

Testimony, Pragmatics, and Plausible Deniability.

1. INTRODUCTION.

In order to assign content to a wide range of context sensitive utterances audiences must rely heavily on their knowledge of the conversational context. Because of this reliance on context there are many ways the recovery of content can go wrong. As a result, speakers are able to make assertions and yet deny responsibility for the proposition asserted, claiming that the audience made a mistake in resolving the context sensitivity. That is, speakers are able to maintain plausible deniability about what is said. Call this the 'deniability problem'. The aim of this paper is to explain why the deniability problem is problematic, and start to identify the range of utterances to which it applies.

Elizabeth Fricker (2012) has pushed a similar line of reasoning, arguing that implicatures fail to carry the epistemic force of outright assertions. She argues that speakers maintain plausible deniability about the implicature, meaning that they fail to undertake the commitments necessary to transmit testimonial knowledge. I start by outlining Fricker's view, and arguing that it potentially applies to a far wider range of utterances than those she considers, including many assertions. I then go into greater detail explaining exactly why this is worrying.

The deniability problem is problematic for three reasons. Firstly, many views of testimony (for example, 'telling' based views of testimony such as Fricker (2006a), Hinchman (2005), Moran (2005a, 2005b), and Ross (1986)) emphasise the role of speaker commitments in the justification of testimonial beliefs. The deniability problem entails that speakers often fail to undertake the types of

commitments emphasised by such views. On telling based views of testimony, when an audience is told p they gain a reason to believe p in virtue of the speaker having publicly taken responsibility for the audience's belief that p . However, when the speaker maintains plausible deniability about what is said no such commitment is undertaken. Secondly, plausible deniability prevents epistemic buck passing. Sanford Goldberg (2006), and Benjamin McMyler (2013) have argued that the ability to pass the epistemic buck in response to challenges to one's belief is a distinctive epistemic right agents gain only in virtue of forming testimonial beliefs. This diminishes the belief holder's responsibility for their testimonial belief as compared to beliefs formed via other methods. However, when the speaker maintains plausible deniability the audience loses the ability to pass the epistemic buck, and thus fails to gain the epistemic rights distinctive of testimonial knowledge. Finally, the ability to maintain plausible deniability blocks one of the primary disincentives to deceptive or careless assertion.

After explaining why the deniability problem is problematic I focus on identifying the range of utterances to which it applies. I outline a puzzle arising from the recent debate over context sensitivity in the philosophy of language, which seems to suggest that the deniability problem extends to a very large number of utterances. The puzzle is as follows - on the one hand it has been argued that there is widespread context sensitivity in natural language, and audiences must rely heavily on their knowledge of the context to recover the speaker's intended meaning. This includes cases which we would intuitively treat as being on a par with normal testimony. On the other hand, Herman Cappelen and Ernie Lepore (2004) have argued that such widespread context sensitivity would make communication miraculous, pointing to the same 'problematic epistemics' which, according to Fricker (2012), give rise to plausible deniability. If context sensitivity is very widespread, but has the same 'problematic epistemics' that Fricker identifies for implicature, then the deniability problem is extremely far reaching. This would be a worrying and radical result.

Thankfully this result can be avoided, or at least weakened. There have been several contextualist responses to Cappelen and Lepore which aim to establish that widespread context sensitivity does not make communication miraculous. These responses don't work as general solutions to the deniability problem, however they do allow us to limit its scope. We end up with a set of criteria for identifying discourses which are particularly susceptible to the deniability problem. In the conclusion I suggest some important discourses which may still face the deniability problem.

2. THE DENIABILITY PROBLEM.

Because of what Fricker calls the 'ambiguous epistemics' of implicature speakers are able to maintain plausible deniability about what they have implied. That is, they maintain the ability to deny that they ever intended to imply what the audience takes them to have implied. I take it that such a denial will involve the construction of a believable narrative in which the speaker's attitudes and expectations at the time of utterance were such that they could reasonably make their utterance without intending to imply what they were taken to have implied. More precisely:

Plausible deniability - An agent has plausible deniability about intending to communicate a proposition p with an utterance u of a sentence s if that agent is able to tell a story (with at least some degree of believability¹) about their attitudes and expectations at the time of utterance such that a reasonable agent with those attitudes and expectations could utter s with no intention to communicate p .

The plausibility of the denial will depend on the plausibility of the story about the agent's attitudes.

In order to be plausible the story will have to be consistent with what the audience knows about the

¹ Plausible deniability is clearly a gradable notion since some denials will be more plausible than others. In section five we will consider some factors which might make some denials more plausible than others.

speaker's attitudes and beliefs. That we have such plausible deniability about implicatures is illustrated by the following example:

Implicature: Matt is running out of fuel and needs some fast. He stops and asks a stranger where he can get some fuel. The stranger says 'there is a gas station around the corner'. The stranger thereby implies (implicitly communicates) that the gas station is open and has fuel. However, if Matt were to get to the gas station and find it closed or out of gas he would have a hard time criticising the stranger for her utterance. She could easily maintain that she never intended to communicate that the gas station was open or that it had gas. She can maintain that she was merely suggesting it as a place to try (she could maintain this even if she in fact knew that the gas station was shut or had no gas).

When Fricker talks of 'ambiguous epistemics' she refers to the way audiences and speakers must rely on what she calls 'knowledge context' (or 'K-context') in order to recover the communicated message. Knowledge context is the audience's representation of factors such as mutual knowledge, past utterances, Gricean norms, the mutual goals of the conversation, and any other information generally relevant to interpretation other than basic knowledge of the syntax and invariant semantic content of the utterance. Because of the many complex ways in which audiences must rely on context in order to recover an implicature there are many ways the process can go wrong. For example, the audience could employ an aspect of knowledge context which the speaker never intended them to employ, or they could fail to employ an aspect of knowledge context that the speaker did intend them to employ. Likewise, they could hold false beliefs about the context, or at least beliefs which were not mutual knowledge, and appeal to such beliefs in their recovery of the implicature. Because the recovery of implicatures can go wrong in so many ways the speaker is able to deny responsibility for the proposition the audience recovers, claiming that there was a

mismatch between what the audience recovered and what the speaker intended (perhaps even claiming the speaker didn't intend any implicature at all)². Fricker's own take on the problem is as follows:

Entirely genuine misunderstandings and mistakings are endemically liable to happen, regarding a supposed message that is conversationally implied, not stated, due to the very complex mutual epistemics of the situation. Given these complex epistemics, it is not epistemically feasible to pin undeniable specific commitment onto a speaker: she can always wriggle out of it. This may be in bad faith; but very often it may not—maybe she miscalculated what her audience would infer; maybe she had not really figured it out.

(Fricker 2012: 87)

Since the claim that she intended to communicate *E* turns on claims about her private intentions and K-context, including her second order representations of others' beliefs, she can always get away with denying that she intended any such thing; even if her denial is made in bad faith. Lies about my own intentions and other mental states may be suspected, but cannot be refuted. In contrast, when someone makes an explicit statement of a fact *P*, what she signs up to in doing so—taking responsibility for the truth of *P*—is a public fact about the situation, determined by semantics and objective features of context. So it cannot

²It has been pointed out by an anonymous reviewer that this problem actually cuts both ways. In some cases contextual indeterminacy seems to grant audiences the freedom to recover the meaning they find most convenient. The reviewer provides the following example : One's partner may say that they are going to have coffee at the local café, and one might interpret them as meaning that they are going to the café *now*, using it as an excuse to drink the last of the coffee. One might do this even if one knows that they might be planning to go later (in which case they may still want some of the remaining coffee). I think this raises some interesting questions about the way in which speech acts make various courses of action permissible. However, there is not enough space to give this issue a proper treatment here.

be incorrigibly denied by the speaker. I can be nailed as having stated that *P*; never as having insinuated that *P*. (ibid: 88-89)

Fricker seems to suggest that the deniability problem only applies to implicatures, suggesting that plausible deniability is never available regarding the primary content of one's utterance. The thought is that one can deny that one intended to imply anything without descending into absurdity. However, if one makes an intelligible declarative utterance one cannot, without descending into absurdity, claim that one never intended to communicate anything. This may be true, but one can still maintain plausible deniability about intending the particular proposition the audience recovered without claiming that one never intended to communicate anything at all. To see that speakers do attempt such conversational manoeuvres with asserted contents as well as implicatures consider the following two examples (the first of which occupies a grey area between what is implied and what is asserted, the second of which clearly concerns asserted content)³:

Scalar implicature: We are planning a group trip to a theme park and deciding how many cars to take. I wish to cause logistical problems because I hate fun, so I say 'Matt has three kids' knowing that he has five. On the day of the trip Matt arrives with his five children and we don't have enough space in the car. You challenge me for saying that Matt had only three children. However, I might attempt to maintain that I didn't mean he had *only* three children, I meant that he had *at least* three children, so we would need at least three additional seats.

³ As we move through the examples from the clear case of implicature to the case of mere context sensitivity the stories the speaker tells start to sound slightly less plausible. They would certainly raise our suspicions, and if a speaker frequently made manoeuvres like these then we would consider them untrustworthy. This is a point to which I will return when discussing responses to the problem. For now it suffices to note that on one-off occasions we would usually let such matters slide and speakers would usually get away with making such conversational manoeuvres. This is not to say that they wouldn't be criticised for being unclear, simply that they would not be held to what they communicated.

The plausibility of such a story will depend on the way I said 'Matt has three kids' and the immediate preceding utterances, however we rarely recall such minute details of the conversational context, so it would not be difficult to construct a plausible story on which I intended to communicate only that Matt has at least three children.

Quantifier domain restriction: It is the start of a new year and we have organised a party for the new graduate students. We have a variety of beers on offer, but there are some special craft lagers I want for myself (even though they were brought for the guests). I have stored most of the beer in the fridge, but I have put the craft lagers outside. Sally, one of the new students, arrives and asks where the beer is, so I tell her 'every beer is in the fridge'. Later on you find the craft lagers outside and ask me why I told Sally that every beer was in the fridge. In response to this challenge I might attempt to construct a story along the following lines: I had heard that Sally was a vegan, and I am aware that craft lagers often contain animal products. So when I said 'every beer is in the fridge' I didn't mean every beer we had purchased for the party, I meant every beer which was safe for Sally, as a vegan, to drink.

In these examples the speaker attempted to construct a narrative concerning their attitudes and representations of the context in which the audience's recovered meaning was not intended. In general, if recovery of an asserted content requires extensive appeal to knowledge context then the speaker will often be able to claim that the audience recovered the wrong proposition, thus disclaiming responsibility for the audience's belief. Therefore, if heavy duty appeal to knowledge context is often required for recovering what is said, and this appeal to knowledge context gives rise to the same possibilities of error to which implicature gives rise, then the deniability problem will apply to a wide range of assertions, not just implicatures.

Before moving on it is worth clarifying what is meant when I talk about 'recovering the content of an utterance' or 'recovering what is said'. This is important both for the sake of appreciating the deniability problem and also in order to eliminate an objection which threatens the project of this paper from the outset⁴. In the examples I have given the hearers come to believe some fine grained proposition, and the speaker is able to maintain plausible deniability by claiming that a different fine grained proposition was intended. It might be thought that in reality we only come to believe far more coarse grained propositions, for example the disjunction of all the propositions the speaker could plausibly claim to have meant. After all, it might be thought that we often fail to consciously assign particular fine grained values to contextual variables. Moreover, the deniability problem breaks down on the view that we recover the disjunction of the propositions that the speaker could plausibly have meant. The speaker cannot plausibly deny having meant to communicate something at least as strong as the proposition recovered if the proposition recovered is the disjunction of propositions the speaker could plausibly claim to have meant.

I think this line of reasoning is problematic because it relies on an unrealistic view of testimonial belief formation. We frequently do form testimonial beliefs in fine grained non-disjunctive propositions when there are other propositions the speaker could have claimed to have meant. For example, consider the above case of scalar implicature: on the disjunctive view the recovered proposition communicated would be 'Matt has either exactly three kids or at least three kids'. The problem is that this just collapses into 'Matt has at least three kids', yet it is plausible that we often take utterances like 'Matt has three kids' to communicate that Matt has exactly three children. On the disjunctive view such readings would be rare. Moreover, the range of situations an imaginative speaker would be able to conjure up in order to claim that a miscommunication has occurred will often be rather wide, and an unimaginative hearer (or just a hearer who is interpreting quickly and unreflectively) is unlikely to consider (consciously or subconsciously) the whole range of cases a

⁴ Thanks to an anonymous reviewer for drawing my attention to this objection.

speaker could construct. Thus hearers will be able to infer little from what they have been told due to the level of uncertainty over exactly what it is that they have been told. Consider the quantifier case again. I mentioned one story the speaker could conjure up in order to maintain plausible deniability, however there are many more. For example, the speaker might attempt to claim that 'beer' and 'lager' are distinct, and that the question Sally asked was asked about beer (some people do consider this to be an important distinction), or the speaker could maintain that the intended interpretation was 'every beer which is ready to drink', claiming that the beer left outside was still warm at the time of utterance. On the disjunctive view the proposition the hearer recovers would be 'every beer for the party is in the fridge, or every vegan friendly beer is in the fridge, or.....!'. It is not clear why Sally would consider all these readings (especially if she is not actually a vegan), and even if she did she would not be entitled to assume that there is no other unconsidered situation the speaker could conjure up in order to claim miscommunication. Thus really all she would be entitled to take from the utterance would be that there is *some* beer in the fridge. However, we frequently take much more from such utterances. This is evidenced by the fact that Sally may well be surprised to find the craft lagers outside, or she may assume they were someone's private stash. She would not be so disposed unless she had come to believe something like 'every beer for the party is in the fridge' (she may not have explicitly considered such a restriction, but this does not rule out forming a dispositional or unconscious belief with such a domain restriction). So I consider the disjunctive view to be implausible.

A alternative (perhaps more plausible) way of spelling out the disjunctive view holds that audiences do not take speakers to mean ' p or q or...', but rather that audiences come to believe 'the speaker means p , or the speaker means q , or....'⁵. This view faces many of the same problems however, as the audience is still only able to draw a disjunctive conclusion about the subject matter of the assertion. For example, suppose I assert 'Matt has three kids', and the audience forms the belief

⁵ Thanks to an anonymous referee for pointing out this alternative version of the disjunctive view.

'Either S meant 'Matt as exactly three kids' or S meant 'Matt has at least three kids'. The audience will only be able to draw a disjunctive conclusion about the number of children Matt has. That is, they will at best be warranted in forming the belief 'Matt has either exactly three kids or at least three kids'. This clashes with the fact that we would often form the belief that Matt has exactly three children upon hearing 'Matt has three kids'. Other problems remain too. For example, it still seems likely that the range of interpretations an imaginative speaker would be able to conjure far outstrip the number of interpretations the audience is likely to consider when forming their disjunctive view about what the speaker might have meant. Thus, although this modification of the view does perhaps seem more realistic, it still faces many of the same problems.

This is not to say that I believe we always recover only a single fine grained proposition. There may often be some degree of indeterminacy or uncertainty about the precise proposition (or propositions) communicated. However, my contention is that in cases where the mechanisms employed in the recovery of what is said give rise to many error possibilities the range of propositions an imaginative speaker can claim to have intended will often be significantly larger than the range of propositions over which a typical hearer's state of understanding will be indeterminate. Thus speakers will often be able to maintain plausible deniability about having intended any of the precisifications of the audience's interpretation⁶. By reviewing some of the literature on linguistic context sensitivity we will be able to assess the extent to which the mechanisms employed in the recovery of what is said give rise to the sorts of error possibilities which allow for the deniability problem. However, before

⁶ A related concern is that audiences needn't settle on a single precise interpretation, they may settle on several. The problem with this response is that it is unclear whether we can differentiate between settling on several different interpretations and settling on a single interpretation identical to the conjunction of those interpretations. Moreover, even if we did form testimonial beliefs in this way it would not prevent the speaker from maintaining plausible deniability over individual interpretations. In fact, although the strategy of settling on multiple interpretations would no doubt increase the chance of the audience settling on the correct interpretation it would also vastly increase the risk that at least one of the several interpretations settled upon would be incorrect.

doing so it is important that we clarify exactly why the deniability problem is such a problem.

3. WHY IS THE DENIABILITY PROBLEM A PROBLEM?

In the previous section I outlined the deniability problem, and argued that there is no principled reason to hold that it applies only to implicature. The problem, as it applies to assertions, can be stated as follows: usually when one makes an assertion with a clear propositional content p one undertakes a commitment to p . That is, one undertakes a commitment to defend one's belief in, and assertion of p , or else retract the assertion. This commitment is made public in the act of assertion (see MacFarlane (2005, 2011), and Rescorla (2009) for views of assertion which place special emphasis on this fact). However, in certain cases one has another option available when challenged - one is able to deny that the audience recovered the intended proposition. Call this alternative conversational move the 'mismatch move'. When the mismatch move is available no public commitment is undertaken to defend one's assertion or provide epistemic justification for the proposition seemingly asserted.

In this section I will explain why we should be worried by the deniability problem. We should be worried about the deniability problem because the commitments speakers undertake plausibly play an important role in justifying our testimonial beliefs, and in shaping the epistemic rights we acquire when we form testimonial beliefs. The presence of the deniability problem indicates that speakers are able to back out of certain commitments, meaning that the commitments are unable to perform their justificatory or rights shaping roles.

I start by discussing the justificatory role played by speaker commitments. I focus primarily on Fricker's own view (in order to further explicate her take on the problem) and the assurance

theorists, who press the role of commitments intentionally incurred. The commitments so emphasised seem especially susceptible to plausible deniability, thus the deniability problem seems to be particularly worrying for the assurance theorist⁷. I will then briefly consider the extent to which the deniability problem might be seen to extend beyond views such as Fricker's or the assurance theorist's. Next I outline the role speaker commitments have been taken to play in shaping the epistemic rights we acquire as a result of forming testimonial beliefs. Sanford Goldberg (2006), and Benjamin McMyler (2013) have argued that the commitments speakers undertake when testifying shape the epistemic rights of the audience with respect to their testimonial belief. The presence of the deniability problem causes problems for the acquisition of these epistemic rights. Finally I argue that the presence of the deniability problem blocks one of the primary disincentives to deception.

3.1 Fricker and the Assurance Theorists.

Fricker frames her discussion in terms of what she takes to be the paradigmatic mode of transmission of knowledge via testimony - the act of telling. She argues that by telling an audience that p speakers vouch for, and take responsibility for, the truth of p . In telling someone that p the speaker presents p as being true in an act the import of which is that the hearer can form a belief in p on her say so (Fricker (2006a)). This act licences the audience to believe that p in virtue of the fact that it is the 'conventionally constituted force of her speech act' that in asserting p the speaker purports to speak from knowledge (Fricker 2006a, p594). This is a commitment in the public sphere, but it is also manifested publicly to the audience. The knowledge norm for tellings follows, Fricker thinks, from the fact that in telling someone p you offer them your word that p , and commit to it. To explain this Fricker draws an analogy with promising - it seemingly follows from the fact that in promising to perform act a you commit to doing a , that you should promise to a only if you

⁷ Thanks to an anonymous reviewer for encouraging me to focus on this aspect of the assurance theorist's view.

intend to a , and a ing is within your power. Similarly you should not commit to p unless you know that p . It is the fact that knowledge is the norm of telling which imbues it with its epistemic force. Audiences are justified in believing a speaker when they are justified in taking the speaker to be trustworthy with respect to the particular telling. The speaker's trustworthiness with respect to p consists in her having the following dispositional property: not easily would she assert that p , or vouch for the truth of p , unless she knew that p (Fricker (2006a) p600). When you are properly justified in taking a speaker to be trustworthy, and you know they have acted in such a way that they are committed to knowing that p , then you are both justified in taking p to be true, and justified in believing that the speaker's support for p is sufficient to yield knowledge. Once the audience knows the speaker has asserted p , and is in a position to know that the speaker is trustworthy⁸, then the audience is thereby in a position to know that p . This is how knowledge is spread via assertion according to Fricker.

So, according Fricker, speakers don't properly vouch for a proposition if they maintain plausible deniability with respect to their intention to communicate that proposition, and 'tellings' only occur when speakers do undertake such a commitment. Thus, any assertion which leaves open the mismatch move is not a telling in Fricker's sense. However, tellings (and the commitments they generate) are central to Fricker's view of testimonial knowledge. It is in virtue of the commitments tellings generate that the knowledge norm applies to them, and it is in virtue of the fact that the knowledge norm applies to them that they constitute evidence for the proposition asserted. Thus, in Fricker's framework, audiences will not be able to achieve testimonial knowledge from any assertion where the mismatch move is available (unless an alternative story about testimonial knowledge is given to supplement Fricker's 'telling' based view). If the deniability problem applies to a wide

⁸ In earlier work Fricker (1994, 1995, 2006b) has maintained that in order to form a justified testimonial belief audiences must monitor speakers for trustworthiness. Presumably it is such monitoring which will grant the audience knowledge level justification of the speaker's trustworthiness.

range of assertions, or if particular discourses are especially susceptible to the deniability problem, then this should sanction either a widespread scepticism, or a targeted scepticism regarding the particular discourses in question.

Fricker's view is similar in many ways to the assurance view of testimony advocated by Richard Moran (2005a), (2005b), Edward Hinchman (2005), and Angus Ross (1986). Like Fricker assurance theorists take tellings to be the paradigmatic speech act by which knowledge is transmitted, and like Fricker they take tellings to generate a special sort of commitment on behalf of the speaker. Unlike Fricker assurance theorists take the reasons to believe provided by tellings to be non-evidential (that is, the epistemic force of testimony does not primarily consist in testimony that p providing evidence that p).

Assurance theorists distinguish between two ways in which we may learn something by believing another's assertion. Firstly, when someone asserts p we may take this as evidence that they believe p (or perhaps gain the right to believe that they believe that p), and that, since they are likely to be a reliable belief former, we should therefore believe that p . This way of forming beliefs is, according to the assurance theorist, in principle no different from coming to believe that p through observing any aspect of an agent's behaviour and judging that they believe that p . Indeed, Moran (2005b) argues that treating a speaker's utterance as evidence about their beliefs may be worse than treating other behaviour as evidence for their beliefs due to the fact that the evidence provided by assertions is, by its very nature, doctored evidence (see Keren (2012) for a response). Assurance theorists also argue that speakers do not intend for their utterances to be treated as evidence in this way. Ross (1986) argues that in order to take the evidential view one must judge another's utterance by reference to further generalisations about their psychology and the conditions under which they are likely or unlikely to utter particular words. You need to view the assertion in a 'detached objective

light, as a natural phenomenon arising from certain causes' (Ross (1986, p72)). Ross observes that we can obviously view other people's utterances this way, but thinks it is far less clear that we can view our own utterances this way.

I cannot at one and the same time see it as up to me what I shall say and see my choice, as an observer equipped with a theory of speech behaviour might see it, as determined or constrained by facts about my own nature. (Ross 1986: 72)

Such an attitude would, according to Ross, be a form of disengagement from one's own actions, similar to Sartre's 'bad faith'. In order to take an utterance as it is intended by the speaker we must not treat it as evidence. The alternative, according to the assurance theorist, is to treat the speaker's taking responsibility for the audience's belief as a reason for the audience to hold that belief. The idea is that in telling the audience that *p* the speaker gives the audience permission to epistemically rely on them. When we treat the behaviour of others as evidence for their beliefs, and then form beliefs about the world on the basis of taking others to be reliable belief formers, we do not gain the ability to hold others epistemically responsible for our new beliefs. But when others tell us that *p*, and we take them at their word, we apparently do. Moran (2005a) summarises the distinction between the two ways of viewing another's testimony as follows:

Corresponding to the difference between what the speaker 'gives' and what the speaker 'gives off' is the difference between what I learn from him and what I may learn from what he does and how he does it. Only in the case of what I learn from him, the person, does my relation to his belief involve the speaker assuming any responsibility for what I believe, and that makes a difference to the type of reason to believe that is obtained in the two cases. (Moran 2005a: 335)

So the speaker's public commitment to defend the audience's testimonial belief takes centre stage in the assurance theorist's view of testimonial knowledge. It is important to emphasise the role of the speaker's intentions in generating these commitments. It is by openly and intentionally communicating p that the speaker takes responsibility for the audience's belief. It is not clear that one can unintentionally take epistemic responsibility for another's belief. After all, consider the types of reactions we typically have to an unintentional communication of a controversial proposition. When the speaker clearly, openly, and intentionally communicates that p we expect them to present epistemic reasons in defence of p , indicating that we expect the speaker to bear epistemic responsibility for the communicated proposition. However, when a speaker unintentionally communicates p we do not expect them to defend the truth of p , rather we expect them to either defend the reasonableness of their preferred interpretation, or apologise for being unclear (we may also expect some form of compensation for any mishap which resulted from the miscommunication). This indicates that although we sometimes hold the speaker to bear some practical responsibility for propositions unintentionally communicated we do not generally hold them to bear any epistemic responsibility for such propositions. The importance of the speaker's intentions in generating these commitments is emphasised by Moran:

Only with respect to what I have called 'personal expression', the intentional action of expressing one's belief, is the person in a position to speak for the meaning or epistemic import of what he is attesting to. With respect to whatever else may express itself in someone's speech or other expressive behaviour, while this may indeed be a source of knowledge for the audience, they are on their own as far as assessing its epistemic significance goes. Since beliefs which are revealed in these ways need not even be known by the speaker himself, the hearer (or observer) cannot assume that the speaker is in a position

to offer support or justification for what may be garnered in this way, nor that he speaks with any authority about the meaning or general significance of the belief which manifests itself in his speech or other behaviour. (Moran 2005a: 342)

As we observed when discussing Fricker's view, the speaker's public commitment to defend the audience's belief is precisely the type of commitment the deniability problem undermines. This is worrying if the mismatch move is widely available, or if there are particular discourses in which it is widely available. In cases where audiences have to rely heavily on knowledge context the audience's attitude toward the speaker seems to be no less one of trust in the speaker than in the assurance theorist's paradigm cases. Therefore, if the assurance theorist is correct that in the paradigm cases we do not treat assertions as evidence then there are a range of utterances which we do not typically treat as evidence, but which can only provide epistemic reasons in virtue of being treated as evidence (since speakers don't undertake the assurance theorist's required type of commitment in making these utterances). Once again this sanctions a scepticism about the discourses in question.

Indeed, it seems that the assurance theorist's emphasis on the speaker's intentions makes the problem especially worrying. This is because the story the speaker must tell in order to maintain plausible deniability about their own intentions will usually require appeal to little more than claims about their own internal mental states (over which they have epistemic authority). In fact a speaker could even maintain that they were simply not paying attention to what they were saying in order to maintain that a rather obvious interpretation was unintended. In such cases we would certainly hold a speaker responsible for some wrong doing, however we would often still accept that they did not intend to communicate what they in fact did communicate. Thus, if the assurance theorist is correct about the role of speaker intentions in generating the types of commitment required for testimonial

justification then it will often be fairly easy for speakers to make the mismatch move, meaning that it will be widely available.

Of course, there are many commitments and responsibilities we undertake when asserting over which we do not have intentional control. For example, we do not have intentional control over which propositions will actually be communicated by our utterances. However, we still have a moral duty not to make an assertion if it is predictable that in doing so we will communicate something false⁹. If we do make an utterance which predictably results in a false belief then we will often be held responsible for the resultant belief. This responsibility is not epistemic, even if we believe the predictably communicated proposition we are not thereby duty bound to defend its truth. Nonetheless, we will often be held practically or morally responsible for any misfortune which results from the false belief. It may be thought that these commitments are sufficient to provide a basis for testimonial justification. We might reason, in a similar way to Fricker, that since we are committed to not predictably communicating falsities we will try to only predictably communicate truths. Thus, if a speaker makes an assertion such that it is predictable that p would be communicated by that assertion this provides evidence that the speaker knows that p . Such a view would not provide the kind of anti-reductionist justification the assurance theorist seeks. However, speakers have far less plausible deniability over what will be predictably communicated by an utterance, so perhaps such a view could at least provide a reductionist account of the role of speaker commitments in generating testimonial justification whilst mainly avoiding the deniability problem.

Ultimately I think that this view contains more than a grain of truth. However, it is not without its problems. Firstly, speakers are generally far more careless about what they might predictably communicate than what they clearly and openly communicate. Thus, it is not clear just how often audiences will be justified in believing that the speaker knows that p on the basis of the fact that the

⁹ Thanks to an anonymous reviewer for pointing this out.

speaker made an utterance such that it is predictable that that utterance would communicate p . This will depend on just how easily predictable the audience's interpretation was, and on the context (for example, speakers are likely to be far more careful about what they might communicate in a high stakes context than a low stakes context). Secondly, although speakers will have less plausible deniability about what was predictably communicated by their utterance they will might still have some. That is, they might have some plausible deniability regarding their own ability to predict the resultant interpretation. This will involve the construction of a story about the speaker's representation of the context at the time of utterance such that a reasonable agent with that representation could make their utterance without being able to easily predict that they would communicate p . Even if such a story fails to establish that the speaker was not in a position to predict the resultant interpretation it might still convince the audience that their own interpretation was less obvious to the speaker than they thought, thus making the speaker less criticisable for their utterance. Thus it is unclear to what extent such a view does avoid the deniability problem.

So far we have seen why the deniability problem is worrying for views which emphasise the role of speaker commitments in generating testimonial justification. In the next section I will consider the role speaker commitments have been taken to play in shaping the audience's epistemic rights with respect to their testimonial belief.

3.2 Deniability and Buck-Passing.

In the previous section I argued that the deniability problem is worrying for those who emphasize the role of speaker commitments in testimonial justification. An agent's right to form a testimonial

belief is, according to such views, at least partly grounded in the fact that speakers take responsibility for the audience's belief. However, some authors have argued that the epistemic role of speaker commitments extends beyond the justification of testimonial beliefs. They argue that testimonial knowledge comes with distinctive epistemic rights which sets it apart from other forms of knowledge. These rights pertain not to our acquisition of testimonial beliefs but rather our retention of them in the face of challenges. As Goldberg (2006) emphasises this illustrates a sense in which whatever one says about testimonial justification (that is, whether one is a reductionist or an anti-reductionist about testimonial justification) testimonial knowledge is epistemically distinctive.

Both Goldberg (2006, 2011) and McMyler (2013), when discussing what is distinctive about testimonial knowledge, have emphasised the fact that audiences have the right to 'pass the buck' in response to challenges. That is, when one forms a testimonial belief on the basis of someone else's say so one has the right, when challenged to retain one's belief and defer to the original testifier. For example, imagine that Sammy tells Lizzie that oats lower cholesterol, and Lizzie then tells Mark that oats lower cholesterol. Mark is sceptical and challenges Lizzie's assertion. At first Lizzie says something vague about low density fibre and bacteria in the intestines, but Mark is not convinced and continues to challenge Lizzie's assertion. At this point Lizzie is able to retain her belief but defer to Sammy, saying 'Well Sammy told me, ask her about it'. That is, Lizzie is able to pass on the responsibility for defending her belief to Sammy, since Sammy was the original testifier. The right to pass the buck in this way is a downstream epistemic right which agents acquire through forming testimonial beliefs. Other belief forming methods (such as perception) do not generate such a right. Thus, the fact that such rights are acquired seems to be a distinctive epistemic feature of testimonial beliefs. Goldberg puts the point as follows:

My main contention is that testimonial knowledge is a distinctive kind of knowledge in that this sort of knowledge, but no other, is associated with a characteristic expansion in the sorts of epistemically relevant moves that can be made by the subject in her attempt to identify the direct epistemic support enjoyed by her belief. (Goldberg 2006: 133-134)

...this feature of testimonial knowledge reflects the fact that there is something epistemologically distinctive about relying on the epistemic authority of another rational being: it is because of what is distinctive in relying on the epistemic authority of another rational being, that there is a characteristic expansion in the sort of moves that can be made in defence of a belief acquired on such authority. The characteristic expansion, I suggest, is that testimonial knowledge gives rise to the hearer's right to *pass the epistemic buck* after her own justificatory resources have been exhausted. (ibid: 134)

As emphasised by Goldberg (2006), one can recognise this no matter what one's view of the original justification for testimonial beliefs. This is because the rights pertain not to the conditions under which it is acceptable to form a testimonial belief, but rather the moves one can make when those beliefs are challenged. Moreover, Goldberg (2011) later argues that these distinctive epistemic rights derive from rather uncontroversial features of testimony and assertion. Assuming that assertion has an epistemic norm (that is, assuming that the propriety of an assertion that p requires that the speaker be in some way positively epistemically situated with respect to p) the act of asserting creates mutual knowledge amongst the speaker and audience that the speaker has performed an action the propriety of which requires that they are epistemically well situated with respect to p . Thus, when challenged the audience knows that there is a further body of information (that possessed by the speaker) to which they can appeal in defence of their belief¹⁰. As emphasised by

¹⁰ However, it is unclear to me how these facts alone generate any form of duty in the speaker to defend the audience's belief, and thus where the audience's right of deferral derives from.

both Goldberg and McMyler these rights are genuinely epistemic since they pertain to the ways in which agents are answerable for their beliefs (that is, the epistemic responsibilities they bear with respect to those beliefs). These are taken to be genuinely epistemic features of testimonial knowledge which are not shared by other forms of knowledge. McMyler provides a succinct statement of his reasoning which is worth quoting in full:

..the responsibilities involved in epistemic buck passing are genuinely epistemic responsibilities, responsibilities that pertain to the way in which cognitive agents are distinctively answerable for their beliefs. Belief, I take it, is a commitment-constituted attitude. To believe that p is to commit oneself to a positive answer to the question whether p . A believer is thus answerable for being so committed. She is open to criticism that bears on the content of her commitment, criticism that bears on the question whether p . Plausibly, one aspect of the way in which subjects are thus answerable for their beliefs concerns their epistemic conduct in the face of reasonable challenges to their beliefs, where reasonable challenges to their beliefs involve the presentation of evidence that counts against their beliefs. Typically, when confronted with such a challenge, a rational epistemic agent ought to either find some way to meet the challenge—some basis upon which to rationally discount the evidence presented—or else give up her belief. When it comes to beliefs that are based on being told something by a speaker, however, an epistemic agent is entitled to maintain her belief without meeting the challenge herself by instead passing the epistemic buck back to the speaker. If an audience comes to believe that p on the basis of a speaker's telling, and if a third party challenges the audience's belief by producing evidence that counts against p , the audience is entitled to defer responsibility for meeting the challenge back to the original speaker, whereupon the original speaker is epistemically responsible for meeting the challenge. (McMyler 2013: 1067-1068)

I take these considerations to establish that, at least in some cases, testimony generates the right to pass the epistemic buck, and that this is a genuine epistemic right. In such cases the believer will be less responsible for their belief than they would otherwise be. Moreover, this right to defer to the someone else in defending one's belief illustrates a sense in which such beliefs, and the rights we have regarding them, are distinctly social. This brings us to another sense in which the deniability problem seems genuinely problematic. The deniability problem undermines these speaker commitments. This is particularly clear on Goldberg's (2011) statement of the view which rests on there being mutual knowledge between the speaker and hearer that the speaker has performed an action the propriety of which requires that they are epistemically well situated to a particular proposition p . When the mismatch move is available the speaker is able to provide a defeater to this mutual knowledge claim, meaning that they can back out of their commitments. If speakers undertake no public obligation to defend the audience's belief then audiences don't gain the right to defer to the speaker when challenged. Thus, when the mismatch move is available audiences will form testimonial beliefs but will not acquire the distinctive epistemic rights which standardly accompany such beliefs. If there are discourses which are particularly susceptible to the deniability problem then audiences will bear more individual responsibility for the beliefs acquired from such discourse than they think they do (which could lead to overly credulous belief forming).

So far I have provided two reasons for worrying about the deniability problem. Before continuing to consider the scope of the deniability problem for assertions I will provide one further consideration to establish that it is genuinely problematic.

3.3. Plausible Deniability and Testimony Policing.

So far I have outlined two reasons to worry about the deniability problem. First, it undermines testimonial knowledge on telling based views. Secondly, when the mismatch move is available audiences will not gain the epistemic rights distinctive of testimonial beliefs, meaning that they will bear more responsibility for their beliefs than they think they do. There is one further reason to worry about the deniability problem. Usually when someone falsely asserts that p we can criticise their assertion, and criticise them as an asserter. People who are caught asserting falsehoods are publicly labelled as liars. This is a strong normative assessment, and the chance of being so labelled provides a disincentive against asserting falsehoods. However, when the mismatch move is available it is far harder to make such an accusation. The speaker can always respond claiming that they never intended to communicate the falsehood in question. Thus, if there are discourses which are particularly susceptible to the deniability problem speakers in such discourses will have far less incentive to speak honestly (Goldberg (2013) presses a similar line in order to explain why we should be sceptical of anonymous internet testimony). Of course, we do have other normative assessments available. We can accuse a speaker of being misleading. However, this accusation carries far less normative force, especially if we cannot establish that the speaker has been intentionally misleading. And to establish such a thing will usually require a great deal of information about, for example, the speaker's motivations and their knowledge of their audience.

This concludes the first half of the paper. I have outlined the deniability problem, argued that there is no principled reason to hold that it applies only to implicatures, and provided three reasons to consider it genuinely problematic. Firstly, it undermines testimonial knowledge if one embraces a telling based view of testimony. Secondly, it undermines the distinctive epistemic rights typically acquired via testimonial belief formation. And finally it removes the disincentive for speakers to be intentionally misleading with their assertions. In the second half of the paper my task is to identify

the scope of the deniability problem for assertions.

4. THE CONTEXTUALIST PUZZLE.

So far I have outlined the deniability problem, arguing that there is no principled reason to hold that it applies only to implicatures, and I have provided three reasons to consider it genuinely problematic. What I have not done is argued that a significant number of assertions actually suffer the deniability problem. It is to that task I now turn. I start by considering a puzzle arising from the recent debate over context sensitivity in natural language, which seemingly suggests that the deniability problem is extremely widespread (far more widespread than we would intuitively think). I then discuss some responses to the puzzle which allow us to narrow the scope of the problem back down. We end up with a loose set of criteria for identifying discourses which will be particularly susceptible to the problem.

The first part of the puzzle is the seemingly widespread context sensitivity of natural language. The list of context sensitive uses of language includes indexicals, demonstratives, gradable adjectives, comparative adjectives, definite descriptions, indefinite descriptions, adverbs, conditionals, modals, quantifiers, predicates of personal taste, possessives, incomplete adjectives, psychological attributions, moral attributions, perhaps knowledge ascriptions, non-sentential assertion, vagueness, metaphor, hyperbole, and loose talk. This constitutes a large portion of our language use. Some theorists (radical contextualists) go even further, holding that most, perhaps even all, of our language use is context sensitive. Such theorists note, for example, that seemingly context insensitive terms can be used in many different (and incompatible) ways in different contexts (see,

for example, Travis (1985), Bezuidenhout (2002)). They also provide general theories of utterance comprehension which entail that we always engage in pragmatic processing to recover what is said. On such views, even when the recovered proposition is the proposition determined by the literal meaning of the terms used, such processing is employed to determine that it is the literal meaning (if there is such a thing (see Recanati (2004))) rather than some alternative 'modulation' which is most appropriate.

So, it seems that there is a lot of context sensitivity in natural language, especially if you believe the radical contextualist. This fact by itself shouldn't worry us. What is worrying is that the mechanisms by which we resolve context sensitivity arguably have the same 'ambiguous epistemics' that Fricker argues give rise to the deniability problem for implicatures. Cappelen and Lepore (2004) argue for the radical view that semantic context sensitivity in natural language is restricted to personal pronouns, demonstratives, temporal location adverbs, and the adjectives 'actual' and 'present'. One of their main arguments against widespread semantic context sensitivity is that the mechanisms which contextualists postulate to explain context sensitive communication would seemingly make communication miraculous. They quote the following list of factors to which audiences appeal in resolving context sensitivity:

- “(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)

Cappelen and Lepore point out that if we have to rely on all this information just to recover the content of simple context sensitive utterances such as 'philosophy is fun' then it seems miraculous that we generally tend to communicate smoothly and successfully with context sensitive terms. The fact that we have to rely on such knowledge means that there are many ways the recovery of a given proposition can go wrong. This, as we saw, was precisely what generated the deniability problem for implicatures. There were many ways the recovery of an implicature could go wrong, thus the speaker could easily claim that something did indeed go wrong, disclaiming responsibility for the audience's belief. To see that we do extensively rely on knowledge context to recover what is said consider the two leading approaches to context sensitivity.

Saturation: For many sentences the phonetically articulated elements don't exhaust the syntactic structure, there are non-phonetically articulated syntactic elements which need to be assigned values in order for a complete proposition to be expressed. This is usually taken to consist in there being hidden variables in the underlying logical form which take particular types of value¹¹. Some variables will be bound by linguistic material from earlier on in the same sentence, meaning that their values will be easily recoverable. Indeed, one of the

¹¹ Not all theorists take saturation to be mandated by variables in the underlying logical form. Some theorists deny the existence of such variables, yet maintain that utterances of the sentences in question fail to express full propositions, arguing that audiences supply unarticulated constituents in order to fill in the gaps by appeal to the knowledge contexts (Carston (2002a), Hall (2008), Perry (2001), Recanati (2004)).

primary reasons for positing hidden variables in clauses such as 'it is raining' is that it looks as if there are longer sentences in which they have a bound reading. However, when such sentences are not embedded the audience must employ knowledge context in order to work out the intended value of the variable (Stanley (2000) pushes this line of reasoning).

Often the values which need to be assigned are complex and fine grained. For example, they include properties to restrict quantifiers, and comparison classes for gradable adjectives. The subtle differences in possible values, and the sensitivity of these values to small changes in the mutual goals and presuppositions of the speaker and hearer can, in many cases, give the imaginative speaker scope to make the mismatch move. Take, for example, quantifiers - in many cases one can claim the intended range of a given quantifier to be a restriction on the range the audience attributes. You merely need to be able to identify a subset of the domain the audience attributes such that the members of that subset have some distinguishing feature, and identify an aim relative to which this feature would be relevant such that you could, with at least some level of plausibility, have taken it to be mutually presupposed that it was a conversational aim. For example, in the quantifier case in section two the distinguishing feature of the selected subset was that it contained only vegan friendly beers, and the mutual goal was to identify beers which Sally would be able to drink. Speakers are frequently slippery with what they say in precisely this way.

Modulation: In cases of modulation a constituent has its meaning adjusted, and thus contributes something new to the truth conditions of the utterance. The concept we end up with will be related to the concept encoded in the constituent before modulation takes place, but will usually serve the speaker's purposes better. For example, a concept may be narrowed and thereby assigned a new meaning which applies to a subset of the original extension (e.g. when Sally, who has just left her timid and underachieving boyfriend says 'I

need a man', the concept encoded by 'man' is narrowed in order to apply only to men with particular features commonly associated with masculinity (Carston (2002a)). A concept may also be loosened to generate a new concept which applies to more things than the original concept. Modulation occurs in response to the audience's perception of the demands of the context. It is controversial how the process of modulation works. Some (e.g. Recanati (2004)) think that potential meanings vie for cognitive activation with the most salient meaning being assigned. Others (e.g. Carston (2002a, 2002b), Carson and Wilson (2007)) think that potential meanings are ranked in order of salience and assessed for relevance (a balance of cognitive effects and cognitive effort) until an expectation of optimal relevance is met, at which point the meaning which meets the expectation is assigned.

It should be clear from this that modulation also relies heavily on knowledge context, that the type of context sensitivity it accounts for is common (consider how mundane most of the examples were), and that there will often be many ways an audience can go wrong in recovering the correct proposition. When discussing saturation it was noted that in cases where the values assigned are complex the precise values assigned will be very dependent on knowledge context, thus creating a lot of scope for genuine miscommunication, and the mismatch move. The same is true of enrichment. Indeed, since enrichment is even less constrained it seems the problem will be even worse. Enrichment occurs only in response to the audience's impression of the demands of the context, and the values involved are less constrained than in cases of saturation, giving even further scope for the speaker to claim mismatch. Additionally, speakers can claim to have intended to be interpreted more or less literally than they were (an option which is not obviously available in cases of saturation). All in all, it seems that the problems which arise for saturation not only arise, but are multiplied in cases of enrichment¹².

¹² I think that many contextualists (especially those who focus on modulation and free enrichment) recognise the fact that their theories involve the sorts of 'problematic epistemics' pointed to by Fricker. Many of these theorists have

We now have the two components of the puzzle generated by contextualism. Firstly there is reason to believe that there is a great deal of context sensitivity in natural language, secondly there is reason to think that this context sensitivity has the same problematic epistemics which gave rise to the deniability problem for implicatures. This suggests that the deniability problem is very widespread. This is a radical conclusion, and one we should be eager to avoid, since it seems quite obvious that speakers don't have plausible deniability about what they say the majority of the time. I will now consider a series of responses to Cappelen and Lepore's challenge which allow us to narrow down the scope of the problem.

5. RESPONDING TO THE CONTEXTUALIST PUZZLE.

I have outlined the deniability problem and argued that we have no principled reason to hold that it applies only to implicatures. Next I provided three reasons to consider it genuinely problematic. I then presented a puzzle which seems to suggest that the deniability problem may be very widespread. Context sensitivity seems to be very widespread in verbal communication, and such communication seems to rely heavily on knowledge context. Cappelen and Lepore have argued that

worried about related problems which their own theories raise for communicative success. This has led many prominent theorists (such as Bezuidenhout (1997), Carston (2002), Heck (2002), Recanati (2004), and Sperber and Wilson (1986)) to conclude that communicative success does not require the precise sharing of contents between speaker and audience, but rather entertainment of similar propositions. They endorse this view partly as a result of their recognition of the epistemic difficulties which arise in the recovery of propositional contents. If communicative success required that speakers and audiences shared identical contents then communicative success would be rare due to the epistemic difficulties involved in recovering a propositional content identical to that intended by the speaker (many of these theorists also raise worries about the sharing of Fregean contents).

heavy reliance on knowledge context gives a lot of scope for miscommunication, making successful communication seemingly miraculous. Thus, where there is such heavy reliance on knowledge context it may appear that there is a great deal of scope for speakers to employ the mismatch move. In this section I consider a series of responses to Cappelen and Lepore¹³. I will argue that these responses fail as general responses to the deniability problem. Nonetheless, they are worth considering because they illustrate a sense in which it may be harder for speakers to make the mismatch move in certain contexts. We will see that a discourse must have certain features in order to block plausible deniability for context sensitive assertions. Thus we will be able to identify discourses which are particularly susceptible to the deniability problem.

Ishani Maitra (2007) has argued that certain contextual values are more natural and frequently applied than others, and that we will assign such values unless we have reason not to. For example, if I say 'elephants are big' it will be more natural for you to assign the comparison class of species rather than some other comparison class, for example the class of large mammals. You would generally only assess the claim relative to the class of large mammals in response to additional information which made that reading more likely. If this is the case then speakers and hearers will usually converge on the same contents because audiences will not need to appeal to knowledge context to a problematic extent.

I think it should be clear that this response won't solve our problem. Even if it does help to explain how interlocutors converge on contents it still leaves a lot of scope for speakers to detach themselves from the proposition recovered. If the audience settles on the more common or natural reading then

¹³ Cappelen and Lepore's own position is that we must separate semantic content and speech act content, and that speech act content but not semantic content is context sensitive. Many of the same problems arise for their theory as arise for the contextualist, because the proposition we will form a testimonial beliefs in will usually not be the minimal semantic content, but rather the context sensitive speech act content.

the speaker may claim that further information should have been appealed to in interpretation. Likewise, if the audience moves away from the more common reading the speaker can claim that the audience drew too much from context and that the common reading was intended. This is not to say that Maitra's response is without use. Certainly in some contexts, and for certain common conversational tasks, standard meanings for terms emerge. And certainly in some contexts there will be a meaning so obviously more natural than all the others that the speaker has no scope for plausible denial concerning what they meant. But it is far from clear that the majority of cases are like this.

A related but more promising response has been provided by François Recanati (2010). Recanati doesn't postulate a set of natural or common meanings, however he does postulate that there are important psychological commonalities which dispose people to converge on the same meanings of terms, and which enable us to recognize how others intend concepts to be modulated. For example, people are disposed to recognise the same similarities between the situation of application for a concept, and other situations to which the concept does not straightforwardly apply. As a result they are able to extend or narrow the use of term a which expresses that concept, perhaps modulating the concept along these dimensions of similarity to apply to the new case (see also Bolinger (1968)). Our interpretation is also taken to be guided by sets of implicit biases which are common across speakers and audiences. Recanati does not provide any examples of such biases. However he does point to some biases postulated by psychologists working on the early acquisition of lexical meanings. For example, it is argued by Bloom (2000) that early acquisition of lexical meaning is guided by a whole object bias (a bias toward taking a whole object rather than the parts of an object to be the referent of a term). It is conceivable that a network of such common biases guides our ordinary interpretation and helps us assign values to context sensitive terms¹⁴.

¹⁴ I suspect that this response only secures convergence on similar rather than identical propositions. This shouldn't worry Recanati since he maintains that similarity of content is sufficient for communicative success (as do

This does not solve the deniability problem. Firstly, this network of biases and abilities may often lead audiences in the right direction, but (especially if such biases are rooted in theory of mind) they will still be dependent on assumptions about the context to which the speaker can appeal to in order to claim miscommunication. Thus the mismatch move will still be available. Secondly, this response has a rather narrow scope. Although there may be similarities, or types of similarity, which humans as a kind are more disposed to recognise, and biases toward objects which can be categorised in certain ways, these are extremely unlikely to exhaust the range of similarity judgements and psychological mechanisms which guide interpretation. It would be very surprising if we were not also guided by similarity judgements and biases which are moulded by our individual experiences. This seems especially true when we are dealing with more abstract concepts. This creates a greater chance of genuine mismatch, and more scope for speakers to make the mismatch move.

A third response draws our attention to communication as a collaborative affair. It has been argued that conversational participants don't allow situations to arise where there is any realistic chance of content mismatch (Perrini (2009), Recanati (2010)). This is because speakers don't just make an utterance, get interpreted, and move on. Rather, there are collaborative checks in place to ensure understanding. Both speakers and audiences track each other's facial expressions, tone, and body language for signs of misunderstanding or mismatch. Additionally, if there is uncertainty about what was said the audience asks for clarification (Clark and Krych (2001)).

In responding to the deniability problem the thought would be this: both audience and speaker collaborate to establish a shared meaning, and audiences can refer back to this when a speaker attempts the mismatch move, thereby blocking plausible deniability. This response also fails. It is

Bezuidenhout (1997), Carston (2002), and Heck (2002)).

only in cases where either the audience is aware of their lack of understanding, or when their subsequent interactions with the local environment indicate misunderstanding, that mistakes are corrected. When the speaker and hearer are not coordinating on a mutual task involving the local macro level environment the hearer will have nothing to refer to when calling out the speaker. Additionally, many similar but epistemically distinct contextual values (e.g. quantifier domain restrictions) will have very similar behavioural consequences, meaning that misunderstanding won't immediately generate behavioural evidence of miscommunication. So the response is somewhat limited in scope.

So far I have surveyed a series of responses to Cappelen and Lepore and found them lacking as responses to the deniability problem. Nonetheless, I think they can teach us something important. These responses draw our attention to a set of resources to which an audience can appeal in certain contexts in an attempt to call out a speaker who is attempting the mismatch move, thereby blocking plausible deniability. For example, if there is a clear common use for a term (or common default contextual value), and the audience reasonably assigns such a value only to be met by the mismatch move later on, then the audience is able to maintain that the speaker should have been more explicit about their intention, maintaining that they are partly responsible for the resultant belief. The same goes for modulation based on similarity relations and biases. This is especially true in cases where there are checks in place related to some mutual task. If someone acts on the basis of their understanding, and at that point the speaker fails to flag any misunderstanding, then this goes some way to confirming the audience's initial interpretation, making it far harder for the speaker to make the mismatch move without descending into absurdity. In general, the more the audience is able to check that they have the correct understanding, and the more obvious the default understandings are, the harder it will be for a speaker to claim mismatch without absurdity.

These resources are rather limited, and in one-off instances they may often prove ineffective. However, if speakers repeatedly try to employ the mismatch move in order to avoid commitment in circumstances where these resources are available then audiences will be able to call them out on their frequent misleading behaviour. It may be plausible that in a one-off case the speaker intended the audience to assign a more esoteric meaning to a term than they did. However, it becomes far less plausible in a long run of cases. Repeat offenders will lose plausible deniability. Moreover, speakers have motive to avoid appealing to the mismatch move in contexts where checks are in place, for if they make the move frequently then they will quickly lose credibility as an informant.

The deniability problem now seems somewhat less worrying. However, it was only weakened for discourses where audiences have the resources to call out the speaker by appeal to standard meanings, very obvious ways of extending a meaning, or checks which serve to reliably confirm understanding. It is a partially empirical question how many discourses actually have these features. However, I think it is likely that some important discourses lack them. The deniability problem still arises with its full force for such discourses. These will be discourses in which context sensitivity (especially more unconstrained context sensitivity such as modulation) is rife, which don't involve coordination on macro level tasks, where the values or modulated concepts are complex or abstract, and where there are no highly standardised or clearly stated contextual values. Such discourses provide speakers with a lot of scope to make the mismatch move without losing much credibility. There will be more resources to which speakers can appeal in order to claim misunderstanding, and fewer checks an audience can appeal to in order to call them out or hold them responsible. It is not the task of this paper to establish conclusively that any particular discourse has such features, however several important discourses do seem to be candidates. For example, religious discourse is arguably rife with context sensitivity (consider the many and varied religious conceptions of salvation, love, and even God (see Alston (2005), and Scott (2005) for useful overviews discussing

the context sensitivity of religious language), it has a highly abstract subject matter, and it is not clear that there are always sufficient efforts put in place to coordinate on precise explicit meanings. Likewise much public political discourse seemingly has such features (for example, consider the possible varied meanings of terms like 'class warfare'). Indeed, this no doubt adds to the stereotype of politicians as slippery and dishonest. Another important candidate seems to be ethical discourse outside of academic settings (where there are often norms which require precision and coordination on standard meanings). Such discourse is abstract, arguably context sensitive, and seemingly lacks a norm requiring explicit joint efforts to coordinate on precise meanings. Of course, it is beyond the scope of this paper to establish that any of these discourses do have the features in question. However, given the importance of such discourses in our everyday lives this seems like a worthy question for further research.

6. CONCLUSION.

I have outlined the deniability problem for assertions, explained why it is genuinely problematic, and presented a line of reasoning which seems to suggest that the problem is very widespread. I then looked at several ways of narrowing the scope of the problem back down. It was found that the deniability problem would be less problematic in discourses with certain features. I suggested a selection of important discourses which may still face the problem.

It is not clear how we should react once we discover that a discourse faces the deniability problem. I suspect that the correct reaction will vary between different discourses. One reaction may be to try and establish a set of precise meanings within the discourse, and eliminate context sensitivity as

much as possible (for example, if certain areas of academic discourse were found to face the deniability problem, this would probably be the more appropriate response). An alternative response would be to give less weight to testimony in the problematic discourse. One could treat knowledge regarding the subject matter as necessarily personal rather than social. One could reconceptualise the role of apparent testimony in the discourse, perhaps taking it to be expressive, or seeing it as intended not to bring about belief but rather reflection or some other attitude. And another alternative would just be to view the discourse with scepticism. The plausibility of any given response will depend on the discourse in question. Finally, one might simply choose to weaken the focus on speaker commitments in one's theory of testimony¹⁵.

REFERENCES.

- Alston, W.** 2005. 'Religious Language'. From J. Wainwright, (ed), *The Oxford Handbook of Philosophy of Religion*, pp. 220-242. Oxford university Press.
- Bezuidenhout, A.** 1997. 'The Communication of De Re Thoughts'. *Nous*, 31:197-225.
- 2002. 'Truth Conditional Pragmatics.' *Philosophical Perspectives*, 16: 105-134.
- Bloom, P.** 2000. *How Children Learn the Meaning of Words*. Cambridge Mass: MIT Press/Bradford Books.
- Bolinger, D.** 1968. *Aspects of Language*. New York: Harcourt, Brace and World.
- Cappelen, H. and Lepore E.** 2004. *Insensitive Semantics: A Defence of Semantic Minimalism and Speech Act Pluralism*. Blackwell Publishers.

¹⁵ Many thanks to Herman Cappelen, an anonymous reviewer for this journal, and participants in the the Arché Evidence, Justification, and Knowledge seminar for comments which greatly improved the quality of this paper. This research was supported by the United Kingdom Arts and Humanities Research Council.

- Carston, R.** 2002a. *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Wiley Blackwell.
- Carston, R.** 2002b. 'Metaphor, Ad Hoc Concepts, and Word Meaning: More Questions than Answers.' *UCL Working Papers in Linguistics*, 14: 83-105.
- Carston, R. and Wilson, D.** 2007. 'A Unitary Approach to Lexical Pragmatics: Relevance, Inference, and Ad Hoc Concepts.' In N. Burton-Roberts (ed), *Pragmatics*, pp. 230-259. Basingstoke and New York Palgrave Macmillan.
- Clark, H. and Krych, A.** 2001. 'Speaking While Monitoring Addressees for Understanding.' *Journal of Memory and Language*, 50: 62-81.
- Fricker, E.** 1994. 'Against Gullibility.' From B.K. Matilal and A. Chakrabarti (eds), *Knowing From Words*, pp 125-163. Kluwer Academic Publishers.
- 1995. 'Telling and Trusting - Reductionism and Anti-Reductionism in the Epistemology of Testimony.' *Mind*, 104: 393-411.
- 2006a. 'Second Hand Knowledge.' *Philosophy and Phenomenological Research*, 73: 592-618.
- 2006b. 'Varieties of Anti-Reductionism About Testimony: A Response to Goldberg and Henderson.' *Philosophy and Phenomenological Research*, 72: 618-628.
- 2012. 'Stating and Insinuating.' *Proceedings of the Aristotelian Society - Supplementary Volume*, 86: 61-95.
- Goldberg, S.** 2006. 'Reductionism and the Distinctiveness of Testimonial Knowledge.' From J. Lackey and E. Sosa (eds), *The Epistemology of Testimony*, pp. 127-144. Oxford University Press.
- 2011. 'Putting the Norm of Assertion to Work: The case of Testimony'. From H. Cappelen and J. Brown (eds), *Assertion: New Philosophical Essays*, pp.175-197. Oxford University Press.
- 2013. 'Anonymous Assertions.' *Episteme*, 10: 135-151.
- Hall, A.** 2008. 'Free Enrichment or Hidden Indexicals.' *Mind and Language*, 23 : 426-456.
- Heck, R.** 2002. 'Do Demonstratives have Senses?' *Philosopher's Imprint*, 2: 1-33.

- Hinchman, E.** 2005. 'Telling as Inviting to Trust.' *Philosophy and Phenomenological Research*, 70: 562-587.
- Keren, A.** 2012. 'On the Alleged Perversity of the Evidential View of Testimony.' *Analysis*, 72: 700-707.
- MacFarlane, J.** 2005. 'Making Sense of Relative Truth.' *Proceedings of the Aristotelian Society*, 105: 321-339.
- 2011. 'What is Assertion?.' From J. Brown and H. Cappelen (eds), *Assertion*, pp. 79-96. Oxford University Press.
- Maitra, I.** 2007. 'How and Why to be a Moderate Contextualist.' From G. Preyer and G. Peter (eds), *Context Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*, pp. 112-133. Oxford University Press.
- McMyler, B.** 2013. 'The Epistemic Significance of Address.' *Synthese*, 190: 1059-1078.
- Moran, R.** 2005a. 'Problems of Sincerity.' *Proceedings of the Aristotelian Society*, 105: 341-361.
- 2005b. 'Getting Told and Being Believed.' *Philosopher's Imprint*, 5: 1-29
- Perrini, E.** 2009. 'Does Contextualism Make Communication a Miracle?' *Manuscripto*, 32: 231-247.
- Perry, J.** 2001. *Reference and Reflexivity*. Stanford: CSLI Publications.
- Recanati, F.** 2004. *Literal Meaning*. Cambridge University Press.
- 2010. *Truth Conditional Pragmatics*. Oxford University Press.
- Rescorla, M.** 2009. 'Assertion and its Constitutive Norms.' *Philosophy and Phenomenological Research*, 79: 98-130.
- Ross, A.** 1986. 'Why Do We Believe What We Are Told.' *Ratio*, 28: 69-88.
- Scott, M.** 2010. 'Religious Language.' *Philosophy Compass*, 5: 505-515.
- Sperber, D. and Wilson, D.** 1986. *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge, MA: Harvard University Press.

Stanley, J. 2000. 'Context and Logical Form.' *Linguistics and Philosophy*, 23: 391-434.

Travis, C. 1985. 'On What is Strictly Speaking True.' *Canadian Journal of Philosophy*, 15: 187-229.