexts has the potential to undermine the acquisition of knowledge from testimony in ways to which epistemologists of testimony seldom attend.¹ Even if the speaker is entirely trustworthy and competent on a subject, and her audience knows this, if the audience

¹ On behalf of all authors, the corresponding author states that there is no conflict of interest.

cannot reliably and accurately identify what the speaker is asserting on this subject, surely that will damage the epistemic status of the belief the audience forms on the basis of the speaker's assertion. But if it does so, how exactly does it do so? Andrew Peet (2016) provides an account of how context-sensitivity can damage the epistemic status of beliefs formed on the basis of testimony. According to Peet's account, context-sensitivity poses a problem because it generates epistemic uncertainty: i.e. the evidence available to the audience doesn't allow the audience to identify precisely which proposition has been asserted by the speaker.

In this paper I argue against Peet's epistemic-uncertainty account of the threat posed by context-sensitivity to testimonial knowledge. I will argue that when the evidence underdetermines what the speaker is asserting, there are nonetheless methods for acquiring knowledge on the basis of the assertion and there's no good reason to doubt that audiences generally adopt (some amongst) these methods. Instead, if we're looking for ways in which context-sensitivity might pose a significant threat to the acquisition of knowledge from testimony, then we should be looking at contexts, not in which there's an absence of evidence, but rather in which there's a presence of factors that actively lead the audience to develop a misplaced psychological certainty that the speaker is asserting a proposition that she is not asserting; e.g. misleading evidence, anxiety, bias or prejudice.

The paper will unfold in the following stages. In section 2, I lay out the contexts which (according to the epistemic uncertainty account) are particularly likely to thwart the acquisition of knowledge from testimony. I call these "recovery problem contexts." In

section 3, I present the epistemic uncertainty account of why we are to believe that these contexts have this feature. I'll argue that these contexts thwart the acquisition of knowledge from testimony only if, despite the availability of methods of belief formation which result in safe belief in these contexts, hearers nonetheless opt to use bad methods of belief formation in these contexts. In section 4, I present reason to grant that hearers typically adopt a particular, good method of belief formation (which I call the e-method) when forming beliefs on the basis of testimony. In section 5, I argue that Peet's attempt to support the contrary view involves a sleight of hand and so does not succeed. In section 6, I respond to an attempt to defend the epistemic uncertainty account by restricting the set of contexts in which it is supposed to be operative to a proper subset of recovery problem contexts. In section 7, I respond to an attempt to defend the epistemic uncertainty account that focuses upon the behaviour of the speaker (rather than the hearer) in recovery problem contexts. In section 8, I present the aforementioned alternative to the epistemic uncertainty account of how it is that context-sensitivity can threaten the acquisition of knowledge from testimony.

2 Recovery Problem Contexts

Let's begin then by describing the communicative contexts that Peet proposes are particularly likely to thwart the acquisition of knowledge from testimony. They have three features.

First, in these contexts appears a special kind of linguistic context-sensitive expression: namely, a complex supplementive (cf. (Peet, 2016, p. 399)). A complex supplementive is a

context-sensitive expression: so its content changes with the context in which it is used. It is complex in that its content is not a referent. It is a supplementive in that there is no obvious rule which allows one to establish its content on the basis of independently specifiable features of the context. Quantifier phrases, for instance, such as “every beer” are complex supplementives. They don't have referents as their contents. By standard accounts, they are second-order predicates. They are context-sensitive: e.g. the range of beers “every beer” quantifies over changes with the context in which it is used. And there is no obvious rule that describes how their contents change with independently specifiable features of the context of their use.

Second, the contexts of interest are those in which the “recovery problem” arises (cf. (Peet, 2016, pp. 400–401)). For the recovery problem to arise is for the evidence available to the audience, in the context, to under-determine which proposition the speaker has asserted.² So, given her evidence, there will be more than one proposition which the speaker could be asserting, for all the audience can tell. Let's call these propositions—those that for all the audience can tell, the speaker has asserted—the audience's “epistemic candidates”, in the context.

Finally, the contexts of interest are those in which 'the proposition the audience comes to believe differs slightly from the proposition the speaker intends.' (Peet, 2016, p. 405). There

² Part of the reason why this happens—according to Peet—is because the speaker is negligent in thinking through how she is likely to be (mis)understood. Even if true, this is irrelevant to most of what follows. But I will return to this issue in section 7.

is a mismatch between what the speaker expresses (assuming she expresses what she intends) and what the audience believes.

So the contexts of interest have three features: testimony is expressed using a complex supplementive, the recovery problem arises and the audience believes a proposition which is distinct from the proposition asserted by the speaker. For ease of reference, I will refer to contexts which have all three features as “recovery problem contexts.”

3 Peet's Account: epistemic uncertainty

Peet reasons as follows³:

in recovery problem cases in which the proposition the audience comes to believe differs slightly from the proposition the speaker intends, there will often be nearby worlds in which the audience forms the same belief but in which the belief is false.

(Peet, 2016, p. 405)

Peet employs Sosa's (1999) definition of safety:

A belief by S that p is 'safe' iff: S would believe that p only if it were so that p. (Sosa, 1999, p. 142)

³Peet also argues that beliefs formed in such contexts will be lucky and insensitive. But for reasons of space I focus on safety. Nonetheless, what will be argued about safety applies equally to luck and sensitivity.

Given this definition of safety, Peet's reasoning amounts to the following: if the audience forms a belief in a recovery problem context, on the basis of the speaker's testimony, then the audience's belief is probably unsafe. That's why he concludes:

we are often at risk of forming false or unsafe testimonial beliefs as a result of our heavy reliance upon limited contextual knowledge when planning and interpreting assertions. (Peet, 2016, p. 414)

Given Sosa's definition of safety, Peet's reasoning is truth-conducive. Sosa's definition includes no explicit relativization to the method by means of which the audience forms her belief. So when assessing a belief formed on the basis of testimony for safety, we are permitted to consider nearby worlds in which the audience forms a belief by a method that differs from the actual method employed. Surely there'll be many nearby possible worlds in which the audience forms the same belief through some method or other (not necessarily the same as the actual one employed), such that the belief is had but not true. For example, even if the audience didn't actually select one of her epistemic candidates at random, there'll be nearby worlds in which for some reason or other she does so. In many such worlds, she'll believe that p when p is false. So her actual belief will be unsafe by Sosa's definition.

But for good reason, definitions of safety tend to be relativized to something like the method of belief formation (cf. (Pritchard, 2005, p. 128), (Sosa, 1999, p. 149) and (Williamson, 2009, pp. 364–365)). Whether or not a belief could easily have been false is not a feature of the

belief in itself, but instead a feature of the belief, given the method by, or conditions under, which it was formed. So a better definition of safety would be:

A belief by S that p formed by method M is 'safe' iff: S would believe that p by method M only if it were so that p.

But if Peet's reasoning is understood in terms of this definition (there is one place where Peet (2016, p. 405) explicitly relativizes safety to methods), then it is not truth-conducive. The truth of Peet's conclusion depends upon which method audiences tend to employ to form beliefs in recovery problem contexts. There are some methods of belief formation which can be deployed in response to recovery problem contexts which generate safe belief. Here is one such method. A proposition which is implied by each of the audience's epistemic candidates is a proposition which is true if the speaker's asserted proposition is true. One method of belief formation that an audience could deploy in response to a recovery problem context is to believe a proposition, p, for which she has reason to believe it is implied by the proposition asserted by the speaker. Assuming that this reason for belief is not misleading (as we will suppose until section 7), p will be implied by each of the audience's epistemic candidates: since for any proposition that doesn't imply p, it couldn't be the proposition the speaker asserts and so couldn't be one of the audience's epistemic candidates.⁴ This method will lead to the formation of safe beliefs, insofar as the speaker is a good testifier (whatever is required for that) and the audience is suitably sensitive to this: for this method will lead

⁴ Note that it is not a part of the e-method that one does any of this consciously. They might, though for reasons canvassed in section 6, this seems unlikely.

to the formation of beliefs that are true if the speaker's asserted proposition is true. This method is also available even in recovery problem contexts: for if the audience's evidence doesn't suffice to identify exactly which proposition the speaker has asserted, that doesn't mean it doesn't suffice to identify features of that proposition e.g. implications of that proposition. For ease of reference, I will call this method of belief formation “the e-method.”

The availability of methods like the e-method means that Peet cannot be—and should not be understood to be—neutral about which methods the audience employs in recovery problem contexts. If he wants to conclude that such contexts pose a significant threat to the acquisition of safe belief from testimony, he must be supposing that audiences will commonly adopt methods of belief formation which generate unsafe beliefs in those contexts. Peet's proposal then, with this commitment made explicit, is that when met with an under-determination of the speaker's asserted proposition by the available evidence, audiences behave epistemically irresponsibly and adopt a method of belief formation that tends to generate unsafe belief.

4 We seem to use the e-method

There is *prima facie* reason to believe that audiences commonly adopt the e-method when forming beliefs on the basis of testimony. There is an uncanny correlation between the evidence the audience has about the content of an assertion and what she'll be inclined to believe on the basis of her evidence. Consider the sentence, “there isn't any food.” This sentence could be used to express a variety of different propositions. Amongst them are:

- (1) There isn't any food that a baby would safely and willingly and capably eat.
- (2) There isn't food that the speaker (in case 2 below) thinks her dog would safely and willingly and capably eat.
- (3) There isn't any food that is no longer edible but which indicates that someone was living in the Plaza until recently.
- (4) There isn't any food that this particular homeless person (case 3 below) might think fine to eat, even though persons who only eat food they haven't found in a dumpster would not.
- (5) There isn't any dog food.

Here are three cases. In each case you are given information about the context of utterance of the sentence "there isn't any food."

Case 1

You find a recording of someone saying the sentence, "there isn't any food." But you have no information about where or when or in what circumstances.

Case 2

You witness a woman clearly talking to her dog, who is clearly begging for something, and looks unhappy. She says to the dog, in a sorry tone, "there isn't any food (I'm afraid)."

Case 3

You witness one homeless woman say, "Damn it: there isn't any food," to another, as they are fishing through a dumpster.

The different information you have across the cases gives you different kinds and amounts of evidence about the content of "there isn't any food" as used in its context of use. In case 1, given your complete ignorance of the context of the sentence's utterance, you know very little about the proposition expressed. You know that it says that something doesn't exist, and you know that this something qualifies as food in some sense. But you don't know exactly which range of items is being said not to exist. You certainly are not in a position to say whether the proposition expressed by the sentence in its context entails any of (1)-(5). In case 2, you have more information about the proposition that the sentence expresses when uttered in the corresponding context. But you are not in a position to know everything there might be to know about it. You can be pretty confident that it entails (2) and that it doesn't entail any of (1), (3) or (4). But it's not clear what one should say about (5). In case 3, you can be sure that the proposition expressed by the sentence entails (4). But you would be unsure about whether the proposition entails any of (1), (2), (3) or (5).

Suppose now that in each case, you know the speaker to be trustworthy and competent and so you are inclined to form some belief on the basis of the relevant utterance of "there isn't any food." Which belief would you be inclined to believe in each case? It seems very likely

that there will be a good correlation between the proposition(s) that you have reason to believe to be implied by the proposition asserted in each case, and the proposition(s) that you will be inclined to believe in each case. In case 1, you're not inclined to believe any of (1)-(5). In case 2, you're inclined to believe (2) but not (1) or (3)-(5). In case 3, you're inclined to believe (4) but not (1)-(3) or (5). If so, then it seems we are naturally inclined to believe propositions that we have reason to believe are implied by the proposition asserted, even when we don't have enough information to establish exactly which proposition has been asserted.⁵ This correlation between what we are inclined to believe and what we have reason to believe to be implied by the asserted proposition, is *prima facie* reason to suppose that audiences commonly use the e-method when forming beliefs on the basis of testimony.

5 Peet's reason for denying that we normally use the e-method

Peet explicitly considers the hypothesis that audiences typically adopt the e-method in recovery problem contexts.⁶ He rejects the hypothesis. He provides what is intended as a paradigmatic example, in which a normal audience plausibly forms a belief with a content that is not implied by each of her epistemic candidates. If we grant the paradigmatic status

⁵ Once again, note that this sensitivity to reasons need not be conscious or inferential. Though it might be.

⁶ Peet (2016, p. 408) considers the possibility that audiences adopt the following method of belief formation: only believe propositions that 'are true at all (or almost all) the worlds at which any of the fine grained epistemic candidates are true (such a proposition would be equivalent to the disjunction of the fine grained epistemic candidates).' He calls such propositions "coarse grained." I assume that the word 'any' is a typo (it should be 'each') since if we read this as it stands then each epistemic candidate will qualify as a proposition that is believed when following the described method. But Peet clearly intends the propositions believed when following the method to be distinct from the epistemic candidates. Note, furthermore (and *contra* Peet), that what appears inside the brackets is not equivalent to what appears outside the brackets. Nonetheless, the proposal Peet considers is that audiences believe propositions which are true given that what the speaker asserts is true, even though the audience cannot identify precisely which proposition the speaker has asserted. He's just mistaken in thinking that such propositions will always be equivalent to the disjunction of the epistemic candidates.

of the example, then it's a *prima facie* compelling reason to grant that audiences are not pursuing the e-method. They pursue some other method which doesn't reliably generate safe belief. In which case, Peet's account of the threat posed by context-sensitivity to testimonial knowledge would have some significant support. In this section, I'll describe the paradigmatic example and then argue that it is not an example of what it needs to be in order to support the contention that audiences do not typically use the e-method (at least in recovery problem contexts).

Here's the paradigmatic example:

Paradigmatic Example

Matt and Sally are at Matt's house. Sally says 'I am hungry, is there any food?', to which Matt responds 'Sorry, there isn't any food, lets order a pizza'. Sally agrees, and they order a pizza. Clearly Matt was not saying that there isn't any food *anywhere*. He was saying that there is no food in some restricted domain. However, there are many similar ways the domain could be restricted, here are a few:

1. There isn't any food *belonging to Matt*.
2. There isn't any food *belonging to Matt or Tom (Matt's house mate)*.
3. There isn't any food *that Matt is willing to share*.
4. There isn't any food *which Sally likes and which meets the above criteria*.
5. etc. (Peet, 2016, p. 401)

This much of the example provides us with the testimony viz. the expression of a proposition by uttering the sentence “there isn't any food.” It also provides us with some of Sally's epistemic candidates in this context (viz. (1)-(4)). Peet then goes on to claim that, quite plausibly, Sally would believe number (2) in response to Matt's testimony. The crucial part is as follows:

Many of our beliefs are revealed through the dispositions we acquire upon forming them. If Tom's open food policy were salient to Sally at the time of Matt's assertion then she would, in many cases, acquire the disposition to be surprised upon finding her favourite foodstuff in the fridge labelled 'Property of Tom'. (Peet, 2016, p. 409)

If Sally's disposition reflects her belief, then we could reasonably conclude that she believes that there isn't any food belonging to Matt or Tom viz. (2). But (2) is not implied by (1) or by (3). Hence it is not implied by each of Sally's epistemic candidates. Hence it's not been formed by the e-method. Given the absence of relevant evidence in the example, Sally must have arrived at this belief in a way that is not properly guided by her evidence. If the example is paradigmatic, then we should grant that it's quite normal for audiences to form unsafe beliefs in recovery problem contexts.

I will now argue that the example is not what it seems. In particular, I will argue that it is not at all obvious that the proposition which Sally believes is not implied by each of her

epistemic candidates. In order to establish whether the proposition that Sally believes is implied by each of her epistemic candidates, we need to establish what evidence she has about what propositions are implied by Matt's asserted proposition. So let's register that evidence.

Notice that Matt is responding to a question by Sally: 'is there any food?' Sally's question is what conversation analysts call a *pre-request* (Levinson, 1983, pp. 245–264). A pre-request is a sub-category of pre-turn. A pre-turn in a conversation checks that preconditions on some possible future conversational action are in place. It invites a response that confirms that they are in place or that they are not. If they are not, then the action itself is not performed. If they are, and if the action is not pre-empted in the response to the pre-turn, then the action is then performed. Besides pre-requests, there are also: pre-invitations, pre-proposals, pre-announcements and others. There are several indications in Paradigmatic Example that Sally's question is a preliminary to a request that will come later: and more specifically, a request for something that she could eat. We can see this by comparing Paradigmatic Example with some modified contexts.

Firstly, compare the utterance of the question with the following:

Context 2

Matt and Sally are surveyors of a shipwreck. They are constructing an inventory of the items that are inside the ship, using a deep-sea robot. Sally goes through the list

of items on her list and Matt responds:

Sally: Is there any radioactive material?

Matt: No.

Sally: Are there any animals?

Matt: No.

Sally: Is there any food?

Matt: Sorry, there isn't any food. Let's order pizza.

In the last line, Matt apologizes and makes a proposal just as he does in Paradigmatic Example. But whereas in Paradigmatic Example this seems quite natural, and is left unchallenged by Sally, if the apology and proposal were given in Context 2, it is both unnatural and likely to be challenged or treated as a joke by Sally. Why? Because it's apparent in Paradigmatic Example that Sally is about to request something and Matt is rejecting that request and it's that rejection for which he is apologizing and it's because he's rejecting it, that he's making his alternative proposal. But in context 2, Sally is not making a request at all. Matt is treating Sally's inventory question as if it were a pre-request, when its context suggests it isn't. So the fact that Matt apologizes and makes his alternative proposal in Paradigmatic Example, and Sally doesn't treat this as a joke or challenge it or correct it, is a strong indication that Sally's question, in that context, is to be taken as a pre-request.

Secondly, not only is there reason to believe that Sally's question is a pre-request, there's

also reason to believe that it's a pre-request for something for Sally to eat. Compare Paradigmatic Example with context 3:

Context 3

Matt and Sally are at Matt's house. Sally says 'my cat's hungry, is there any food?', to which Matt responds 'Sorry, there isn't any food, lets order a pizza'.

In this context, the question seems most naturally read as a pre-request for food for Sally's cat to eat, and not for her to eat: Matt even seems to come off as misunderstanding Sally's question when he suggests ordering in pizza — not the most sure-fire way to ensure a happy cat. This suggests that, in Paradigmatic Example, Sally's pre-request is not just any old pre-request. It is a preliminary to a request for food for *her* to *eat* and not for some other purpose.

If Matt's utterance of 'there isn't any food' is indeed a response to Sally's particular pre-request (and more specifically, if it's a denial that a pre-condition for a subsequent request for some food for Sally to eat is satisfied), that gives us (including Sally) evidence about the content of the proposition that Matt asserts by uttering this sentence. It gives us evidence that the sentence is being used to deny the existence of food that Sally could eat. It's not being used to deny the existence of food for the purposes of producing an inventory of what's in Matt's house (no matter how rotten, how soaked in water it might be etc.). It's not being used to deny the existence of food that would satisfy some other request (e.g. for Sally's cat to eat). What matters in this context is whether there's anything that Sally could

eat. We should surely acknowledge that the context doesn't enable us to establish, for *any* food, whether it counts as food that Sally could eat. But despite that, there are *some* paradigmatic items that should not exist if the proposition Matt asserts in responding to Sally's particular pre-request is true. In particular, food 1 should not exist if Matt's asserted proposition is true:

Food 1

- It is available for Sally to eat without paying for it or stealing it from a stranger.
- Matt owns the food.
- It is not dangerous to eat—it's not gone bad, etc.
- It is stored in a fairly standard and nearby way—e.g. in tupperware in the fridge, and not in a plastic bag in the trash can.
- It falls squarely within the cultural norms of the interlocutors (e.g. it's not any kind of insect).
- It is not disliked by Sally.

But then it seems that the information we have about the speech act which Matt is performing—a lot of which we garner from the position of his utterance within a broader conversation—serves to provide us with some information about the content of the quantifier phrase 'any food' when used to perform that speech act in this context.

The constraints that location within a conversational sequence can place upon content has

not gone unnoticed. Paradigmatic Example is presumably fictional. But it bears an uncanny resemblance to a non-fictional example of pre-sequences and their effects on the content of a quantifier phrase provided by Stephen Levinson.

A: Whatcha doin'?

B: Nothin'

A: Wanna drink?

(Levinson, 1983, p. 346)

Here there is a pre-invitation ('Whatcha doin?'), a response to that pre-invitation using a quantifier phrase ('Nothing') and then the offer itself, 'Wanna drink?' Levinson notes the following about B's utterance of the quantifier 'nothing':

Nothing in [the example] can be read as 'nothing that would make the offer of an evening's entertainment irrelevant' or the like... (*ibid*)

The fact that 'Nothin' is a response to a particular pre-request shapes what it is reasonable to take the quantifier to be quantifying over. Should B's scratching his foot count as doing something? In this context, obviously not (though in another context, it could). Should B's having made other arrangements to visit the cinema, exclusively with a friend, count as something? In this context, obviously yes (though in another context, it would not). Given that the quantifier is used in answering a particular pre-invitation, we learn some things

about how best to understand what the quantifier quantifies over in this context. In Peet's example, we see the same interaction between the role played by the quantifier phrase in a conversational sequence and the most reasonable assignment of content to that quantifier phrase.

In light of this, we can now see that given what information Sally has available to her in Paradigmatic Example, Sally has good reason to believe that the following proposition is implied by the proposition Matt asserts—whatever it is: *there isn't any food 1*. This proposition will be implied by each of Sally's epistemic candidates. But this much doesn't show that (2) is implied by each of Sally's epistemic candidates—for this proposition doesn't imply (2).

There is one additional piece of information that Peet provides us about the context and which he makes sure Sally is aware of: Tom shares his food. This is salient to Sally. So I assume this means she is aware of it. Let's say that food 2 is as follows (i.e. just like food 1, except Tom owns it instead of Matt):

Food 2

- It is available for Sally to eat without paying for it or stealing it from a stranger.
- Tom owns the food.
- It is not dangerous to eat—it's not gone bad, etc.
- It is stored in a fairly standard and nearby way—e.g. In tupperware in the fridge, and

not in a plastic bag in the trash can.

- It falls squarely within the cultural norms of the interlocutors (e.g. it's not any kind of insect).
- It is not disliked by Sally.

Given all the information we're told Sally has, she has good reason to believe that the following proposition is implied by Matt's asserted proposition (whatever it is): *there isn't any food 2*.

We can conclude from this examination of Sally's evidence that Sally has reason to believe that the following proposition is implied by the proposition that Matt asserts: *there isn't any food 1 and there isn't any food 2*. If so, then this proposition must be implied by each of her epistemic candidates. Moreover, this proposition could easily be expressed with the sentence that Peet uses to describe Sally's belief, viz. "There isn't any food belonging to Matt or Tom." All we need is for the quantifier phrase "any food" to be contextually restricted to admit only items that satisfy either the conditions of food 1 or the conditions of food 2. But then the example is not a clear case in which Sally forms a belief in a proposition that is not implied by each of her epistemic candidates.

One might try to resist this conclusion in the following way. Surely Matt could have been asserting only that Tom *generally* shares his food, but not that he shares all of his food whatever the circumstances. For example, he might have a birthday cake for a birthday the

following day which he's not willing to share. But Sally's belief is not entailed by this proposition. So surely we *do* have an example in which Sally would easily form a belief in a proposition that she has no reason to think is entailed by Matt's asserted proposition i.e. that is implied by each of her epistemic candidates. However, the question to ask is: when faced with such a birthday cake, and when informed of its status, would Sally consider what Matt said to have been false? Answer: No. The implication? What Sally believes, in believing that there isn't any food 2, doesn't require that *such* food be available for her to eat (i.e. satisfy the first condition for being food 2). Obviously when she knows that Tom is willing to share his food, she wouldn't reasonably believe that Tom will share it no matter the circumstances. Sally's belief will be nuanced in this respect (as one would expect of a belief expressed using a verb with habitual or generic aspect). Here's the moral: whenever you think you have found circumstances in which Matt's proposition could be false while Sally's belief is true, think through whether Sally would consider such circumstances to show Matt's assertion false. If not, then she doesn't believe the proposition you're supposing she does.

It seems that Peet has stated what he thinks Sally's epistemic candidates are before he finishes specifying what information Sally has about the proposition asserted by Matt. However, changes to the information Sally has about Matt's asserted proposition will change what her epistemic candidates are in the context, and thus, what is common to all of them. The example, therefore, generates the firm appearance that Sally would believe a proposition that is not permitted by the e-method only with the help of a sleight of hand.

The problem here does not arise just because Peet has chosen the wrong example. The problem with Peet's particular example is likely to arise for any attempt to provide a compelling example in which the audience clearly believes a proposition that is not implied by each of her epistemic candidates. To provide a compelling example in which an audience believes a given proposition *p*, we should give the audience reason to believe that *p* is true. But if *p* is the content of a belief formed on the basis of testimony, then the reason to believe that *p* is true which the audience will have, will be a reason to believe that the speaker has asserted something that entails (or at least, strongly implies) *p*. This is indeed what we see Peet doing, when he talks about what is salient to Sally when trying to convince us that Sally would easily believe that there isn't any food belonging to Matt or Tom. But if the audience has such a reason, then the audience obviously does have reason to believe that *p* is entailed by the speaker's asserted content. But then the proposition will—contrary to the point of the example—be implied by each of the audience's epistemic candidates.

6 Forming beliefs fast and without conscious consideration of the evidence

Although recovery problem contexts don't imply that audiences will use a bad method when forming beliefs on the basis of testimony in such contexts, perhaps there is a subset of such contexts in which a bad method will be employed. Peet (*ibid*, p. 400, p. 404, p. 410, p. 412, p. 414) distinguishes between casual and non-casual contexts. A casual context is one in which audiences 'interpret extremely quickly, and without much conscious consideration of the evidence' (*ibid*, p. 400). Examples of such contexts include: 'when we socialise at the pub, relax at home with our partners, or engage in passing small talk in the department.'

(*ibid*, p. 406) Perhaps then, although there's no reason to grant that audiences generally use a bad method of belief formation in recovery problem contexts, we're on safe ground in supposing that audiences typically use a bad method of belief formation in casual recovery problem contexts.

But it is doubtful that in casual recovery problem contexts, audiences are less likely to be using the e-method (or some other method that generates safe beliefs). Firstly, there's nothing about the e-method which, by definition, requires slow conscious consideration of the evidence. The method requires only that audiences form beliefs in ways that are responsive to the relevant evidence. Secondly, and more generally, it isn't true that in order to form beliefs in a way that is responsive to one's evidence, one must consciously think through the evidence and form one's belief slowly. There are many belief forming processes which are fast, involve little or no conscious consideration of the evidence, but which generate beliefs that are reliably sensitive to the available evidence. Consider an acquired perceptual expertise such as the ability to recognize finger prints, to read radiographs, fMRI scans, or to recognize the species of a bird (cf. references in (Stokes, unpublished)). Beliefs formed by persons who have acquired such expertise are more likely to be true than those formed by non-experts, but are formed much faster and with less conscious consideration of the evidence than non-experts are likely to consider. The possibility that speech comprehension is a form of perceptual expertise is not at all implausible. Until we have reason to believe that seemingly competent communicators in a language lack such expertise, the fact that someone has formed a belief fast and without conscious consideration

of the evidence is not a reason to think that their beliefs, so formed, don't track the evidence (and are therefore unsafe).

Casual contexts in which beliefs are formed quickly and without conscious consideration of the evidence, then, are not contexts in which audiences are, on account of that, any less likely to employ a method that generally generates safe belief (e.g. the e-method).

7 Do recovery problem contexts make speakers incompetent?

We have been assuming that the proposition the speaker asserts and intends to assert is within the set of the hearer's epistemic candidates. Peet himself seems to take this for granted. When he discusses the e-method (cf. footnote 6 above), he doesn't criticize it on the ground that a belief formed with this method would still not be implied by the speaker's asserted proposition (which it wouldn't be if the epistemic candidates didn't include the speaker's asserted proposition), but instead, as we have seen, on the ground that it's empirically implausible that hearers use this method. However, at this stage, a defender of Peet's account might be tempted to ditch the assumption and adopt the hypothesis that in recovery problem contexts, a speaker will tend to create the misleading impression that she intends to assert a proposition which is not implied by the proposition she intends to assert. A version of Paradigmatic Example is a case in point. At one stage, Peet (2016, p. 404) asks us to suppose that, when it's apparent to Sally that Tom shares his food (and so it is reasonable for her to interpret Matt as asserting a proposition that implies that there's no food owned by Tom), Matt nonetheless intends to assert that only Matt has no food.

Assuming that Matt only has support for the proposition he intends to assert, it follows that Sally's belief is still (probably) going to be unsafe, even if she uses the e-method in this context: for the proposition she believes probably won't be supported by the basis upon which Matt makes his assertion. One might take this to support Peet's contention that beliefs acquired from testimony in recovery problem contexts are very likely to be unsafe.

However, for this to be so, the under-determination of the speaker's assertion by the hearer's evidence must make it likely that the speaker will create the misleading impression that she intends to assert a proposition that is not implied by the proposition she actually intends to assert. For of course mistakes by hearers in interpretation, and mistakes by speakers in their choice of words, are made—we don't need a paper to be convinced of that. The interest of Peet's paper, however, is in the prospect of identifying a kind of context in which these mistakes are particularly likely and so in which the acquisition of knowledge from testimony is particularly unlikely. Peet thinks that the under-determination of the speaker's asserted proposition by the context that is discernible to the hearer makes this likely. We've already seen that this is not true *vis-a-vis* the hearer's interpretation of the speaker's assertion. But it is equally not true *vis-a-vis* the speaker's choice of words. The fact that the precise proposition she intends to assert is not discernible by the hearer doesn't in any way imply that the speaker will tend to create the impression that she intends to assert a proposition that is not implied by the proposition she intends to assert. It's quite consistent for us to suppose that in all contexts wherein there is this kind of under-determination, the speaker manages to express herself clearly enough that the proposition she intends to assert falls

within the hearer's epistemic candidates. If a defender of Peet wants to say otherwise, then we need some additional argument.

One such argument parallels the line reviewed in the previous section. For Peet (2016, p. 403) stresses that speakers—just like hearers—don't tend to think too much about the reasonable interpretations of their utterances that are available in the context of utterance as they speak. So perhaps we should think that speakers will be likely to generate, not just under-determining evidence, but positively misleading evidence about the propositions they intend to assert in recovery problem contexts because they don't think too much about this sort of thing. However, the problems with this proposal parallel those with the proposal discussed in the previous section. To become adept at speaking, is to be able to do it without having to think, slowly and consciously, about the different ways one might be being interpreted. As with other learned skills, even if a speaker doesn't tend think too much about how she could be misinterpreted as she speaks, that doesn't imply that she speaks badly i.e. tends to provide positively misleading evidence about the proposition she intends to assert.

8 An alternative account: misplaced psychological certainty

One way that context-sensitivity might threaten the acquisition of knowledge from testimony is that it gives rise to an under-determination of the speaker's asserted proposition by the evidence available to the audience. What we can now see is that such epistemic under-determination threatens the acquisition of knowledge only if, when in contexts that give rise to such under-determination, speakers are typically incompetent at

making sure that their intended contents fall within the audience's epistemic candidates, or alternatively, audiences typically adopt a bad method of belief formation in response to evidence that under-determines which proposition a speaker has asserted. But as we have also seen, there is nothing about such recovery problem contexts which will typically cause speakers or audiences to behave in these ways.

What this suggests is that a focus on epistemic under-determination is slightly to the side of where we should be focusing. If we're looking to explain how context-sensitivity might threaten knowledge by testimony, we would do better to focus on contexts in which there are factors that actively mislead even generally competent audiences about which proposition has been asserted: i.e. contexts in which even if the audience uses a method like the e-method, she is either systematically caused to misapply the method, or alternatively, the method itself (even when applied properly) is systematically corrupted. Although epistemic under-determination isn't itself such a factor, contexts in which such factors are likely are contexts in which a testimonial belief is likely to be unsafe. I will now describe two kinds of such context.

Firstly, there are contexts in which the speaker has good reason to provide the audience with misleading evidence about the content that she asserts. This will happen if the speaker has multiple audiences and a strong reason to ensure that a given content is perspicuous only to one of those audiences. For example:

John's parents have been preparing him for boarding school. But they are hiding this departure from John's sister—who will be very upset when she finds out. John is upstairs getting ready for soccer practice. John's sister sits between John's parents, by the door as they wait for John. John's mother asks John's father, in a sullen tone, inquiring into whether John is now in a position to make the entrance exam for the boarding school, "Is John ready?" John's father, recognizing that this is what his wife is asking, replies, "yes, John's ready." John's sister has no idea that John is being prepared for boarding school. So she thinks—as is reasonable—that the remark concerns John's being ready for soccer practice.

The nature of the assertion "John's ready" is deliberately concealed from John's sister. John's sister's resulting belief is unsafe: even when formed using the e-method, in many nearby possible worlds (so in which the context of belief formation includes the same or similar intentionally provided misleading evidence) in which she forms the same belief in the same way, her belief will be false.⁷ Generally, if one is in a context where the speaker has good reason to mislead one about what she's saying, one will probably form one's belief on the basis of misleading evidence rendering one's belief unsafe.⁸

Secondly, there are contexts in which, although the speaker isn't producing misleading evidence about what it is she is asserting, there are factors in the context which actively

⁷This kind of context encompasses those in which a politician has reason to use an overt intentional dog-whistle (Saul, forthcoming).

⁸It's possible that there are contexts in which the speaker is likely to provide misleading evidence accidentally. But I'm not sure in which contexts this is likely to happen.

cause a normally competent audience to reliably misinterpret the supportive force of her evidence without her noticing. Here are three examples of such a context.

Firstly, social anxiety leads one to negatively interpret social situations, including others' speech, even when a more sober assessment of the evidence would point to a different interpretation of the evidence (cf. (Beard & Amir, 2009), (Blanchette & Richards, 2010) and (Kingsbury & Coplan, 2016)). She says "of course you're nice, you have to be," and she was saying that you come from a nice and decent family, so you're bound to be nice. But because you suffer from social anxiety, you think it's reasonable, given the available evidence, to interpret her as saying that you have nothing else going for you, so you have no choice but to be nice to people (i.e. a difference in interpretation of the modal base of the modal verb "have"). If you form a belief on the basis of this testimony, it'll have been formed using a method (a method warped by your social anxiety) that in many nearby possible worlds will generate false beliefs. Generally: if you are in a context wherein you suffer from social anxiety, then the beliefs you form on the basis of testimony from others about yourself (even if you're normally quite adept at using the e-method) are likely to be unsafe, because your social anxiety causes you to wrongly assess your evidence about what the speaker has asserted.

Secondly, one might not suffer from social anxiety. But one might nonetheless, for unfair reasons, think rather poorly of somebody, and that might influence how one interprets their behaviour (including their speech). In Russian there is an idiom, "Всякое лыко в строку,"

which translates as *every strip of bark into the line*. It's used to describe the tendency to interpret everything someone does in such a way that that person can be blamed. If one has such an unfair tendency toward a particular speaker, then it's likely to lead one to think one has reason to believe she asserted something that entails that p , when really one has no such reason. Here's an example.

Part of Lucy's job is to distribute information to everyone in the company. Owing to a warped interpretation of his past experience with Lucy, John falsely believes that Lucy tends to distribute information to her friends and not others: a set of people that doesn't coincide with the set of people who ought to receive the information. One day John wants to know whether today's information has been distributed to everyone that the information may be useful to. He asks Lucy, "has today's information been distributed?" Lucy replies, "Yes, the information has been distributed to everyone." Given John's false belief about Lucy, John thinks that Lucy's "everyone" must have been used by Lucy to quantify over everyone in Lucy's preferred circle. So this is what John comes to believe on the basis of her testimony: believing himself to possess evidence that this is entailed by the proposition she asserted.

John's belief will be unsafe. For given the method by which he formed his belief (viz. on the basis of an assessment of his evidence (concerning what Lucy asserted) that is influenced by his unfair dislike of Lucy), there will be nearby possible worlds in which he forms the same

belief in the same way, but in which his belief is false. Generally, if someone is speaking to you for whom you weave *every strip of bark into the line*, then you are likely to form unsafe beliefs on the basis of that speaker's testimony, even if you're ordinarily competent at using the e-method.

Thirdly, the same phenomenon may arise because the speaker is a member of a *group* against which one is prejudiced. In another paper, Peet (forthcoming) presents two examples of this (to show how the speaker may suffer epistemic injustice, and not to show how the audience's beliefs will, as a result, be unsafe). Here is a paraphrase of one of Peet's examples:

Amia is the manager of a restaurant. But her kitchen staff have various false views about a woman's competence to manage a restaurant. Amia is concerned that all front of house staff are women. So she says, "I need a man." Amia was expressing the proposition that she needs a man who is competent at front of house duties, to carry out such duties. But when the kitchen staff hear this, because of their false beliefs about women, they think it most reasonable to interpret Amia's asserted content as entailing that she needs a man to help her run the restaurant, and not a competent man to take up a front of house position.

A member of the kitchen staff forms a belief that is unsafe because their prejudice gives them a misplaced psychological confidence in a particular interpretation of Amia's words. In many nearby possible worlds in which they form the same belief using the same prejudice-

affected method, their belief will be false. Generally, if a speaker is a member of a group against which one is prejudiced, and this influences testimony interpretation, then it is likely that one will form unsafe beliefs on the basis of their testimony.

8 Summary

Linguistic context-sensitivity might well typically give rise to an epistemic under-determination of the speaker's asserted proposition by the audience's evidence. But, I have argued, this in itself is not a threat to the acquisition of knowledge from testimony: there are methods that audiences plausibly use to form beliefs in response to testimony which generate safe belief, even given such under-determination. However, context-sensitivity also opens up testimonial transactions to influence from factors that cause audiences to be confidently mistaken about which proposition a speaker has asserted: either because such factors corrupt methods of belief formation which otherwise typically deliver safe testimony-based beliefs, or because such factors render audiences (in such contexts) locally incapable of properly applying methods of belief formation which otherwise typically deliver safe testimony-based beliefs. Context-sensitivity endangers the acquisition of knowledge from testimony not because context doesn't provide enough to guide the audience in her belief formation (she can easily adapt to that) but rather because context can actively misguide an audience in her belief formation.⁹

⁹ I'd like to thank audiences at the University of Helsinki, Department of Philosophy in Spring 2017 and at the 2017 Joint Session in Edinburgh for very helpful comments and criticisms. I'd like to thank a reviewer for this journal for the same. I'd like to thank Andrew Peet for helping me improve upon an earlier version of this paper. The research that led to this paper was funded by the Estonian Science Agency through grants PUT632 and IUT20-5, and by grant PHVF117925 from the University of Tartu.

References

- Beard, C., & Amir, N. (2009). Interpretation in Social Anxiety: When Meaning Precedes Ambiguity. *Cognitive Therapy and Research*, 33(4), 406–415.
- Blanchette, I., & Richards, A. (2010). The influence of affect on higher level cognition: A review of research on interpretation, judgement, decision making and reasoning. *Cognition and Emotion*, 24(4), 561–595.
- Kingsbury, M., & Coplan, R. (2016). RU mad @ me? Social anxiety and interpretation of ambiguous text messages. *Computers in Human Behavior*, 54, 368–379.
- Levinson, S. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Peet, A. (forthcoming). Epistemic injustice in utterance interpretation. *Synthese*.
- Peet, A. (2016). Testimony and the epistemic uncertainty of interpretation. *Philosophical Studies*, 173(2), 395–416.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Saul, J. (forthcoming). Dogwhistles, Political Manipulation and the Philosophy of Language. In D. Fogal, M. Cross, & D. Harris (Eds.), *New Work on Speech Acts*. Oxford: Oxford University Press.
- Sosa, E. (1999). How to Defeat Opposition to Moore. *Philosophical Perspectives*, 13, 141–153.
- Stokes, D. (unpublished). On Perceptual Expertise.
- Williamson, T. (2009). Reply to Ernest Sosa. In P. Greenough & D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 364–369). Oxford: Oxford University Press.