

This is a penultimate draft of a paper forthcoming in Synthese. Please cite the final version.

Epistemic Injustice in Utterance Interpretation¹.

1. Introduction.

There has been much recent discussion of the harmful role prejudicial stereotypes play in our communicative exchanges. For example, Miranda Fricker (2007) explores a type of injustice (testimonial injustice) which arises when the credibility judgements we make about speakers are informed by prejudicial stereotypes. One might, according to Fricker, wrong someone by assigning them a low credibility due to their race or gender. This discussion has produced many important insights. However, it has so far focused on the role stereotypes play in our epistemic assessments of communicative actions, rather than our interpretations of such actions (interpretations which underlay our epistemic assessments)². Yet, the same prejudicial stereotypes that infect credibility judgements can also infect our interpretation. For example, imagine a black man wrongly suspected of rape in depression era Alabama. Suppose that he is overheard saying 'she seemed vulnerable' whilst describing his motivations for being at his accuser's home. Given the context and the prejudices of the time (with black men being seen as predatory toward white women), he would likely be misinterpreted as stating that he saw her as an easy target. Call this 'Interpretative Injustice'. This paper explores the sources of interpretative injustice, and considers some of the harms to which it gives rise. There are several harms caused by interpretative injustice. Firstly, it silences. It prevents certain groups from being able to efficiently communicate knowledge to other (perhaps more powerful) groups. Secondly, it results in speakers being held epistemically responsible for

1 This paper has benefited greatly from comments and discussions with Sebastian Becker, Herman Cappelen, Josh Habgood-Coote, Elizabeth Fricker, Patrick Greenough, Elizabeth Marr, Matthew McKeever, Leonard Randall, Caroline Touborg, Brian Weatherston, and two anonymous referees for this journal. I would also like audiences at the St Andrews Friday Seminar, and the Arché work in progress seminar where this work was presented. This research was supported by the United Kingdom Arts and Humanities Research Council, and a Royal Institute of Philosophy Bursary.

2 This is not to say that the topic of miscommunication in general has been ignored in this literature. For example, José Medina (2013) urges a focus on the dynamics and mechanisms of communication, but focuses on the larger scale dynamics of intergroup communication, as well as the differing and dynamic availability of hermeneutical resources within and between different social groups. Rather, it is the role of prejudicial stereotypes in generating miscommunication which has been ignored. My interest here lies with smaller scale interactions, and the cognitive mechanisms underlying particular instances of communication.

propositions they never intended to communicate. And thirdly, it contributes to the illusion that prejudicial low credibility judgements are epistemically justified. I close by arguing that if Miranda Fricker's strategy for treating testimonial injustice is implemented in absence of a treatment of interpretative injustice then we risk epistemically harming hearers with little benefit to speakers. Thus testimonial injustice and interpretative injustice are best treated in tandem.

2. Interpretative Injustice.

Fricker's brand of epistemic injustice arises as a result of the way in which our credibility judgements are guided by potentially prejudicial stereotypes. Credibility judgements occur late in the process of testimonial belief formation. Other stages in this process are likewise guided by stereotypes. In particular, our interpretation of utterances themselves will often be guided by stereotypes³. This is especially true in cases involving context sensitivity, loose talk, unfamiliar dialects or accents, noisy environments, and implicature. In such situations we must appeal to our knowledge of the context, including what we know about the speaker (their likely goals, interests, beliefs, background, and intelligence) in order to reach a verdict on what proposition they are intending to communicate.

Interpretative injustice is the phenomenon whereby a hearer's employment of prejudicial stereotypes results in the hearer attributing a message to the speaker when the speaker never intended to convey that message. As we will see, these prejudicial stereotypes may affect both speech perception and the assignment of meaning. Thus, interpretative injustice is a broad phenomenon. I mostly follow Fricker's use of the term 'prejudicial stereotype' to mean, roughly, a generalisation which embodies a judgement about a social group which is not properly evidence responsive⁴. That is, one harbours a prejudicial stereotype if one harbours a (usually negative) generalisation (or set of associations which embody a generalisation) about a particular social group, and one's harbouring of

3 My focus here will be on our understanding of the content of utterances. However, our understanding of the types of act performed can also be shaped by misleading stereotypes. For example, we may mistakenly interpret an order as a request. Such cases are discussed in Kukla (forthcoming).

4 The processes underlying our understanding and credibility judgements are likely to be associative. However, we can think of sets of associations as embodying generalisations. For example, the association of 'black' with 'crime', 'gun', and 'drugs' would embody a generalisation about black people e.g. 'black people are criminals'.

that generalisation is not sensitive to the available evidence. If an aspect of our social practice systematically, unfairly, and disproportionately harms certain social groups, then in cases where members of these groups have been harmed by the social practice they have been victims of injustice. Interpretative injustice gives rise to numerous harms (discussed in section five). These harms are experienced disproportionately by particular disadvantaged social groups. Moreover, in cases where these harms are experienced as a result of interpretative injustice it will be the hearer's lack of evidence responsiveness which gives rise to the harm. Thus it seems that the harm is unfairly inflicted upon the speaker (the miscommunication does not result from either a fault in the speaker or mere bad luck, it arises out of a fault in the hearer⁵). Moreover, these harms will constitute part of a wider web of harms which the disadvantaged individual is systematically subjected to.

For the sake of clarity, it is worth considering some related concepts and cases to see what does and does not fall under the umbrella of interpretative injustice. Firstly, I do not consider interpretative injustice, as discussed here, to include cases where a hearer reaches the correct judgement about what a speaker intended to communicate, but does so on the basis of a misleading stereotype. Such cases seem to be instances of lucky interpretation, and may be interesting in their own right. However, I will not be discussing them further, as there is no misinterpretation and the audience is not harmed.

Also of potential interest is the more general case where someone misinterprets a communicative act on the basis of a false belief or presupposition. Interpretative injustice is a subspecies of this phenomenon, and many of the things I say about the harms of interpretative injustice will carry over. However, I am concerned primarily with misinterpretation caused by prejudicial stereotypes or associations, because this is a more systematic phenomenon and patterns with other forms of injustice.

What about cases in which the speaker may not have been understood even by an unprejudiced audience? One might think that in such cases the hearer's prejudice does not harm the speaker, meaning that no injustice occurs. This would be a mistake.

⁵ This is not to say that audiences in cases of interpretative injustice are always completely free of responsibility for the miscommunication. An already unclear speaker whose chances of being understood are further undermined by the audience's prejudice is still at an unfair communicative disadvantage as a result of the audience's prejudice.

Consider the following case: A person of colour tries to subtly communicate p , and does so in such a way that it might be unclear even to a reasonable audience that p was intended. However, the audience considers persons of colour to be incapable of subtlety, and so straightforwardly interprets the speaker as asserting q . In this case although the speaker's chances of being understood were already reasonably low the audience's prejudices lower the chances even further, in a way which is out of the speaker's control. This is analogous to throwing away someone's lottery ticket on the basis of their race. It is unlikely that they would have won even if the ticket had not been thrown away, however their chances of winning are still substantially and unfairly diminished. Essentially, the act of subtle communication is ruled out (or made far more difficult) for the person of colour. One can imagine similar cases in which one is unable to have humour recognised as humour, or in which artistic divergences from ordinary language in one's poetry are treated as mere linguistic incompetence. In all such cases one is restricted in one's ability to use language in a particular way due to the hearer's prejudices⁶.

It is also worth considering the relationship between interpretative injustice and silencing at this point (we will return to the topic of silencing in section five). Many cases of interpretative injustice are cases of silencing, and many of the harms of interpretative injustice arise as a result of its ability to silence. Silencing occurs when an individual or group is prevented from carrying out a communicative action, either by being prevented from attempting the action (for example, through intimidation) or as a result of the action being rendered unsuccessful. Hornsby and Langton (1998) characterise the latter forms of silencing in terms of the illocutionary or perlocutionary force of the act being blocked. That is, either the communicative act is prevented from producing its intended outcome (perlocutionary silencing) or it is prevented from even constituting the intended type of action (illocutionary silencing). Hornsby and Langton maintain that in order for an illocutionary act (for example, an assertion or a refusal) to be carried out the fact that it is intended to be an assertion or refusal must be recognised by the audience. So when the audience fails to recognise the illocutionary act the speaker is attempting the speaker is illocutionarily silenced. More recently Ishani Maitra⁷ (2009) has characterised silencing

6 Indeed, the ability of certain groups to use language creatively by coining new words may also be interpreted as incompetence. Here we see interpretative injustice contributing to what Fricker calls 'hermeneutical injustice'.

7 Maitra is not aiming to capture cases of silencing where the speaker is prevented from attempting the

in Gricean terms. Grice's conditions on speaker meaning are as follows:

'A speaker *S* means something by uttering *x* iff, for some audience *A*, *S* utters *x* intending:

- (i) *A* to produce a response *r*;
- (ii) *A* to think (recognise) that *S* intends (i); and,
- (iii) *A*'s fulfilment of (ii) to give him a reason to fulfil (i)'

Grice (1989): 92.

Maitra then characterises silencing as follows:

'In my view, a speaker is communicatively disabled iff she is unable to fully successfully perform her intended communicative act, because her intended audience fails to satisfy either the second or the third of her (Gricean) intentions.'

Maitra (2009): 327-328

Most cases of interpretative injustice are cases of silencing by these definitions, since if one misinterprets the content of an utterance then one thereby fails to recognise the illocutionary act being performed, or the speaker's communicative intention⁸. Indeed the paradigmatic case of ethically problematic silencing is arguably a case of interpretative injustice. In the paradigmatic case of silencing a woman tries to refuse a man's sexual advances by saying 'no'. However, as a result of his misleading conception of women as wanting to avoid the appearance of promiscuity, yet usually desiring sex, he misinterprets the refusal. He thereby fails to recognise her communicative intention. Thus the woman is prevented from performing the illocutionary act of refusal. In this case the man's interpretation was guided by a misrepresentation of the speaker's interests, desires, and intentions based on a harmful stereotype of women. To use Kristie Dotson's (2011) phrase, this would be an 'instance' of silencing. Dotson draws a distinction between instances of silencing and practices of silencing:

8 These are all forms of what José Medina (2013) calls communicative forms of silencing. They prevent communication from occurring (or being reciprocated). Medina contrasts communicative and epistemic forms of silencing, where epistemic forms of silencing arise when communication occurs, but the audience fails to treat the speaker as a knower (cases of testimonial injustice in Fricker's sense are cases of epistemic silencing). Medina argues that the two notions are complimentary, and I agree. Indeed, the final section of this paper discusses the interaction between interpretative injustice (which gives rise to a form of communicative silencing), and Fricker's treatment of testimonial injustice (a form of epistemic silencing).

'An instance of silencing concerns a single, non-repetitive instance of an audience failing to meet the dependencies of a speaker, whereas a practice of silencing, on my account, concerns a repetitive, reliable occurrence of an audience failing to meet the dependencies of a speaker that finds its origin in a more pervasive ignorance' Dotson, (2011): 241.

It appears that, since interpretative injustice silences, and arises systematically as a result of shared social stereotypes and the mechanisms of linguistic interpretation, interpretative injustice might also be thought to produce a practice of silencing⁹. It is not clear that all cases of interpretative injustice are instances of silencing however, since silencing is usually characterised in terms of the intended audience failing to properly reciprocate, and interpretative injustice can occur in cases where the intended audience fully reciprocates, but an eavesdropper misinterprets the utterance with harmful consequences¹⁰. As we shall later see, such misinterpretation can still be harmful (as there are several harms of interpretative injustice which do not derive from its ability to silence).

Since most cases of interpretative injustice seem to be cases of silencing¹¹ it may be unclear why interpretative injustice is a worthy topic of investigation in its own right. I think there are several reasons why it is important to study interpretative injustice as an independent phenomenon. Firstly, as we will see in section five, interpretative injustice is associated with its own harms, not brought about simply through its ability to silence¹². Secondly, interpretative injustice complicates our response to epistemic injustice. This is

9 It is not clear how reliably the silencing must occur for a 'practice of silencing' in Dotson's sense, to arise. I do not believe that misinterpretation occurs in all, or even most cases in which marginalised speaker's interact with non-marginalised speakers. Rather, my claim will be that marginalised testifiers are subjected to systematic and disproportionate misinterpretation from a common source, and that this is, in itself, harmful.

10 It is plausible that in some such cases, where the speaker realises that there is a risk of misinterpretation by someone other than the intended audience, silencing still arises as a result of self-censorship. That is, the speaker may refrain from making an utterance for fear of being misinterpreted by an eavesdropper. Dotson (2011) refers to this style of self censorship as 'smothering'. This will be discussed further in section four.

11 Perhaps all cases, depending on how the notion of silencing is developed to deal with eavesdropper cases.

12 This is not to say that the harms discussed in section five are never caused by other forms of silencing. Rather, interpretative injustice, by its very nature (and unlike silencing in general), seems particularly strongly (and systematically) associated with these particular harms.

discussed in the final section of the paper. The final (and more general) point is that although silencing is a unified phenomenon, different types of silencing have different causes. In order to be able to find solutions to the problems caused by silencing, and in order to properly understand the way in which different social practices silence, it is important to have a clear taxonomy of the types of silencing not just in terms of their effects (i.e. whether they are perlocutionary, illocutionary etc.), but in terms of their causes. The phenomenon of interpretative injustice, insofar as it is a form of silencing, fits into such a taxonomy because it has a distinctive cause - underlying prejudicial biases shaping the way we interpret people's speech. Other forms of silencing have different causes, for example the same biases altering our credibility assessments, or more overt prejudices causing us to discount, or try to actively prevent the testimony of certain groups. Some types of silencing are unified by, for example, being caused by prejudicial stereotypes. However, these stereotypes can act in different ways on different levels of cognition, giving rise to importantly different forms of silencing. A solution to one type of silencing will not necessarily generalise to all.

Finally, it is worth distinguishing interpretative injustice from Fricker's notion of hermeneutical injustice. Hermeneutical injustice is the phenomenon whereby a subjugated group is unable to render their experience intelligible either to themselves or to others, due to their lacking the concepts with which to do so¹³. Fricker uses the concept of sexual harassment as an example. Until the notion of sexual harassment entered into public discourse women who experienced it were restricted in their ability to communicate the species of wrong they were being subjected to, or even identify it themselves in a clear way (Fricker (2007)). Interpretative injustice does not concern the existence or the availability of the concepts required to render our social experiences intelligible. Rather, interpretative injustice occurs when the wrong content is assigned as a result of prejudicial stereotypes influencing interpretation, regardless of what concepts are available in our public language. The existence of interpretative injustice is therefore

¹³ Fricker's is not the only definition of hermeneutical injustice, Medina (2013) offers the following alternative definition:

'hermeneutical injustice will be treated, roughly, as the kind of injustice that appears when there are wrongful interpretative obstacles that affect people differently in how they are *silenced*, that is, in their inability to express themselves and so to be understood.' Medina (2013): 91.

Interpretative injustice does fall under Medina's definition of hermeneutical injustice, at least in cases in which it silences.

entirely consistent with the absence of hermeneutical injustice.

This is not to say that hermeneutical injustice and interpretative injustice are not importantly related in practice. As Gaile Pohlhaus Jr (2012) argues, instances of hermeneutical injustice are often more complex than they seem on Fricker's model. The availability of concepts throughout a linguistic community is not uniform. A concept might be developed to render an experience intelligible, and this concept may be available to one subset of the population and not others. This might arise through what Pohlhaus Jr calls 'wilful hermeneutical ignorance', whereby privileged groups refuse to draw on the hermeneutical resources developed by marginalised groups in order to render their experiences intelligible, or it may arise through the marginalised group's own efforts. In such cases interpretative injustice is likely to interact with hermeneutical injustice in important ways. When a speaker invokes a concept which is only available to a subset of the community their testimony will likely be misinterpreted by anyone who does not have a grasp on the concept. This misinterpretation arguably has a separate source from interpretative injustice (as discussed here), as such miscommunication (which can still be harmful) could arise without the actual evocation of prejudicial stereotypes in interpretation. However, in cases where it is known that the language or conceptual resources are identified with a particular marginalised group, it seems likely that stereotypes about that group will influence the way in which the communicative practice is interpreted by outsiders. In this case the resultant misinterpretation would arise out of a combination of hermeneutical ignorance and interpretative injustice. Moreover, in many cases of interpretative injustice the conceptual and linguistic resources are at least present for the speaker to correct the audience's misinterpretation. However, in cases where hermeneutical ignorance is in play speakers will generally lack the linguistic resources with which to correct misinterpretations, thus these misinterpretations are more likely to stick.

3. The Sources of Interpretative Injustice.

It is commonplace to think that we judge a context sensitive term to receive a particular value in context, a particular proposition to be implied, or loose talk to be resolved in a particular way, by making judgements about the conversational common ground and mutual salience. That is, we make judgements about what information the speaker is

trying to communicate, and the information to which they expect to appeal, on the basis of our representation of the speaker's representation of our mutual goals, knowledge, interests, intelligence, salience judgements etc. However, it is rare that we explicitly reason about the common ground, or about mutual salience. Rather, like our credibility judgements, our judgements about what is said or implied are unreflective and intuitive. We don't explicitly consider the common ground and work out the speaker's meaning in an effortful step by step manner, rather we rely on quick heuristics and stereotypes in order to quickly make a judgement about the common ground. As in the case of credibility judgements a hearer's stereotypes concerning social identity will have an impact on the judgements reached. For example, a stereotype concerning social identity could alter the audience's representations of the speaker's likely interests or background knowledge.

There is empirical support for the notion that social identity judgements play a role in semantic processing. For example, Van Berkum et al (2008) found that anomalies with respect to the social identity of the speaker (such as 'I have a large *tattoo* on my back', spoken in an upper class accent, or 'if only I looked like *Britney Spears* in her latest video' spoken in a male voice) yielded the same type of neural response as semantic anomalies (such as 'the earth revolves around the *trouble* in a year'). This suggests that speaker identity plays a role in semantic interpretation even at the earliest stages. More recently Boland and Clark (MS) found that in linguistic contexts which promote predictive processing audiences were able to more quickly assign a referent to an ambiguous term when the referent was associated with the social identity of the speaker (for example, fingernails were more quickly identified as the referent of 'nails' when the sentence was spoken in a female voice). Finally, Gieselman and Bellezza (1977) found that audiences judged the same sentence to be more or less 'potent' (a measure of connotative meaning) when spoken by male or female respectively. These studies support the view that amongst the stereotypes upon which we rely to recover content are stereotypes concerning social identity (e.g. class, gender, race etc.).

It is worth considering a specific account of the way in which we assign contextual values in order to clarify precisely how implicit stereotypes are able to play a role. I use relevance theory as an example because it is a clear and relatively well known account of the mechanisms underlying the resolution of context sensitivity. According to relevance

theory we assess interpretations for relevance, and select the first interpretation to meet our expectation of optimal relevance. 'Relevance' is a term of art denoting a feature of inputs to cognitive processes which strikes the greatest balance between relevant cognitive effects (new derivable information, and strengthening or weakening of previous assumptions) and cognitive effort (Sperber and Wilson (1986), Carston (2002)). That is, when we assign a meaning to a term we rank concepts on the basis of cognitive activation and assess each for relevance. The first concept to meet the audience's expectation of relevance is assigned as the meaning of the term. Consider the sentence 'Betty walked to the edge of the cliff and jumped'. Imagine that it is used in a discussion of Betty's suicide. The audience may start by assessing the literal meaning of 'jumped', but the resultant proposition would not meet their expectation of relevance, so they would move on and assess other candidate meanings. The topic of conversation (suicide) will make the notion of 'jumping of a cliff' salient and will thus be amongst the potential meanings assessed for relevance. It would meet the audience's expectation of optimal relevance, and so would be assigned as the meaning of 'jumped'. There are multiple points at which stereotypes and generalisations based on the speaker's social identity can enter into and affect this process. For example, an audience's view of the speaker will affect the level of activation of the various concepts which are ranked in order to be assessed for relevance. That is, if an audience member associates a particular concept or topic with a particular race, gender, or social class then a candidate meaning related to that concept or topic may be assigned a higher initial ranking when the speaker is a member of that race gender, or social class (consider the example discussed in the introduction where the stereotype of predatory black male sexuality affects the interpretation of 'vulnerable'). Likewise, the cognitive effects taken to be relevant will be affected by the audience's expectations regarding the aims, interests, and beliefs of the speaker.

Indeed, such factors don't just guide us in our resolution of context sensitivity, loose talk, and implicature. By observing our informants and making judgements about their beliefs, backgrounds, and interests, we build up a picture of the speaker which also guides our expectations regarding what words they are likely to use. There is reason to think that social identity judgements play a role in speech perception - in our judgements about the very words spoken. I am sure every reader will have found themselves in a situation where they mistook one word for another on the basis of expectations about

what the speaker was going to say. For example, picture yourself in a noisy bar talking to a well groomed, well spoken man in an expensive looking suit. Suppose he is actually a social scientist (you don't know this), and he says 'I work on current opinions on markets'. Due to the noisy environment you could easily mishear him and, partly on the basis of the expectations you have developed as a result of his appearance, hear him as saying 'I work on currency options markets'. If he were less well dressed and had a more working class accent you may mishear him as saying 'I work currently in an open market' (as in, open air market)¹⁴. These are clear cases of social identity judgements leading to miscommunication (for psychological and linguistic research into the impact of social information (including social identity judgements) on sentence processing see Casasanto (2008), Campbell-Kibler (2010), Creel and Bergman (2011), and Summner et al (2014))¹⁵.

Subpersonal reliance on stereotypes is an important feature of interpretation. We have already seen that it can lead us astray. One can imagine many ways in which prejudicial stereotypes in particular can lead to problems. For example, Payne (2002) observed that white subjects primed with images of black faces were more likely to misclassify an image of a pair of pliers as an image of a gun. You might imagine a similar situation arising with speech perception. For example, one might mistakenly hear a black speaker as saying 'I've got a gun ', when saying 'I want some gum'. It is not hard to see how such misinterpretation, if common enough could be very problematic (for example, in cases involving trigger happy police officers). One can also imagine cases where someone's testimony is not understood and, as a result of prejudicial stereotypes, the audience assumes their testimony was irrelevant. For example, imagine an unintelligent and prejudiced supply teacher covering a social studies class. A black male student invokes the notion of hegemony whilst making a point. However, the teacher is not aware of the concept of hegemony, she just assumes that the student is referring to some aspect of urban culture which has little bearing on the class. She thereby writes off his testimony. This would be a case of interpretative injustice since it pertains to the type of content attributed to the student rather than the teacher's assessment of the student's credibility. She may take him to be perfectly credible informant with respect to urban culture, but is

14 To be clear, I don't intend this to be taken as a case of interpretative injustice, for no harm is brought about and the biases leading to the misinterpretation are not prejudicial.

15 This phenomenon can be seen as analogous to that of cognitive penetration discussed by Siegel (2012, 2013).

simply not interested in such information.

Clearly the problems just outlined are important - the first due to its ability to disproportionately subject certain groups to potentially harmful situations, and the second because it constitutes a form of silencing. However, I would like to focus primarily on a more general problem raised by prejudicial stereotypes in interpretation. The problem is simply that, as a result of systematic misleading stereotypes, utterances by members of certain groups are going to be misinterpreted far more often than the utterances of other groups. Accurate stereotypes (or, stereotypes which serve as reliable heuristics) will generally aid hearers in recovering what the speaker intends to communicate, in the same way that reasoning based on accurate generalisations will usually lead us to the truth. However, if there are popular stereotypes about certain groups which get things radically wrong then the application of these stereotypes in utterance interpretation will lead to these groups being disproportionately misinterpreted. Reasoning on the basis of faulty generalisations will generally lead to faulty beliefs.

The problems don't end here though. In her discussion of credibility judgements Fricker notes that stereotypes of disadvantaged groups have often included negative generalisations about intelligence, rationality, and knowledgeableness. Disadvantaged groups are often perceived as unintelligent, irrational, and ignorant. Judgements about these attributes will determine the charitability of the audience's interpretation. If one's interpretation is based on the assumption that the audience is unintelligent and has various false beliefs then one's interpretation will be less charitable than if one interpreted on the basis of the assumption that the speaker is a rational and reliable belief former. The result is that members of disadvantaged groups are more likely to be interpreted uncharitably. The fact that certain groups are disproportionately subjected to uncharitable interpretation is evidenced by a recent study by Alison Brooks et al (Forthcoming), which found that potential investors systematically preferred entrepreneurial pitches by men over identical pitches by women, and that they preferred pitches by attractive men over identical pitches by unattractive men. Participants were asked to rate how persuasive, fact based, and logical the presentations were, and found presentations by men (especially attractive men) to rate higher on each scale. Since the extent to which a presentation is logical or fact based is dependent on the content of the presentation it seems reasonable to conclude that the gender of the entrepreneur had an

impact on the way they were interpreted, with males being interpreted more charitably. Charitability will be especially relevant when we consider Fricker's response to testimonial injustice. First, however, it is worth spending some time considering some cases of interpretative injustice, and identifying some harms which it brings about.

4. Instances of Interpretative Injustice.

In section two it was argued that interpretative injustice constitutes an injustice because it systematically and unfairly harms members of disadvantaged social groups. We also saw how the phenomenon arises. However, we have not yet explored any examples of interpretative injustice. This section discusses several examples.

Firstly, let us consider entrepreneurial pitches, since we have seen that there is empirical evidence that interpretative injustice occurs in such scenarios. Brookes et al (forthcoming) showed that investors preferred identical entrepreneurial pitches by men (attractive men especially), judging such pitches to be more fact based and logical. The fact that identical scripts were read, yet different pitches were interpreted to be more or less fact based/logical suggests that different contents were attributed in each case. But what sorts of different content attributions could plausibly account for such a preference for male over female pitches? Let us consider two of the sorts of sentence which may be spoken in this context. The entrepreneur is likely to utter sentences such as 'we aim to achieve x by 2018' or 'Consumer data suggests x'. These sentences can each be modulated to mean importantly different things. In the first sentence 'aim to achieve' could be interpreted as differently as 'we will achieve' to 'we would like to achieve'. In the second, 'suggests that' could be interpreted strongly, so as to mean something like 'demonstrates/shows that', or weakly, similarly to 'points in the direction of...'. It is easy to see how pitches where interpretations of the former kind are favoured would be preferable to investors, as there appears to be far less uncertainty involved. That is, such pitches would appear to be based on solid claims rather than speculation or wishful thinking. This contrasts with pitches where the latter form of interpretation is more common, where one's investment would seem far more uncertain. We can also see how investors may be more prone to the latter form of interpretation when viewing pitches by females. Where masculinity is often associated with independence, confidence, and groundedness in reason/fact, femininity is typically associated with intuitive and

emotional reasoning, and with vulnerability or helplessness (the 'damsel in distress' stereotype). Moreover, if prejudice is common in such contexts some women may be less confident in the environment. The result is likely to be that the hedged nature of the assertions will be more salient to audiences viewing pitches by female entrepreneurs, as female entrepreneurs may be viewed, in part, as seeking help (rather than offering an investment opportunity), and their positive claims may be viewed with a greater degree of initial skepticism. This gives male entrepreneurs a significant communicative advantage. In order to communicate the same information a female entrepreneur would have to hedge her assertions far less, and be far more explicit. However, the ability to hedge one's assertions in such scenarios offers a clear advantage, because one is not forced to explicitly commit to things about which there is some degree of genuine uncertainty, and it is harder for one to be held to precise claims in the future.

A second example is as follows: consider a female restaurant manager, Amia, discussing front of house recruitment with her assistant manager. This particular restaurant has developed a highly misogynistic kitchen culture, with many of the chefs (all of whom as male) questioning the ability of a woman to run a restaurant. Amia is concerned about the fact that all of the current front of house staff are female, and so she says 'I need a man'. The word 'man' here will need to be modulated (narrowed) by the audience, as not any old man will do. Amia's intended audience will no doubt grasp her intended meaning, and modulate 'man' to mean something like 'male with front of house restaurant competence'. However, imagine that Amia's utterance is overheard by one of the chefs, who is already skeptical of her ability, as a woman, to run a restaurant. There is a very real possibility that he will modulate 'man' differently when interpreting Amia's utterance. That is, drawing on his representation of her as a struggling incompetent and out of place female manager, he may take her to be saying that she needs a male to help her run the restaurant, someone to help take charge and add a 'male touch'. In such a case the concept 'man' will be narrowed even further to incorporate stereotypically male qualities associated with strong leadership. This sort of misinterpretation is likely to further enforce the kitchen's negative impression of Amia, and make her job even harder.

Finally, let us consider one of Fricker's central cases: the trial of Tom Robinson in Harper Lee's 'To Kill a Mocking Bird'. Robinson is a black man falsely accused of rape in depression era Alabama. Robinson is wrongly convicted despite the fact that a

convincing case for his innocence has been put forward. Fricker uses this as an illustrative example in which an audience (the jury) assigns an inaccurate degree of credibility to the speaker as a result of their prejudicial biases, and the speaker is harmed as a result. As Fricker notes, in the racial climate of the trial it is very difficult for the jury to take the word of a black man over that of a white woman, especially in a rape case. This was, after all, a period in which there was still very widespread and open anxiety about the sanctity of white southern womanhood and the supposed threat of black male sexuality. Considering the trial of Tom Robinson, and its historical context, we can come up with clear cases of interpretative injustice. Fricker draws our attention to a particular passage in which Robinson is asked why he visited his accuser's home. Robinson did so because he felt sorry for his accuser, she seemed to live a lonely sad life, and he wanted to help. He expresses this by saying that he felt sorry for her. This goes down very badly. In the climate of the trial the very notion that a black man could feel sorry for a white woman would be found shocking, and the audience would likely read into his utterance that he felt superior to her. This, in and of itself, might be seen as a case of interpretative injustice. However I think we can modify the case to get a clearer example. Let us suppose that Robinson did not state that he felt sorry for his accuser, but rather that he visited her regularly because she seemed vulnerable. Here the term 'vulnerable' could be understood in several different ways. As used by Robinson it could have meant 'in need to help/assistance'. However, in the context of the racist south where the dominant conception of black masculinity involved sexual aggression (especially toward white women), as well as being uncontrolled, animalistic, predatory, uncivilised, and morally inferior, the description of the apparent rape victim as 'vulnerable' could easily be seen as him accidentally revealing that he saw her as an easy target¹⁶. Indeed, if he realised his misstep and 'shifted uncomfortably in his chair' (as he does in the book), this could be taken as further confirmation that he accidentally revealed his true intentions.

It might be worried that such a modification of the Tom Robinson case is implausible, since marginalised groups will generally develop a sensitivity to situations in which their testimony will fail to secure proper uptake, and where such testimony may be harmful¹⁷.

¹⁶ One can see how this would fit into the relevance theoretic story given earlier. The relevant associations of predatory animalistic sexuality will be highly salient to the audience, thus the resultant concept of vulnerability will receive a higher level of activation and be ranked above Robinson's intended meaning.

¹⁷ Thanks to an anonymous referee for this point.

As a result marginalised testifiers will refrain from making such problematic utterances. Kristie Dotson (2011) calls this form of self censorship 'smothering', and argues that it is a form of self silencing. Dotson writes:

'A linguistic exchange that might prompt testimonial smothering concerns situations where unsafe testimony, which is testimony that an audience can easily fail to find fully intelligible, runs the risk of leading to the formation of false beliefs that can cause social, political, and/or material harm. In testimonial smothering, testimony is omitted that is both unsafe and carries a risk of causing negative effects by virtue of being unsafe' Dotson (2011): 244.

Robinson's utterance certainly seems to fit this description. The result is that Robinson never would have uttered a sentence such as 'she seemed vulnerable' in the trial context, because he would have been well aware that such an utterance would be badly misinterpreted. Of course, the phenomenon of smothering is compatible with marginalised speakers occasionally slipping up, as Robinson does in Lee's novel. However, 'she seemed vulnerable' is arguably a far worse slip up than 'I felt sorry for her', so it is less likely to occur. I do not think that this renders the modified Robinson case valueless however. Although it does seem correct that Robinson would be unlikely to make such an assertion in the context of his trial, it does not seem unlikely that he would assert the same sentence in different contexts, where he felt more secure. For example, suppose Robinson is discussing his case with someone he trusted, but is overheard (or is being eavesdropped on) by a prejudicial interpreter. In such a situation Robinson is far less likely to be on the guard against misinterpretation, yet misinterpretation by the eavesdropper could be equally harmful. Moreover, not all marginalised speakers will be equally sensitive to the ways in which they are likely to be misinterpreted. For example, children, the cognitively impaired, or outsiders who are not accustomed to the culture of discrimination they are inhabiting will be far less sensitive to the ways in which they are likely to be misinterpreted¹⁸. Thus, such speakers will be especially vulnerable, even in cases where many marginalised speakers would censor their own speech. Having considered some cases of interpretative injustice, let us now

¹⁸ The case of Emmett Till springs to mind here. Till was a 14 year old African American boy visiting Mississippi from Chicago in 1955. Till was brutally murdered, and the accounts of the events leading up to his murder are mixed. According to some he whistled at a white woman, evidently not realising the significance of such an act in the context. According to others, he stuttered and spoke with a lisp, and was thus misheard as whistling.

turn to consider some of its distinctive harms.

5. The Harms of Interpretative Injustice.

Firstly, as discussed in section two, interpretative injustice silences. For example, the paradigmatic case of problematic silencing, in which a woman is unable to communicate her refusal of a man's sexual advances, appears to be a case of interpretative injustice, as the man misinterprets her use of 'no'. As emphasised by Ishani Maitra, the harms of silencing do not end here. She writes:

'Speech has, at the very least, great instrumental value. It enables us to get what we want and need, for ourselves and others. It constitutes our first line of defence against a variety of injuries, from unwanted sexual overtures to tyrannical governmental action. And it is essential both to the propagation of knowledge, and to the proper functioning of a democratic society. When a speaker is communicatively disabled, she is thereby deprived of these (and other) benefits that speech can offer' Maitra (2009): 331.

As Maitra points out, when one is silenced one is restricted in one's ability to share knowledge. Fricker, in her discussion of epistemic injustice, emphasised the importance of the ability to convey knowledge. She argues that the ability to share knowledge is a fundamental human value, and that when a group is limited in their ability to share knowledge they are thereby limited in their ability to engage in a practice fundamental to human value, and to human society more generally. If prejudicial stereotypes lead to some groups being misinterpreted disproportionately the prejudicial stereotypes make it harder for these groups to enter into the practice of sharing knowledge. Thus, if Fricker is right, these stereotype will make it harder for certain groups to engage in social practices fundamental to human value. Moreover, it will make it harder for these groups to pursue their ends. Brookes et al (forthcoming) illustrates one particular way in which this could occur - female entrepreneurs are at a distinct career disadvantage, seemingly as a result of their tendency to be interpreted uncharitably (at least in certain situations). It is easy too imagine how the same phenomenon could occur in job interviews, marketing pitches, and academic research presentations etc. In such cases the female speaker will be prevented from communicating her intended message, and will be treated as if she has

attempted to communicate something different.

Related harms occur when some groups are interpreted overly charitably. In these cases it is not clear that individuals are directly harmed. However, a practice of systematically interpreting some groups overly charitably will lead to those groups having specific advantages over other groups in a way that produces many of the same harms as silencing. For example, silencing makes it harder for members of certain groups to advance in their careers. Overly charitable interpretation of one group will likewise make it harder for members of other groups to advance in their careers, because it will be harder for them to compete. A possible example of this occurring in philosophy might be those from prestigious institutions being interpreted more charitably than philosophers from lower ranking institutions. It is easy to see how this could occur. Given the difficulty of strictly interpreting even the clearest analytic philosophy, there are usually a few ways an argument can be read. If our knowledge of the prestige of the author influences our interpretation of their argument this will place the author at an advantage over others who we may interpret less charitably (and perhaps more accurately) as a result of the status of their institution (race and gender could conceivably lead to the same form of overly charitable interpretation)¹⁹.

The second harm systematically associated with interpretative injustice involves unjust attributions of responsibility. When a speaker is interpreted as having asserted a proposition p they are held publicly accountable for defending p . This fact is emphasised by MacFarlane (2005, 2011), Hinchman (2005) Goldberg (2006), Rescorla (2009), and McMyler (2013). Hinchman puts the point as follows:

'If you've told someone that p , that person is now under certain conditions - for example, if he's challenged whether p - entitled to hold you accountable for producing a reason to believe that p .' Hinchman (2005): 568.

Of course, in cases of interpretative injustice the speaker does not actually say what they are interpreted to have said. So the audience will not actually be entitled to hold the speaker accountable for having said p . Nonetheless, if the audience thinks the speaker

¹⁹ Of course, judgements based on institutional affiliation may not be completely unreliable, since philosophical ability has at least some role in candidate job placement. I do not wish to take a stand on the usefulness of institutional affiliation as a guide to ability here.

has said that p they will hold the speaker responsible for p , even if they are not entitled to do so. If prejudicial stereotypes result in certain groups being misinterpreted more frequently than others then members of these groups will also be disproportionately held responsible for communicating things they never intended to communicate. Being held responsible for something you never did constitutes a harm.

In combination with other prejudices this could lead to further problems down the line. In her discussion of credibility judgements Fricker notes that some groups have been perceived as inherently dishonest. If this is the case then the marginalised speaker's ability to correct the audience's mistake may also be harmed. They could be perceived as being dishonest and trying to slip out of a commitment. Indeed, if some groups find themselves having to correct mistakes in interpretation more often than others this could add to the perception of that group being sneaky, dishonest, or too ignorant to express themselves clearly. This in turn will contribute to the stereotype that the group is dishonest or stupid. Thus, it will give rise to lowered credibility judgements. It is even conceivable that such factors could contribute to stereotype threat - the phenomenon whereby individuals under perform in line with a stereotype when that stereotype is made salient to them²⁰. That is, individuals may under perform as communicators when stereotypes regarding their honesty or communicative abilities are salient. The plausibility of such a claim will depend on the precise mechanisms of stereotype threat, and the types of under performance at issue. For example, if stereotype threat arises primarily through additional strain on working memory then it is unlikely that marginalised speakers will be less truthful testifiers in cases where the relevant stereotypes are mutually salient, but perhaps they will be less clear (assuming that communicating clearly places demands on working memory). However, if factors such as motivation loss and reduction in effort play an important role in stereotype threat then it is perhaps plausible that speakers may put less effort into ensuring that they are making truthful contributions when confronted with negative stereotypes concerning honesty. Ultimately this is an empirical question.

Relatedly, victims of interpretative injustice will be forced to defend themselves against misinterpretation (both during and following communicative exchanges). This experience of having questions raised against oneself, having one's credibility brought

²⁰ Thanks to an anonymous referee for encouraging me to engage with the idea of stereotype threat here.

into question, and having to actively re-shape the audience's conception of oneself, will often be cognitively and emotionally draining, and sometimes humiliating. This will be especially true of cases in which one's words, which are one's primary means of shaping the audience's conception of oneself, are at risk of being misinterpreted and actually contributing to the conception one is trying to correct²¹. We can imagine a case similar to that of Tom Robinson's in which this occurs. Imagine that Robinson is an outsider to the culture of discrimination in which he finds himself, and describes his reasons for being at his accuser's home by saying 'she seemed vulnerable'. Asked to clarify what he means by 'vulnerable' he might reply 'well, she seemed like she needed a man'. It is unclear what should be built into the concept of 'man' here (presumably not any male will do), and it raises the question of the purpose for which a man is needed. In this case Robinson would have meant that his accuser needed the help of someone who was capable of performing manual labour around the homestead. However, in making this utterance he would once again be at risk of misinterpretation (due to the very same biases which lead to the misinterpretation of 'vulnerable'). The audience's conception of black men as animalistic sexual predators, and their preconceptions of Robinson's likely aims in the scenario, could lead to him being misinterpreted as stating that his accuser desired a man fitting the common black male stereotype of the time (perhaps even with the implication that she desired the rape). This, of course, is far from what Robinson intended. However, after his two misinterpreted utterances any further attempts to clarify his assertions would likely be taken as attempts to wriggle out of what he has said. I'm sure it is clear to the reader how humiliating, frustrating, and emotionally draining this experience would be. We can also imagine a similar continuation of the restaurant manager case. Suppose Amia says the following to explain her utterance to the chef: 'I can't just have a team of girls'. If the chef does not appreciate the need for gender balance then he is unlikely to grasp Amia's clarification. In such a case his interpretation of the fact that Amia considers a front of house team incomplete without any males is likely to be influenced by, and further re-enforce his views about the competencies of women in the workplace. The cumulative cognitive and emotional stress of living in a society in which one struggles to be heard, and in which it is more difficult for one to use language in particular ways (e.g. use of subtlety, artistic language use, or the use of commands), is a

²¹ Of course, the humiliation of trying to defend oneself, and the cognitive drain associated with being heard, will be systematically associated with other forms of silencing. What is not systematically correlated with other types of silencing is the cause (being unfairly held responsible for something one never intended to communicate), or the fact that one's only means for improving the situation (one's words) can be twisted in such a way as to worsen the situation.

harm which arises from the silencing aspect of interpretative injustice. Thus, the first two harms are in this way intertwined.

Finally, if members of certain groups are interpreted uncharitably more frequently than others then this will contribute to the stereotype that members of those groups are poor informants, thereby feeding into the stereotypes which give rise to prejudicial credibility judgements. Consider the following: we have two informants, Max and Sally. If we frequently interpret Sally uncharitably so that, for example, we take her to be asserting that p where p is less plausible than some alternative interpretation q which we would assign to Max in the same circumstances, then the following situations will arise more frequently with respect to Sally than with respect to Max: A) we find Sally's utterance so initially implausible that we lower our judgement of her credibility²², and B) we believe her only to later discover that the proposition we took her to be asserting was false, which once again leads to us to lower our judgement of her credibility. Therefore, as a result of our uncharitable interpretations of Sally we will find ourselves with what seem like good reasons for assigning Max a higher credibility than Sally, even though he may be no better as an informant.

If audiences make credibility judgements on the basis of the speaker's social identity, and the speaker's social identity is a partial determiner of the charitability of the audience's interpretation, then that audience will find themselves in situations A and B more frequently with respect to the testimony of certain social groups than others. This will thereby strengthen the stereotype that particular groups have a low credibility, and thus contribute to prejudicial credibility judgements. We can imagine this occurring in the context of entrepreneurial pitches as discussed above. If an investor frequently has the impression that female entrepreneurs present weaker, less coherent and more illogical pitches, then they may come to the belief that female entrepreneurs are generally not very intelligent or reliable. As a result they may assign a lower credibility to certain

²² It might be thought that we would not assign initially implausible interpretations if we are taking the Gricean maxims to be in effect. This would be a mistake. When applying Gricean norms you need to represent what would be a cooperative contribution given the speaker's representation of the situation. For example, imagine you know that p , and believe that the speaker doesn't know that p . If the speaker were to make an assertion which would imply q only if p was common knowledge, then you would not take them to be implying q . If you take the speaker's representation of the communicative situation to be defective then the contribution you take the speaker to be making will not be the most cooperative given the actual facts, but rather the most cooperative given a particular set of faulty background assumptions. And this contribution may seem implausible.

assertions made by female entrepreneurs, for example assertions which concern their ability to carry out particular projects. Thus interpretative injustice contributes to epistemic injustice.

So far everything I have said can be seen as supplementary to Fricker's account of epistemic injustice. In the final section I argue that the two phenomena interact in an important way. More precisely, I argue that adopting Fricker's proposed solution to testimonial injustice without also treating interpretative injustice will cause the hearer to be epistemically harmed, with only minimal benefit to the speaker.

6. Interpretative Injustice and Credibility Adjustments.

Fricker does not merely identify and describe the phenomenon of epistemic injustice. She also offers a strategy for overcoming it. She argues that the appropriate response to epistemic injustice is to develop the virtue of testimonial justice, whereby we gain a sensitivity to the sorts of situations in which our credibility judgements may be biased, and re-consider our judgements in such circumstances. In such situations we should not rely on our quick intuitive judgements, but should try to make an accurate and unprejudiced judgement. Fricker puts the point as follows:

'When the hearer suspects prejudice in her credibility judgement-whether through sensing cognitive dissonance between her perception, beliefs, and emotional responses, or whether through self-conscious reflection - she should shift intellectual gear out of spontaneous, unreflective mode and into active critical reflection in order to identify how far the suspected prejudice has influenced her judgement. If she finds that the low credibility judgement she has made of the speaker is due in part to prejudice, then she can correct this by revising the credibility upwards to compensate. There can be no algorithm for her to use in determining how much it should be revised upwards, but there is a clear guiding ideal. The guiding ideal is to neutralize any negative impact of prejudice in one's credibility judgements by compensating upwards to reach the degree of credibility that would have been given were it not for the prejudice.' Fricker (2007): 91-92.

The hope is that we eventually become more reliable habitual judges of credibility, and that prejudicial stereotypes will no longer affect our credibility judgements. At this point we will display an instinctive sensitivity to actual signs that the speaker has low credibility and not judge speakers to have low credibility on the basis of factors such as race and gender.

It might be thought that developing the virtue of testimonial justice would make us more reliable belief formers. For example, suppose that on reflection I realise that I have an underlying bias against a particular race, and that I assign members of that race a lower credibility than I should. Next time I have an interaction with a member of that race I actively reconsider my intuitive credibility judgement, and compensate upwards. The effect of this is that the speaker is no longer prevented from communicating their knowledge to me, and I no longer miss out on knowledge I can acquire from the speaker.

This solution becomes less straightforward once interpretative injustice is taken into account. Consider the case of Max and Sally from the previous section. However, imagine that in this case we have not only been interpreting Sally uncharitably, but have also been assigning her a low credibility, and thus not trusting her testimony. After reading Fricker's 'Epistemic Injustice' and reflecting on the way we assign credibility we realise that we have been assigning Sally a low credibility on the basis of prejudicial stereotypes. As a result we no longer trust our intuitive credibility judgements. Instead we actively reason about her credibility. Sometimes this results in our rejecting her testimony for reasons other than prejudicial bias. However, on other occasions it results in us compensating upwards and assigning her a high degree of credibility. That is, in many cases when Sally makes an utterance and we take her to be asserting a proposition p , we go against our instinctive (but prejudicial) credibility judgement and assign her a higher credibility, thereby leading us to believe p . Certainly we will find ourselves believing what we take Sally to have said more often than we would have done otherwise.

It should be clear why this will lead to problems if we do not also adjust for interpretative injustice. Consciously assigning Sally a higher credibility will not necessarily affect the charitability of our interpretation²³. The practice of assigning high

23 It is perhaps a matter of contingent empirical fact that the practice of actively adjusting one's credibility

credibility to uncharitably interpreted utterances will lead to unreliable belief formation, and will thus be harmful to the hearer. Moreover, in the cases where Sally is misinterpreted she will still be blocked from sharing her knowledge. The audience will bear an epistemic cost which is of little benefit to the speaker. This result applies more generally. If we treat epistemic injustice in the way Fricker suggests, without also treating interpretative injustice, then in many cases we run the risk of not only continuing to harm the speaker, but also harming ourselves.

To be clear, I am not claiming that Fricker's solution fails to treat testimonial injustice. If we raise our credibility assignments then we will avoid disrespecting the speaker as a knower. Thus one of the main harms of testimonial injustice will be removed. Rather, my point is cautionary. Treating testimonial injustice in the way Fricker suggests without also treating interpretative injustice will, in cases where both interpretative and testimonial injustice are present, put the hearer at epistemic risk, and only make a minimal positive difference to the speaker.

It might be thought that this is unproblematic, since the virtue of testimonial justice can be applied earlier on in the process of testimonial belief formation. This will involve applying one's sensitivity to one's own underlying biases in order to recognise not only situations in which one's credibility judgements are misleading, but also cases in which one's interpretation may be rendered defective. Once such a sensitivity has been developed one can adjust one's interpretation accordingly. This will involve shifting gear from being a passive spontaneous interpreter to being an active interpreter, and perhaps applying a heuristic similar to Fricker's 'adjust credibility upward' heuristic.

The first point to note here is that the problem raised above was only intended to illustrate the risks for treating testimonial injustice without also treating interpretative injustice. If one shifted gear earlier on then one would not be ignoring interpretative injustice. Thus, there is no tension with the point above. However, there are two further points which indicate that it may be less straightforward to shift gears with respect to one's interpretation than it is to do so with respect to one's credibility judgements. The first (minor) point is simply that although the processes underlying both interpretation

judgements will eventually change one's overall perspective of the speaker, which will also solve the problem of interpretative injustice. However it is an empirical question (to which we don't have an answer) whether this will be the case.

and credibility judgements are subpersonal, credibility judgements seem more easily accessible to consciousness. We are more used to actively reasoning about a speaker's credibility during conversational exchanges than we are to reasoning about how to interpret them.

The second point is that credibility judgements seem to be scalar in a way in which interpretation is not. That is, it might be thought that our credibility judgements fall somewhere on a scale, and in order to adjust our judgements we must merely shift where we locate the speaker's utterance on that scale. With interpretation there are multiple possible meanings that need to be ranked for plausibility given our knowledge of the context. Thus, the task of adjusting interpretation cannot be reduced to a simple heuristic such as 'shift credibility upward'. This is not to say that such heuristics are of no use at all. For example, our credibility adjustments may affect the charitability of our interpretation. Thus, employment of a heuristic such as 'adjust credibility upward' if made before the utterance actually takes place, may lead to more charitable interpretation (at least in cases where the audience's negative credibility judgement would have led to misinterpretation had it not been corrected). However, due to the complex and multifaceted nature of interpretation such simple heuristics will only take us so far.

If a heuristic based approach is to be adopted, we may be better off adopting an approach similar to that offered by Karen Jones (2002). Jones notes that we should separate credibility judgements from our judgements regarding the prior probability of the asserted contents, as the credibility we assign can affect the probability we judge the asserted content to have. The result is that in cases where audiences make prejudicial credibility judgements they are likely to also assign a lower prior probability to the content of the testimony. The same seems to be true with interpretative injustice. The credibility we assign can have an impact on the way we interpret the speaker. Indeed, this may be one of the mechanisms by which credibility assessments affect judgements of prior probability. That is, there are two ways in which we might think of the interaction between credibility assessments and the assignment of prior probability to asserted contents. We could hold the asserted content p fixed and maintain that credibility assessments affect our judgement regarding the prior probability of p (this seems possible, and I take it to be what Jones has in mind), or a low credibility judgement could cause us to interpret the speaker as saying something to which we

antecedently assign a low prior probability. This could occur as a result of interpretative injustice (indeed, it may be that the probabilities we assign to asserted contents are sometimes affected in both these ways). However, credibility judgements are only one potential factor which may bias interpretation. Moreover, accurate credibility judgements, and some social identity judgements, may be fruitfully employed in interpretation. Thus merely separating credibility assessments and interpretation does not seem a promising strategy. Another of Jones's heuristics seems to offer us a more promising route. Jones's third heuristic for making credibility judgements calls for one to seek more corroborating evidence when it is less reasonable for one to trust one's own ability to judge the credibility of particular sorts of speaker. A similar strategy might be employed to deal with interpretative injustice. That is, one might attempt to develop a sensitivity to the kind of case in which it is less reasonable for one to trust one's own interpretation, and employ more evidence (or, more active reflective reasoning) when interpreting in such cases. This will not entirely solve the problem as the employment of such heuristics will place higher cognitive demands on audiences when interpreting marginalised speakers. As a result, employment of such heuristics will still leave certain groups at a communicative disadvantage compared to those for whom no adjustment is needed. However, such an approach may be the lesser of two evils.

Conclusion.

I have introduced the notion of interpretive injustice and explained the numerous harms to which it gives rise. It gives rise to silencing, unjust attributions of responsibility, and it contributes to the stereotypes which give rise to epistemic injustice in our credibility judgements. Finally, I discussed the relationship between interpretative injustice and Fricker's proposed treatment of testimonial injustice. This discussion is instructive as we start to think about how we might treat interpretative injustice. Fricker's solution to testimonial injustice may be thought of as having two components. Firstly, she recommends that we develop a sensitivity to the biases which play a role in our credibility judgements, and secondly she recommends that we adjust our credibility judgements upward in situations where we suspect that our credibility judgements may be influenced by such biases. I have argued that interpretative injustice complicates Fricker's proposed solution. I agree that developing a sensitivity to the biases which guide one's behaviour (both in interpretation and credibility judgements) will be an

important factor in a treatment of both interpretative and testimonial injustice. However, simply raising the credibility one assigns to the speaker, in cases where one's interpretation is uncharitable, fails to prevent silencing, and it also epistemically harms the audience. Raising one's credibility judgement before interpretation has taken place will help insofar as misinterpretation arises from stereotypes associated with the credibility of the informant. However, the factors which contribute to interpretation are not identical to those which contribute to credibility assignments, and they may vary between contexts. So application of Fricker's proposed heuristic prior to interpretation will not solve the problem of interpretative injustice either (indeed, it was not designed to do so). Thus, interpretative and testimonial injustice should be treated together, and the treatment of interpretative injustice requires a new heuristic. Moreover, it seems that whatever heuristic we apply will have to be somewhat different in form to Fricker's 'adjust upward' heuristic, as interpretation is multifaceted and non-scalar. A heuristic such as 'evoke more evidence in interpretation' (similar to Jones's (2002) third rule for credibility assignments) seems more appropriate. However, heuristics along these lines differ from Fricker's heuristic in an important way - they are far more cognitively demanding. Thus, application of the heuristic will be, in some ways, detrimental to the audience²⁴, and may also be detrimental to speakers. That is, if audiences from non-marginalised groups are under a heavier cognitive burden when communicating with marginalised speakers than with speakers from their own community then this itself will put marginalised speakers at a disadvantage. This may simply be the lesser of two evils. However, if we develop a sensitivity not only to the types of bias which guide interpretation, but also to the sorts of situation in which misinterpretation is likely to be particularly harmful, then we may be able to avoid some of the most harmful forms of misinterpretation without placing cognitive demands on audiences that are more detrimental than helpful. That is, we must adjust the amount of evidence appealed to in interpretation not just as a function of the chance of error, but also as a function of the costs of error²⁵.

24 See Szabó Gendler (2011) for a discussion of the cognitive and epistemic costs of adjusting for implicit bias.

25 Developing such a sensitivity will not only involve considering direct harms which may result from particular misinterpretations, but also considering the ways in which particular subtle and widespread misinterpretations contribute to systematic disadvantage. Doing so may not be simple, as the harms of such misinterpretations will often be hidden to those who inflict them. Thus, as suggested by Medina (2013), a development of such a sensitivity may involve active engagement with different perspectives and social movements.

Bibliography.

- Boland, J., Clark, S. (Manuscript). Homophone Disambiguation and Vocal Stereotypes.
- Brookes, A., Huang, L., Kearney, A., Murray, F. (Forthcoming). Investors Prefer Entrepreneurial Ventures Pitched by Attractive Men. *Proceedings of the National Academy of Sciences of the United States of America*. 1-5.
- Campbell-Kibler, C. (2010). Sociolinguistics and Perception. *Language and Linguistics Compass*. 4 (6). 377-389.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Wiley Blackwell.
- Casasanto, L. (2008). Does Social Information Influence Sentence Processing? *30th Annual Meeting of the Cognitive Science Society*. 1-6.
- Creel, S., Bergman, M. (2011). How Talker Identity Relates to Language Processing. *Language and Linguistics Compass*. 5 (5). 190-204.
- Dotson, K. (2011). Tracking Epistemic Violence, Tracking Practices of Silencing. *Hypatia*. 26 (2). 236-257.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press.
- Geiselman, R. Bellezza, F. (1977). Incidental Retention of Speaker's Voice. *Memory and Cognition*. 5 (6). 658-665.
- Goldberg, S. (2006). Reductionism and the Distinctiveness of Testimonial Knowledge. From Lackey and Sosa (Eds). *The Epistemology of Testimony*. Oxford University Press. 127-144.
- Grice, H.P. (1989). *Studies in the Way of Words*. Cambridge MA: Harvard University Press.
- Hinchman, E. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research*. 70 (3). 562-587.
- Hornsby, J., Langton, R. (1998). Free Speech and Illocution. *Legal Theory*. 4 21-37.
- Jones, K. (2002). The Politics of Credibility. From Antony and Charlotte (Eds). *A Mind of One's Own: Feminist Essays on Reason and Objectivity, 2nd edition*. Boulder: Westview. 154-176.
- Kukla, R. (Forthcoming). Performative Force, Convention, and Discursive Injustice. *Hypatia*. 1-18.
- Lee, H. (1960). *To Kill a Mockingbird*. J.B. Lippincott & Co.

- MacFarlane, J. (2005). Making Sense of Relative Truth. *Proceedings of the Aristotelian Society*. 105. 321-339.
- MacFarlane, J. (2011). What is Assertion? From Brown and Cappelen (Eds) *Assertion*. Oxford University Press. 79-96.
- Maitra, I. (2009). Silencing Speech. *Canadian Journal of Philosophy*. 39 (2). p 209-338.
- McMyler, B. (2013). The Epistemic Significance of Address. *Synthese*. 190. 1059-1078.
- Medina, J. (2013). *The Epistemology of Resistance: Gender and Racial Oppression, Epistemic Injustice, and the Social Imagination*. Oxford University Press.
- Payne, K. (2001). Prejudice and Perception: The Role of Automatic and Controlled Processes in Misperceiving a Weapon. *Journal of Personality and Social Psychology*. 81 (2). 181-192.
- Pohlhaus Jr, G. (2012). Relational Knowing and Epistemic Injustice: Toward a Theory of Wilful Hermeneutical Ignorance. *Hypatia*. 27 (4). 715-735.
- Rescorla, M. (2009). Assertion and its Constitutive Norms. *Philosophy and Phenomenological Research*. 79 (1). 98-130.
- Siegel, S. (2012). Cognitive Penetrability and Perceptual Justification. *Noûs*. 46 (2). 201-222.
- Siegel, S. (2013). The Epistemic Impact of the Etiology of Experience. *Philosophical Studies*. 162. 697-722.
- Sperber, D., Wilson, D. (1986) *Relevance: Communication and Cognition*. Oxford: Blackwell; Cambridge, MA: Harvard University Press.
- Sumner, M., Kim, S., King, E., McGowan, K. (2014). The Socially Weighted Encoding of Spoken Words: A Dual-Route Approach to Speech Perception. *Frontiers in Psychology*. 4. 1-13.
- Szabó Gendler, T. (2011) . On the Epistemic Cost of Implicit Bias. *Philosophical Studies*. 156 (1). 33-63.
- Van Berkum, A., Van Den Brink, D., Tesink, C., Kos, M., Hagoort, P. (2008). The neural Integration of Speaker and Message. *Journal of Cognitive Neuroscience*. 20 (4). 580-591.