

Preprint: This paper is forthcoming in Pacific Philosophical Quarterly, please cite published version.

Testimonial Knowledge Without Knowledge of What is Said.

Andrew Peet

Abstract

This paper discusses the following question: what epistemic relation must audiences bear to the contents of assertions in order to gain testimonial knowledge? There is a brief discussion of why this issue is of importance, followed by two counterexamples to the most intuitive answer: that in order for an audience to gain testimonial knowledge that *p* they must know that the speaker has asserted *p*. It is then suggested that the argument generalises and can be made to work on different sets of assumptions about the conditions for knowledge, and the conditions under which a proposition is asserted.

1. Audiences and Asserted Contents.

Standard formulations of the conditions for testimonial knowledge rarely make any mention of the relation audiences must bear to an asserted content (or the fact that it has been asserted) in order to gain testimonial knowledge. Rather, focus is usually on the relation the audience must bear to the speaker (i.e. must audiences have reason to trust speakers?), and the relation the speaker must bear to the asserted content. This is perhaps due to the fact that content recovery is usually taken to be unproblematic by epistemologists. As Sanford Goldberg puts it:

'...so prevalent is the assumption that the comprehension dimension is unproblematic - that hearers reliably recover the propositions attested to - that there is virtually no discussion of the comprehension processes in the epistemological literature' (Goldberg 2007: 54)

However, in the philosophy of language content recovery is not seen as epistemically unproblematic. Theorists concerned with context-sensitivity have long recognised that audiences sometimes struggle to recover the precise proposition intended by the speaker. Anne Bezuidenhout lists the following factors as contributing to audiences' interpretations of context-sensitive utterances. All of these factors leave space for misinterpretation:

- (i) Knowledge that has already been activated from the prior discourse context (if any).
- (ii) Knowledge that is available based on who one's conversational partner is and on what community memberships one shares with that person.
- (iii) Knowledge that is available through observation of the mutual perceptual environment.
- (iv) Any stereotypical knowledge or scripts or frames that are associatively triggered by accessing the semantic potential of any of the expressions currently being used.
- (v) Knowledge of the purposes and abilities of one's conversational partner (e.g. whether the person is being deceitful or sincere, whether the person tends to verbosity or is a person of few words etc.).
- (vi) Knowledge one has of the general principles governing conversational exchanges (perhaps including Grice's conversational maxims, culturally specific norms of politeness, etc.). (Bezuidenhout 2002: 117)

If audiences must rely on all this knowledge in order to correctly recover what is said then it is easy

to worry that recovery might often go wrong. Indeed, the epistemic challenges raised by interpretation of context-sensitive utterances have led some theorists (such as Bezuidenhout 1997, Recanati 2004, Heck 2002, Carston 2002, and Sperber and Wilson 1986) to conclude that speakers and hearers needn't even share the same contents in order for communication to be successful, and others (such as Cappelen and Lepore 2004) to reject the view that there is widespread context-sensitivity in natural language. The thought being that successful communication would be rare if context sensitivity is rife and audiences need to rely on so much contextual knowledge to recover the contents of context sensitive utterances. This controversy illustrates that, despite epistemologists' assumptions to the contrary, it is not clear that comprehension is unproblematic. Thus the question of the audience's relation to what is said is more pressing than has previously been recognised. Moreover, if audiences must bear a demanding relation (such as knowledge) to the objects of assertion then theories of communication which seem to entail that audiences often struggle to achieve knowledge of what is said will face a problem: they will entail skepticism about much of our testimonial knowledge. Thus, the issue is not only of importance to epistemologists.

Perhaps there is a simple answer to this question though: Hearers must know (or be in a position to know) that the speaker has said that p in order to gain testimonial knowledge that p ⁱ. After all, if the audience doesn't even know that the speaker has said that p then why worry about whether or not the speaker is trustworthy or reliable with respect to p ? This seems like a highly intuitive requirement on testimonial knowledge. Moreover, when we are told that p , and we are asked why we believe that p , we will usually state ' S said that p ' as our reason. If the fact that a speaker said that p is part of our reason for believing that p then surely knowledge that the speaker said that p is necessary for testimonial knowledge that p . Anna-Sara Malmgren (2006) employs this thought in her argument against Tyler Burge's claim that there is a priori testimonial knowledge. She claims that knowledge of what is said is a posteriori, and that knowledge of what is said plays an epistemic role in the acquisition of testimonial knowledge, meaning that a posteriori knowledge plays an epistemic role in the acquisition of testimonial knowledge. If Malmgren is correct then there cannot be a priori testimonial knowledge. Malmgren sums up the thought as follows:

'However, the following consideration gives us a prima facie reason to think that it plays an epistemic role: suppose John tells me that it is raining, and that I thereby come to know that it is raining—that is, suppose that I gain knowledge by (John's) testimony that it is raining. If you asked me *how* I know that it is raining, then presumably part of my (pretheoretical) answer would be: "John told me," "John said so," or "John said that it is raining." Here is a natural thought about *what I am doing* in giving this answer: I am citing part of my (epistemic) *reason* for believing that it is raining, part of what makes me *warranted* in believing that this is the case. What my answer brings out is that part of my reason for believing that it is raining is that John said so. But that is just to say that my warranted (or knowledgeable) belief about what John said plays an epistemic role in the formation of my knowledge that it is raining.' Malmgren, (2006), 225.

Furthermore, it might be worried that if an audience doesn't know that p has been asserted then they would have believed p even if p had not been asserted (and thus in cases in which p is false)ⁱⁱ. A knowledge requirement is also likely to appeal to the accessibility internalist. Such theorists hold that in order to know that p one must be able to come to know the basis for one's knowledge that p . It is commonly thought that the fact that the a speaker has testified that p is at least part of the basis for one's testimonial knowledge that p . Thus, in order to gain testimonial knowledge that p one must at least be in a position to know that the speaker has said that p .

Despite its initial plausibility the knowledge of what is said requirement is incorrect. Audiences needn't be in a position to know that a speaker has asserted that p in order to gain testimonial knowledge that p .

2. Against the Knowledge Requirement.

The main argument of this paper rests on two primary assumptions. Firstly I assume a roughly Gricean intention based view of communication. That is, I assume that in order for a speaker to say that p they must intend to say that p (at least for utterances containing context-sensitive terms such as demonstratives and definite descriptions). The second assumption depends on the first: if an audience member is not in a position to know whether the speaker intended to say that p then the audience is not in a position to know that the speaker said that p . If, for all the audience knows, the speaker did not intend to say that p , then since the speaker's intention to say that p is necessary for their having said that p , for all the audience knows the speaker did not say that p . For the majority of the paper I also assume that safety is a necessary condition on knowledge. I adopt these assumptions because the Gricean view of 'what is said' is the dominant view in the philosophy of language, and the safety principle on knowledge is still the dominant anti-luck condition in epistemology. Moreover, I am sympathetic to both views. However, as I note in 2.1.2 and the conclusion, these assumptions are not essential. This is important since it is easy to worry that I am attacking a straw man. That is, it is easy to read the cases presented here and feel that very few people are committed to the view that testimonial knowledge requires knowledge of what is said where 'what is said' is spelled out in the terms presented here. This may be correct (although the Gricean view certainly seems to be the default position on 'what is said' for many). However, as long as one employs a notion of 'what is said' in the conditions one gives for testimonial knowledge one will need to give an account of the conditions under which a proposition counts as 'what is said'. As I explain in the conclusion cases similar to the ones given here will be available for sets of conditions which leave space for the audience to be mistaken about whether or not p has been said.

I will present two cases in which the audience appears to gain testimonial knowledge that p despite not being in a position to know that the speaker intended to communicate that p . If the two aforementioned assumptions are correct then these will also be cases in which the audience gains testimonial knowledge that p without knowledge that the speaker said that p . There are several reasons to discuss multiple cases. Firstly, considering several cases allows us to identify a common structure for counter examples to the knowledge of what is said requirement. This allows us to see how we might generate further counter examples based different assumptions about the conditions for what is said, and the conditions for knowledge. Secondly, intuitions about cases are notoriously unreliable, especially intuitions about cases involving far fetched sci-fi scenarios, or cases with very complex set-ups. Thus it is worth considering several cases in order to ensure that we are not led astray by some misleading feature of one case. This is especially important in this instance, since the second case I present involves the sort of problem case discussed by philosophers of language. As a result, there is a risk that intuitions about the case will be influenced by prior theoretical commitments of the reader. Thus, it is worth starting with a case which abstracts away from these issues before considering a case which bears closer parallels to more normal cases in which audiences lack knowledge of what is said. Finally, each case gives rise to a different intuitive line of response. These responses are considered and rejected. However, the reader might not find every counter-response convincing. Thankfully, they need only find one of the counter-responses convincing in order for the project of this paper to be successful.

2.1.1. Case One.

MAD SCIENTIST^{iii,iv}: The philosophical mad scientist is at it again. His victim is

Sally, a car enthusiast. This time, instead of envying his victim, he has implanted a special chip in her brain. This chip causes her to sometimes say 'that is a fuel efficient car', but only when she is in the presence of fuel efficient cars, and only when she does not intend to do so. It works as follows: whenever Sally is in the presence of a fuel efficient car it turns on and randomly selects one of two values. If it selects value 1 it switches off and becomes inactive again. However, if it selects value 2 it has the following effect: If Sally doesn't intend, and doesn't gain the intention, to comment on the fuel efficiency of the car, then it forces her to utter the sentence 'that is a fuel efficient car'.

One day Sally and Matt are walking through the city when Sally sees a particularly fuel efficient car. She considers commenting on its fuel efficiency, but hesitates because she doesn't know if Matt has any interest in cars. She decides on a whim to just go for it, and says 'that is a fuel efficient car!'. On the basis of Sally's assertion Matt forms the true belief that the car is fuel efficient. What neither Sally nor Matt know is that this was a case in which the chip selected value 2, so Sally would have uttered 'that car is fuel efficient' even if she did not intend to.

I submit that Matt gains testimonial knowledge that the car is fuel efficient. His belief about the fuel efficiency of the car was formed on the basis of Sally's testimony, and it is safe and sensitive (he forms the same belief in all nearby worlds, it is true in all nearby worlds, and in the closest worlds where it is false neither Sally's intentions nor the chip cause her to utter the sentence 'that is a fuel efficient car'). His belief is causally related to its truthmaker, and Sally undertook all the usual commitments one undertakes when one provides testimony. Indeed, as emphasised by Luzzi (2010), if cases like this don't count as cases of knowledge then they constitute interesting new forms of Gettier cases not answerable in any of the standard ways.

However, Matt does not know (nor is he in a position to know) that Sally intended to communicate that the car was fuel efficient. This is because his belief that she intended to communicate that the car was fuel efficient was neither safe nor sensitive^v. In the closest worlds in which Sally doesn't have the intention to say that the car is fuel efficient he still forms the belief that she does intend to say that it is fuel efficient, because the chip makes her say it. Moreover, since Sally was very hesitant about saying that the car was fuel efficient and only decided to do so on a random whim there are many nearby worlds in which she doesn't form the intention to comment on the fuel efficiency of the car. Thus, if assumptions 1 and 2 are correct then Matt gains testimonial knowledge that the car is fuel efficient despite not knowing that Sally said that the car is fuel efficient.

2.1.2. Response to Case One.

It might be argued that in this case Matt does actually have knowledge about what is said. There are two ways this might be done. Firstly, one might maintain that although Matt's belief is unsafe it still counts as knowledge. Secondly, one might maintain that Matt's belief is in fact safe. I will discuss these options in turn.

The first response holds that rather than being a case of testimonial knowledge without knowledge of what is said, MAD SCIENTIIST is a counter example to the safety principle on knowledge. That is, it might be argued that although Matt's belief about Sally's intention is unsafe he nonetheless knows that Sally intended to communicate that the car was fuel efficient. Indeed, this case does bear some resemblance to a series of cases which have been employed as counterexamples to the safety condition on knowledge (see Neta and Rohrbaugh (2004) and Comesaña (2005))^{vi}. For example, Neta and Rohrbaugh present the following case:

EXPERIMENT: 'I am participating in a psychological experiment, in which I am to report the number of flashes I recall being shown. Before being shown the stimuli, I consume a glass of liquid at the request of the experimenter. Unbeknownst to either of us, I have been randomly assigned to the control group, and the glass contains ordinary orange juice. Other experimental groups receive juice mixed with one of a variety of chemicals which hinder the functioning of memory without a detectable phenomenological difference. I am shown seven flashes and judge, truly and knowingly, that I have been shown seven flashes. Had I been a member of one of the experimental groups to which I was almost assigned, I would have been shown only six flashes but still believed that I had been shown seven flashes due to the effects of the drug. It seems that in the actual case I know that the number of flashes is seven despite the envisaged possibility of my being wrong. And yet these possibilities are as similar in other respects as they would have to be for the experiment to be well designed and properly executed.' Neta and Rohrbaugh, 2004, p 400.

I find this case (and others like it) unconvincing, since I find myself unable to shake the feeling that the agent in such cases does not genuinely know^{vii}. However, the issues at play in this debate are subtle, and I do not want to rest my case on my own (no doubt theory laden) intuitions. Luckily I don't have to. Neta and Rohrbaugh present a modified version of their case (analogous to Goldman's (1976) Fake Barns case) in which it seems that the agent clearly does not know that they have been shown seven flashes. The case is as follows:

EXPERIMENT*: 'I am taking part in a long series of psychological experiments, in each of which I am to report the number of flashes I recall being shown to me after ingesting a glass of liquid. In this one case, I have been assigned to the control group and the liquid is ordinary orange juice. I am shown seven flashes and judge, truly, that I have been shown seven flashes. In some of the other trials in which I have participated, I have been assigned to an experimental group in which the liquid also contains a drug which interferes with memory, and the beliefs I formed on those trials were false'. Neta and Rohrbaugh, 2004, p 402.

We can modify the mad scientist case in a similar way in order to secure the intuition that Matt does not know that Sally intended to comment on the car's fuel efficiency. Imagine that Sally and Matt have spent the afternoon together, and Sally has commented on the fuel efficiency of many cars. In some of these cases she has been caused to vocalise the sentence 'that is a fuel efficient car' as a result of the chip's interference, in other cases she did so of her own volition. We can even imagine that Sally's hesitation to comment on the fuel efficiency of the car was due to her noticing Matt's seeming lack of interest in her previous comments on fuel efficiency^{viii}. In this version of the case Matt will have formed many false beliefs about Sally's communicative intentions, and he will have been lucky in those cases where the beliefs he formed were true. This is clearly a situation in which his faculties for assigning communicative intentions are not reliable. Thus it seems highly implausible to claim that he knows that Sally intended to comment on the car's fuel efficiency.

Nonetheless, Matt's actual belief about the fuel efficiency of the car is still safe and sensitive. After all, the chip never makes Sally utter 'that is a fuel efficient car' when she is not in the presence of a fuel efficient car. Additionally, his belief is still caused by its truthmaker, and the causal chain which actually lead to his belief was not deviant in any way. Finally, Sally undertook all the usual commitments when she asserted. Thus, it seems that Matt still gained testimonial knowledge, despite the fact that he did not know Sally's communicative intentions.

A second way of maintaining that Matt does know Sally's intention is to hold that his belief about her intention is safe after all^{ix}. As we have just seen, MAD SCIENTIST bears parallels to the cases given by Neta and Rohrbaugh (2004) as counterexamples to safety. Tomas Bogardus (2012) has

responded to Neta and Rohrbaugh, maintaining that their cases are not really examples of unsafe knowledge, since, according to Bogardus, the agent's belief is in fact safe. Bogardus argues that Neta and Rohrbaugh (as well as Comesaña (2005) and Kelp (2009)) make the mistake of moving from the fact that an agent was at epistemic risk just before they formed their belief to the claim that they were at epistemic risk whilst forming their belief. For example, in EXPERIMENT the agent was at epistemic risk until it was decided that they would be in the control group. After this point the risk is eliminated. The same line of response seems viable here - before Sally made the snap decision to comment on the car Matt was at epistemic risk. However, once the decision is made the risk is eliminated. Bogardus maintains that when assessing the safety of a belief we should move away from intuitive judgements about similarity of worlds and instead focus on the ordinary language statements of safety principles in the literature, asking^x:

1. Is it the case that were S to believe thusly, she would believe truly?
2. Would S not easily believe that p without it being the case that p?
3. Did S form her belief in a way that could very easily have delivered error?

Bogardus maintains that when we apply such heuristics to the cases given by Neta and Rohrbaugh it is natural to judge that the belief formed was in fact safe. For example, in EXPERIMENT, given the circumstances as they stand (with the agent's faculties for counting flashes working normally) they would not easily have been led to error.

It is less clear that Bogardus's heuristic yields a safe belief when applied to MAD SCIENTIST. As Bogardus notes, we must be careful not to hold all the facts of the case fixed when judging the relevant counterfactuals. Most importantly, we must not hold fixed the truth of the agent's belief. Otherwise Bogardus's tests would render every true belief safe. The belief in question here is Matt's belief that Sally intended to say that a particular car is fuel efficient. If we hold this aspect of the case fixed then clearly Matt's belief will come out as safe. That is, given that Sally did intend to comment on a particular car it is true that Matt would not have easily been led to error. However, if we don't hold Sally's intention fixed it becomes far less clear that Matt could not easily have been led to error. After all, if Sally had not intended to refer to a particular car Matt still would have formed the belief that she did have this intention, and this situation very nearly arose. This stands in contrast to Neta and Rohrbaugh's case where it is fairly natural to judge that the agent could not easily have been led astray, even if the belief the agent formed in the actual world were false. Thus, I don't think that it is obvious that Matt's belief comes out as safe when we apply Bogardus's heuristics. However, this really depends on what we hold fixed about the case, meaning that intuitions are likely to vary^{x1}.

Thankfully there is a second way to avoid this worry. As noted in response to Neta and Rohrbaugh we can modify MAD SCIENTIST in such a way that it parallels Goldman's fake barn cases. In the modified fake barn case the chip has already caused Sally to comment on the fuel efficiency of several cars, and as a result Matt has already formed several false beliefs about Sally's intentions. Indeed, we can imagine that in this case the majority of Sally's vocalisations of 'that is a fuel efficient car' are caused by the chip. In this situation Matt's ability to correctly judge whether or not Sally has intended to comment on the fuel efficiency of a car are clearly unreliable. Thus, his belief about her intentions, even when true, does not constitute knowledge.

A final worry about this case is that when the chip in Sally's brain causes her to say 'that is a fuel efficient car' she is not really testifying. Rather, in such cases she is acting as an instrument. If this were the case then Matt's belief would not be formed via the same method in each case. This is not a problem. In all the worlds in which Matt does form a belief about the fuel efficiency of the car on the basis of Sally's actual intentional utterance his belief is still true. So his belief about the fuel efficiency of the car is still safe. Matt might be thought to be lucky in one sense - he was lucky that

his belief was testimonial. But he was not lucky that it was true. Moreover, although he might be thought to gain knowledge about the car via a different method in the worlds in which Sally never forms an intention to assert that the car is fuel efficient, his belief about Sally's actual intentions is formed via the same method, for his evidence regarding Sally's intentions is the same in each case (his evidence being the fact that she opened her mouth and the words 'that is a fuel efficient car' came out). So the chip worlds are relevant to the assessment of the safety of his belief about Sally's intentions, even if they are not relevant to the assessment of the safety of his testimonial belief.

2.2.1. *Case Two.*

I take MAD SCIENTIST to establish that we can have testimonial knowledge without knowledge of what is said. However, it is a rather complex and far fetched case. Thus, I imagine that the dedicated proponent of the knowledge condition might well remain unconvinced. Additionally, the question of the audience's knowledge of what is said is of interest due to worries raised in the philosophy of language about the successful recovery of the speaker's intended meaning. Yet, MAD SCIENTIST bears little similarity to the problematic cases discussed in the philosophy of language. Thus it is worth considering a second case which bears directly on problems of context-sensitivity^{xiii}.

The general problem is that in many cases of context-sensitivity there will be multiple very similar values which can be assigned to the context-sensitive term, each determining a different proposition^{xiii}. The evidence provided by the utterance and the context will often be insufficient to grant the audience knowledge that the speaker intended one of these propositions rather than another. It is worries along these lines (along with worries about our ability to share Fregean contents) that led the theorists mentioned in the first section to reject the view that speakers and audiences must share contents in order for communication to be successful. However, as the final case indicates, these epistemic difficulties do not always rule out testimonial knowledge

The case I am about to present harks back to Wettstein's 1981 argument against the Russellian theory of definite descriptions. Wettstein noted that in most attributive uses of a definite description the explicitly articulated description will not identify a unique referent. Consider the sentence 'the book is red' - there are many books in the world, indeed there are many red books, thus the sentence alone does not identify the unique red book intended. The description must be supplemented in order to designate a unique book. However, there will be many potential supplementations which would uniquely identify the intended book. Wettstein asks us how one supplementation could be determinately selected, and he offers us an alternative theory of definite descriptions on which the content is determinate. However, it has since been observed that similar problems arise for many context-sensitive terms. For example, Schiffer 1992 notes that it arises for belief reports, and Buchanan 2012 notes that it arises for quantifiers and non-sentential assertions. Thus, similar cases can be given involving quantifier domain restriction and non-sentential assertion. The case, adapted from Donnellan 1966, is as follows:

MURDERER: Sally and Matt are investigating a chain of murders they know to be related. Sally is standing over the mangled remains of Frank - the latest victim. Frank has clearly been murdered in a rather brutal way, and the scene has been covered in insane etchings and other clear indicators that the murderer is insane. Additionally, both Sally and Matt know that Frank was murdered by the murderer they are investigating (that is, the individual who committed the previous murders). Sally phones Matt and says 'the murderer is clearly insane'. The description used here is incomplete, and could be rendered complete in a number of different ways. For example, Sally could have intended the proposition 'the murderer (of Frank) is clearly insane', or 'the murderer (who we are investigating) is clearly insane'. Matt recovers the proposition 'the murderer (who we are investigating) is clearly insane' (which happened to be the proposition intended by Sally), and comes to believe it on

the basis of Sally's testimony.

Once again I think that it is clear that Matt gains testimonial knowledge in this case. The belief he forms is true, safe, sensitive, and caused by its truth maker. Additionally, Sally undertook all the usual commitments to the proposition that Matt comes to believe. However, since as far as Matt can tell Sally might have intended the other (very similar) proposition, Matt does not know that Sally intended the proposition he recovered.

2.2.2. *Response to Case Two.*

There are several possible responses to this case, however each response fails. Firstly one might attempt to maintain that this case involved a referential use of the definite description. If this were the case then there would not be multiple competing candidate meanings, there would be one single proposition expressed. This response does not seem promising since neither Sally nor Matt know who the murderer is, they merely know that there is a single murderer responsible for all the murders they are investigating. Thus, it does not appear that Matt or Sally have the resources with which to singularly refer to the murderer. Nonetheless, this may just push the proponent of the knowledge of what is said condition to adopt a very liberal view of the conditions required for singular reference. Thankfully we needn't rest on the claim that this case involved an attributive use of the definite description, for a directly analogous case can be given with quantifiers. Consider the following:

MURDERER*: In this case Sally and Matt are once again investigating a chain of murders they know to be related. Every victim is a member of 'The secret Society of Evil', and the same insane etchings have been found at each crime scene. However, in this case there is significant DNA evidence which suggests that (and which leads Matt and Sally to believe that) different murderers committed each crime. This time Sally phones Matt and says 'Every murderer is clearly insane'. Clearly Sally is not saying that every murderer in the world is insane, rather she is saying that a certain group of murderers are insane. Thus, in this case a property must be supplied to restrict the quantifier. However, there are several equivalent properties which Sally could intend. For example, she could intend 'Every murderer (we are investigating) is clearly insane', or 'Every murderer (of a secret society of evil member) is clearly insane'. Matt recovers 'Every murderer (we are investigating) is clearly insane', and as it happens this is the proposition Sally intended.

Once again it seems clear that Matt gains testimonial knowledge. However, it also seems clear that he did not know Sally's communicative intention, for she could, for all he knows, easily have intended to communicate 'Every murderer (of a secret society of evil member) is clearly insane'. However, this time the referential use response is not available^{xiv}.

A second response might be to hold that in a case like this Sally really asserted both propositions. One might hold that in typical cases a speaker actually asserts many different propositions (for a view like this see Cappelen and Lepore 2004). This seems intuitive, since it certainly seems strange to imagine Sally specifically intending one of the above readings and not the other. They seem to come together. I do not believe that this undermines the case however. The case is a simplification, and we can see that when the case is rendered more realistic the problem remains. It is unlikely that Sally intended or asserted only one proposition. It is more likely that there is a set of propositions consistent with Sally's communicative intentions. However, just as there seemed to be a problem with Matt recovering the precise proposition Sally intended there seems to be a problem with him recovering the precise set of propositions which are consistent with Sally's intentions, or with him recovering only propositions which are in this set. We can conceive of cases in which, for all Matt knows, the proposition he recovers might not be in the set of propositions consistent with Sally's

intentions. For example, imagine that Sally and Matt both have access to a huge amount of information about the murderer, and this causes them to think about the murderer in different terms at different times. Matt has spent all day looking at the connections between the murders and several gangs, and has been thinking of the murderer as 'the gang affiliated murderer'. He has not been thinking about the murderer in terms of the many financial crimes which the murderer has committed. Although he knows the murderer has committed a set of financial crimes he often has to do some cognitive work to bring this information to mind. Sally on the other hand has spent all day thinking about the financial ties in the case. However, just before she phones Matt she sees a gang sign on the victim's arm and this raises the gang affiliations back to salience. As a result the resolutions of the description which relate to gang affiliation do happen to fall under her intention, and thus the proposition Matt recovers does fall under Sally's intention. However, it easily could have failed to do so. And in this situation Matt still would resolved the context sensitivity in the same way. So, for all Matt knows the proposition he recovers is not in the set of propositions consistent with Sally's communicative intentions. Yet, it still seems he gains testimonial knowledge.

3. Concluding Remarks.

It has been argued that despite its intuitive appeal, the claim that audiences must know what has been said in order to gain testimonial knowledge is false. The argument presented here assumed that a speaker must intend to say that p in order to say that p (at least in cases involving context sensitivity). A safety condition on knowledge was also assumed through most of the paper. These assumptions were adopted because the Gricean view of 'what is said' is the dominant view in the philosophy of language, and the safety principle on knowledge is still the dominant anti-luck condition in epistemology. However, we are now in a position to see that these assumptions were not essential. The cases presented all have the following features: a speaker asserts a proposition p , the speaker meets all the normal requirements for the transmission of knowledge (e.g. the speaker knows that p , and commits to p etc.), the audience recovers p through the usual means, is justified in believing that p has been asserted, and forms a belief that p . Moreover, the audience's belief that p meets the usual conditions on knowledge (e.g. safety, sensitivity, being caused by its truthmaker, local reliability etc.). However, there is some necessary condition C for a speaker's saying that p such that the audience is not in a position to know that C is met.

In the cases presented C was a speaker intention requirement. However, as long as it is possible for a speaker to say that p without the audience thereby coming to know that p has been asserted it seems that it will be possible to generate similar cases with different C conditions. For example, if we confined ourselves to cases involving demonstratives and adopted a view on which a demonstration is required to pick out the referent of a demonstrative (for example, Reimer 1991), then we could produce cases in which an audience is not in a position to know what object the demonstration picks out, but in which the testimonial belief they form nevertheless meets the conditions on knowledge. Alternatively we might suppose that very small differences in the patterns of use give rise to subtle changes in meaning (see Dorr and Hawthorne 2014 for a discussion of the consequences of such views). Then we might consider cases in which the audience is not in a position to know that the pattern of use is such as to determine the meaning they assign, but in which they still seem to meet the conditions for knowledge with respect to the object of their testimonial belief. These will all be cases in which any belief an audience might potentially form regarding condition C will fail to meet some condition on knowledge (be it safety, sensitivity, local reliability, some causal condition etc.), but in which some feature of the environment guarantees that their belief in the proposition asserted does meet all the conditions for knowledge. This is important, since it is easy to worry that the argument presented in this paper attacks a straw man. That is, it is easy to worry that few epistemologists endorse the view that testimonial knowledge requires knowledge of what is said where p must be intended in order to count as 'said'. This might

be true (although such epistemologists would be at odds with the standard view in the philosophy of language), but as long as epistemologists of testimony appeal to a notion of 'what is said' they will need some account of the conditions under which p counts as being said. And an argument along the lines of the one presented here will be available for any set of conditions which leave open the option of p being said without the audience to be in a position to know that p has been said.

Thus, attempts to solve the problem by adopting a different metasemantics for context sensitivity, or by rejecting the safety condition on knowledge, seem unpromising. One might reject the assumption that a speaker can say that p without the audience being in a position to know that they have said it (a similar condition is endorsed by King 2014, forthcoming). However, in order to do so one would have to not only deny that testimonial knowledge was gained in any of the cases outlined in this paper, but also deny that anything was said. This seems like a radical move. So, it seems the best response is to embrace the notion that knowledge of what is said is not required for testimonial knowledge. It is not clear what alternative relation (if any) an agent must bear to the object of testimony in order to gain testimonial knowledge. However, in order for a proposal to be immune to the sorts of counterexamples presented here it seems the best strategy will be to search for a relation which is entailed by the collective presence of the other preconditions for testimonial knowledge (including the speaker's having said that p). Otherwise there will be space to present a case in which the other conditions for testimonial knowledge are met (including the causal, anti-luck, and reliability conditions on knowledge) in which the condition does not apply. In such a case we will need strong reasons to deny that testimonial knowledge is gained^{xv}.

Arché Philosophical Research Centre

University of St Andrews

References.

- Baumann, P. (2014). 'No Luck With Knowledge? On a Dogma of Epistemology'. *Philosophy and Phenomenological Research* 89, pp. 523-551.
- Bezuidenhout, A. (1997). 'The Communication of De Re Thoughts'. *Nous* 31, pp. 197-225.
- Bezuidenhout, A. (2002). 'Truth Conditional Pragmatics'. *Philosophical Perspectives* 16, pp. 105-134.
- Bogardus, T. (2014). 'Knowledge Under Threat'. *Philosophy and Phenomenological Research*. 88, pp. 289-313.
- Buchanan, R. (2012). 'Is Belief a Propositional Attitude'. *Philosophers Imprint* 12, pp. 1 - 20.
- Cappelen, H. and Lepore, E. (2004). *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Oxford: Blackwell.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Blackwell.
- Comesaña, J. (2005). 'Unsafe Knowledge'. *Synthese* 146, pp. 395-404.

- Donnellan, K. (1966). 'Reference and Definite Descriptions'. *Philosophical Review* 75, pp. 281-304.
- Dorr, C., Hawthorne, J. (2014). 'Semantic Plasticity and Speech Reports'. *Philosophical Review* 123, pp. 281-338.
- Fricker, E. (1994). 'Against Gullibility'. In B. K. Matilal and A. Chakrabarti (eds.), *Knowing from Words*. Amsterdam: Kluwer Academic Publishers, pp. 125-161.
- Fricker, E. (1995). 'Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony'. *Mind* 104, pp. 393-411.
- Fricker, E. (2003). 'Understanding and Knowledge of What is Said'. From A. Barber (ed) *Epistemology of Language*. Oxford: Oxford University Press, pp. 325-367.
- Fricker, E. (2006). 'Second Hand Knowledge'. *Philosophy and Phenomenological Research* 73, pp. 592-618.
- Goldberg, S. (2007). *Anti-Individualism - Mind and Language, Knowledge and Justification*. Cambridge: Cambridge University Press.
- Goldman, A. (1976). 'Discrimination and Perceptual Knowledge'. *Journal of Philosophy* 73, pp. 771-791.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Heck, R. (2002). 'Do Demonstratives Have Senses?'. *Philosopher's Imprint* 2, pp. 1-33.
- Hinchman, E. (2005). 'Telling as Inviting to Trust'. *Philosophy and Phenomenological Research* 70, pp. 562-587.
- Kelp, C. (2009). 'Knowledge and Safety'. *Journal of Philosophical Research* 34, pp. 21-31.
- King, J. (2014). 'Speaker Intentions in Context'. *Noûs* 48, pp. 219-237.
- King, J. (2014). 'The Metasemantics of Contextual Sensitivity'. From A. Burgess and B. Sherman (eds.) *Metasemantics: New Essays on the Foundations of Meaning*. Oxford: Oxford University Press, pp. 97-118.
- Luzzi, F. (2010). 'Counter-Closure'. *Australasian Journal of Philosophy* 88, pp. 673-683.
- Malmgren, A. (2006). 'Is There A Priori Knowledge by Testimony?'. *Philosophical Review* 115, pp. 199-241.
- Moran, R. (2005a). 'Problems of Sincerity'. *Proceedings of the Aristotelian Society* 105, pp. 341-361.
- Moran, R. (2005b). 'Getting Told and Being Believed'. *Philosopher's Imprint* 5, pp. 1-29
- Neta, R. and Rohrbaugh, G. (2004). 'Luminosity and the Safety of Knowledge'. *Pacific Philosophical Quarterly* 85, pp. 396-406.

- Pritchard, D. (2005). *Epistemic Luck*. Oxford: Oxford University Press.
- Recanati, F. (2004). *Literal Meaning*. Cambridge: Cambridge University Press.
- Reimer, M. (1991). Demonstratives, Demonstrations, and Demonstrata. *Philosophical Studies* 63, pp. 187-202.
- Schiffer, S. (1992). 'Belief Ascription'. *Journal of Philosophy* 89, pp. 499-521.
- Sosa, E. (1999). 'How to Defeat Opposition to Moore'. In J. Tomberlin (ed.), *Philosophical Perspectives* 13, pp. 141-54.
- Sperber, D. and Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford/Cambridge, MA: Blackwell/Harvard University Press.
- Stanley, J. and Szabó, Z. (2000). On Quantifier Domain Restriction. *Mind and Language* 15, pp. 219-261.
- Wettstein, H. (1981). Demonstrative Reference and Definite Descriptions. *Philosophical Studies* 40, pp. 241-257.

- i This answer is already endorsed by Fricker (1994, 1995, and 2003), Christensen and Kornblith (1997), Malmgren (2006), and Goldberg (2007).
- ii Goldberg (2007) pursues this line of argument.
- iii Incidentally, this case also constitutes a counter example to Jeff King's (2014) claim that in order to say that p a speaker must intend to say that p and the audience must be in a position to know that the speaker intends to say that p . This is because the speaker in this case clearly says that p even though the audience member is not in a position to know that the speaker intended to say that p .
- iv The first case is similar to those provided by Frederico Luzzi 2010 as counterexamples to counter-closure (the view that in order to gain knowledge that p by inference from q one must know q). Indeed, if one held a view of testimony on which audiences make a sub-personal inferences from beliefs like 'S said that p , and S is trustworthy' to 'I should believe that p ', then one would be forced to reject counter closure on the basis of the cases presented in this paper.
- v One might worry that audiences do not generally form beliefs about speaker intentions when they form testimonial beliefs. This worry would be consistent with the view that audiences must be in a position to know the speaker's intentions in order to gain testimonial knowledge. However, these cases illustrate that the audience need not be in a position to know the speaker's communicative intentions, since any belief they may potentially form about the speaker's communicative intentions will fail to be safe or sensitive.
- vi There are further examples provided in Kelp 2009, Baumann 2012, and Bogardus 2014. However, these cases bear less resemblance to the cases presented here.
- vii I feel differently about Comesaña's case. In Comesaña's case Juan is planning to dress up as Michael for a Halloween party, but decides against it at the last minute. Judy, has been employed to give directions to the party. She will tell everyone it is in the same location, but if she sees Michael (or someone she thinks is Michael) she will ring the organiser and the party location will be moved. So although someone looking like Michael would receive the same testimony as everyone else, Judy would arrange for that testimony to be false. Since Juan almost dressed up as Michael there is a sense in which he almost received false testimony. However, it strikes me that Juan's belief in this case is not unsafe. The world at which he receives the false testimony is just too different from actuality. Indeed, it is not even clear to me that the worlds in which Juan receives the false testimony are worlds in which the same belief forming method is employed. This is because when we form testimonial beliefs we rely on the speaker to be honest to *us* - the audience. When Juan comes as himself he relies on Judy to be honest to him - to Juan. However, if he dresses up as Michael he relies on Sally to be honest to Michael. Indeed, we could even imagine that Michael knows of the plan to keep him from the party, and decides to dress up as Juan in order to receive truthful testimony as to the party's location. In this case Michael is clearly and intentionally employing a different method than he would if he approached Judy undisguised. Thus a different belief forming method is employed in each case. Thus, since safety is always relativised to a method, Juan's belief is safe (for views which place a special emphasis on this interpersonal relation between the speaker and the hearer see Moran (2005a, 2005b), Hinchman (2005), and Fricker (2006)).
- viii We can imagine that the chip retroactively implants in Sally the belief that she intended to comment on the car's fuel efficiency when she is forced by the chip to do so.
- ix Thanks to an anonymous referee for pointing out this line of response.
- x These questions correspond to the statements of safety given by Pritchard (2005), Sosa (1999), and Hawthorne (2004).
- xi Indeed, it is hard not to get caught up in a web of Sobel sequences and reverse Sobel sequences when applying Bogardus's heuristics. That is, once we hold a particular feature fixed when assessing a given counterfactual it is hard to then unfix that feature in future assessments of the counterfactual, meaning that the order in which we assess counterfactuals with certain factors held fixed will have an impact on our intuitions about their truth. This is not just a problem for the application of Bogardus's heuristics, but something we must be conscious of whenever we assess multiple variations on a philosophical thought experiment involving counterfactual scenarios. The implications of this context sensitivity for initial judgements about cases are less clear.
- xii MAD SCIENTIST is also a case of context sensitivity. However, it is a case of reference assignment, and reference assignment seems less epistemically problematic than many other forms of context sensitivity. In cases of reference assignment there are usually fewer candidate meanings for the audience to choose between. In the cases which are generally seen to be problematic there are a large number of similar meanings for the audience to choose between. It is seen as mysterious how the audience could select the correct meaning from amongst the many available candidates.
- xiii This will not generally be the case when context merely supplies a referent (as is the case with, for example, demonstratives or referential uses of definite descriptions). However, it will be the case when context must supply a more complex value such as a property, comparison class, modal base, ordering relation on worlds, standard of taste etc., and in cases of loose talk or non-sentential assertion.
- xiv One might worry that context does not supply a property to restrict the quantifier, but rather supplies a the domain restriction in the form of a set. This would be a version of the referentialism response for quantifier domain restriction. Indeed, through much of Stanley and Szabó's seminal 2000 paper on quantifier domain restriction they speak as if this is their view. However, as they explain, this is a simplification. They ask us to consider the following case: John buys 70 bottles of beer every time he goes to the supermarket. This time there are only 70

bottles of beer on the shelf, and so he buys every bottle of beer. However, someone could truly utter 'If there were a few more bottles on the shelf John would not have brought every bottle'. If quantifier domain restriction merely provided a set this sentence could not be truly uttered, because 'every beer' would pick out the same set of beers in worlds where there were more than 70 beers on the shelf (Stanley and Szabó, 2000, p 252) .

^{xv} Thanks to Sebastian Becker, Herman Cappelen, Josh Habgood-Coote, Julia Driver, Katherine Hawley, Matt McKeever, Brian Weatherson, and an anonymous referee at this journal for helpful comments and discussion. Thanks also to the attendees of the the Scots Philosophical Association annual meeting in St Andrews, where an early version of this paper was presented.