

Defective Contexts

Andrew Peet

1. Introduction

Successful communication depends not only on our knowledge of language, but also our knowledge of context. If a speaker utters the sentence “he is going to get burnt,” we will have to rely on our knowledge of the context in order to grasp what proposition they are trying to express. If there is a mutually salient individual in front of us, whose trousers have just caught on fire, then we will know that this salient individual is the intended referent. If we were talking about our mutual friend Frank, and somebody has just asked how Frank’s latest business deal is going, it will be clear that Frank is the intended referent. Two completely different propositions are expressed in these situations, and without contextual knowledge, we would not be able to tell which proposition was expressed.

Contexts can be fruitfully modelled as bodies of shared information presupposed by the conversational participants (Stalnaker (1973, 1978, 2002)). A speaker presupposes a proposition if they take it for granted for the sake of the conversation. We can define defective and non-defective contexts in these terms (Cf. Stalnaker 1978):

NON DEFECTIVE CONTEXT: One in which all the presuppositions of the participants are the same.

DEFECTIVE CONTEXT: One in which the presuppositions of the participants differ.

It is commonly assumed that most communicative interactions take place within non-defective contexts (or contexts which are close enough to being non-defective for it not to matter). Stalnaker (1978) offers the following line of reasoning in support of this assumption:

A defective context will exhibit a kind of instability, and will tend to adjust to the equilibrium position of a non-defective context. Because hearers will interpret the purposes and content of what is said in terms of their own presuppositions, any unnoticed discrepancies between the presuppositions of speakers and addressees is likely to lead to a failure of communication. Since communication is the point of the enterprise, everyone will have a motive to try and keep the presuppositions the same. And because, in the course of a conversation, many clues are dropped about what is presupposed, participants will usually be able to tell that divergences exist if they do. So, it is not unreasonable, I think, to assume that in the normal case contexts are non-defective, or at least close enough to being non-defective.

A context is CLOSE ENOUGH to being non-defective if the divergences do not affect the issues that actually arise in the conversation. Suppose, for example, that you know that Jones has won the election, believe mistakenly that I know it as well, and are prepared to take the truth of this proposition for granted if the occasion should arise, say by using it as a suppressed premise in an argument, or by using the description *the man who won the election* to refer to Jones. On my dispositional account of speaker presupposition, if you are prepared to use the proposition in this way, then you DO presuppose that Jones won the election, even if you never have a reason to display this disposition because the subject never comes up. Since I do not know that Jones won the election I do NOT presuppose it, and so the context is defective. But the defect may be harmless.

It will not necessarily be harmless: If the news is of sufficiently urgent interest, your failure to raise the subject may count as a display of your disposition to take it for granted. There will not exactly be a failure of communication, but there will be a misperception of the situation if I infer from the fact that you do not tell me who won that you do not know either.
Stalnaker, 1978, p. 85-86.

If this line of reasoning is sound then defective contexts (or, at least, contexts which are defective in some consequential manner) will constitute deviations from normality. They will be the sort of marginal cases we can harmlessly ignore in much (though not necessarily all) of our linguistic theorising. After all, communication will typically occur in non-defective contexts. So, if we want to understand how communication functions then these are the contexts on which we should focus.

I believe that defective contexts should occupy a more central place in our linguistic theorising. This is because I reject the claim that defective contexts constitute deviations from normality. I believe that it is fairly standard for our communicative interactions take place within defective contexts. Thus, to better understand how linguistic communication functions we must understand how it works in defective as well as non-defective contexts.

In this chapter I hope to persuade you that defective contexts are more ubiquitous than we typically assume. In doing, so I will draw attention to a number of pressing social and theoretical issues which arise once we start to consider defective contexts. I will proceed by pointing to a number of ways in which defective contexts can emerge without self-correcting in the manner envisioned by Stalnaker. First I will consider situations in which some, but not all interlocutors recognise that the context is defective. I will then consider cases of opaquely defective contexts, and the question of what it is for a context to be “close enough” to being non-defective. I will close by suggesting some further avenues for exploration.

2. Transparently Defective Contexts

Stalnaker writes that when one party recognises a failure of alignment in presuppositions, their presuppositions will be adjusted so as to render the context non-defective. This can happen two ways. Firstly, an audience might make explicit the fact that they do not share the speaker’s presuppositions. The speaker will, if cooperative, respond by removing the offending item from their presuppositions. Alternatively, the audience may adjust their own set of presuppositions to accommodate the speaker’s presupposition.

We will consider each approach, and see that they can (and in some important cases will) fail to render the context non-defective. In the former case, an interlocutor’s attempt to make the defect explicit might fail to be recognised or understood, in which case the defect will remain. In the latter case, we will see that accommodation is not an all or nothing affair. An audience might accommodate p for certain purposes, but not for all the purposes the speaker has in mind. This will not only fail to render the context non-defective, but it may also damage the speaker’s epistemic position with respect to the offending presupposition.

2.1 Hermeneutical Impasses

Our ability to fix a defective context by making a defect explicit depends on our ability to make mutually apparent the existence and nature of the defect. That is, if we recognise and reject an agent’s presupposition of p , we must be able to make them recognise that they are presupposing p , that we don’t presuppose p , and that we refuse to accommodate their presupposition. Furthermore, to render the context non-defective, we must be able to do this without introducing further defects. That is, it is no good removing an agent’s presupposition of p if this simply leads them to adopt an alternative but equally problematic presupposition.

It is not always easy to make problematic presuppositions mutually apparent in this way. Firstly, not all presuppositions are easily articulated, meaning that attempts to make them explicit can lead to further confusion. Secondly, certain presuppositions are more central to the presupposer's understanding of the world than others. These central presuppositions are, in a sense, part of the bedrock against which the agents understand and interpret the rest of the world. When an interlocutor attempts to reject a presupposition their act must be interpreted by their audience. Yet, if the presupposition is central to the audience's understanding of the world then that presupposition will likely form part of the basis for their attempt to interpret the speaker's act. The speaker's attempt to reject the presupposition and fix the context will, thus, be unsuccessful.

As an illustrative but unrealistic example, we might consider how difficult it would be for a solipsist to make transparent their rejection of an ordinary (non-philosophical) interlocutor's presupposition that there is an external world occupied by many minds. This presupposition is so central to the way most of us understand and interpret the world that we cannot be easily shifted from it. A basic presupposition of the way we understand communicative interactions is that they are interactions between minded individuals in a public space. Thus, the solipsist's attempts to make explicit and remove this presupposition will usually result in failure. Call cases like this 'Hermeneutical Impasses' (I take this terminology from Anderson (2017)).

The case of the solipsist is hypothetical. It illustrates the possibility and structure of hermeneutical impasses, but it is the sort of case which can be harmlessly ignored in serious theorising about language. Yet hermeneutical impasses are not uncommon. All they require is for a presupposition (or set of presuppositions) to be central to one interlocutor's understanding of the world, and yet utterly objectionable from the point of view of another. This kind of situation arises with a worrying frequency. Salient examples arise with respect to the differing assumptions about race which have shaped dialogue between African Americans and white Americans. It is worth considering some such examples.

I will start with a much discussed case; 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. He is in an impossible position. He is facing an all white jury which harbours the common racial beliefs of the time: that black people are violent, untrustworthy, and by their very nature inferior to whites. These beliefs are not fleeting or easy to shift. They are central to the way in which the jury members understand the social world they inhabit. Without these beliefs they would not be able to live the lives they live; they would not be able to sustain their position of dominance. Their failure to fully grasp the humanity of black people forms a sort of self-serving ignorance which is essential for the maintenance of the social order from which they benefit (Mills (2007)). These beliefs also form the background presuppositions against which they interpret Robinson's utterances, and assess his credibility as an informant (see Fricker (2007) on this latter point). Pohlhaus Jr. (2012) calls our attention to a particular instance in the text where these presuppositions manifest themselves and bring about misunderstanding: During his cross examination, Robinson is asked why he was at his accuser's home. He tells the jury that he felt sorry for her. Having done so, he immediately realizes his mistake, and shifts uncomfortably in his seat. The notion that a black man could feel empathy for a white woman runs contrary to the racist background assumptions which form the basis for the jury's interpretation of his speech. As a result, his utterance is interpreted as suggesting that he felt superior to his white accuser (see Peet (2017) for an account of how such a misinterpretation can arise). Following his utterance, it now becomes a presupposition of the jury, something they take to be common ground, that Robinson felt superior to his accuser. It is hard to see how Robinson could correct this misinterpretation. Even if he somehow managed to make it clear that he didn't feel this way, the jury still needs some way to interpret his utterance. And the correct interpretation, that he felt empathy for her and understood her situation, is one they are prevented from reaching by the racist presuppositions that are so central to their understanding of the

social world.¹ In order for the context to be rendered non-defective the jury must be parted from their racist beliefs (an impossible task in this setting).

Anderson (2017) draws our attention to a similar situation which has arisen with respect to the Black Lives Matter movement: To be properly understood, the slogan “Black Lives Matter” must be understood against the background presupposition that, at present, black people in America, unlike white people, are treated as if their lives don’t matter (especially by the police). It has to be understood that white people are in a privileged position and that, at present, black people often (and rationally) fear for their lives at the hands of the police. This is hard for many white Americans to accept. Even for those who recognise that racial inequality still exists in the U.S., a full recognition of the extent of such inequality can be hard to grasp and come to terms with. This ignorance is, much like the southern racists’ ignorance, self-serving. It is central to the ability of many white Americans to get on with their lives, and benefit from the social order as it stands. This ignorance leads to misinterpretation. A common response to “Black Lives Matter” has been that “all lives matter,” as if this second claim was inconsistent with that of the Black Lives Matter activists. The suggestion is that Black Lives Matter activists are claiming that only Black Lives Matter, or that Black Lives Matter in some special sense. Such an interpretation only makes sense against the assumption that African Americans are not treated any differently from white Americans (at least with respect to the value of their lives). Fixing this context by making the defect salient would require forcing those who reply that “all lives matter” to recognise their privileged position relative to black people, and the realities of the situation of black people in the U.S. Given the role that such ignorance plays, it will not be straight forward to remedy.

Hermeneutical impasses don’t arise only when a presupposition is central to an interlocutor’s understanding of the world. Sometimes the concepts required to render explicit the rejection of a presupposition are not available as part of the common conceptual repertoire of the interlocutors. In such cases conceptual engineering may be required. The clearest examples involve hermeneutical injustice: the phenomenon whereby a deficiency in our shared conceptual repertoire renders members of marginalized social groups unable to make sense of or articulate to others important social experiences. Fricker (2007) presents the example of a woman trying to make sense of and articulate her experience of what we would now call ‘sexual harassment’ before this concept was first coined. The perpetrator of the harassment would likely have thought of their behaviour as flirting (or some other socially acceptable form of behaviour). The victim rejects this characterisation of the interaction. But it is hard to do so convincingly without an alternative concept to fill the gap.

Hermeneutical impasses illustrate one clear and important way in which defective contexts can fail to be self-correcting. In these cases, the interlocutors will, at least sometimes, share a mutual interest in communicating. It will be in their interests to render the context non-defective by making the defective presupposition mutually apparent, and removing it from everyone’s representation of the common ground. Yet doing so will be close to impossible. The only other alternative is for the problematic presupposition to be accommodated. But this would, in cases like those discussed above, involve giving up too much. Moreover, it is not clear that accommodation will, in practice, always render a context non-defective. This is the theme of the next section.

2.2. Accommodation and Omissive Implicature

When interlocutors recognise that their presuppositions do not align they have a number of options. The first is to object to the presupposition. We have just seen that this won’t always fix the context. Moreover, this is an extreme measure. It will often feel confrontational, and can involve a certain degree of risk. Another alternative is to accommodate the presupposition. That is, when one

¹ Imagine, for example, that he simply states “I don’t mean that I felt superior to her”. This would force the audience to identify a satisfactory re-interpretation of his original utterance which doesn’t carry the implication of felt superiority. That is, it would require that they conceive of a black person as being able to feel sorry for a white woman without felt superiority. But if this possibility is in tension with their central presuppositions then no such interpretation will be forthcoming, and his attempted clarification will come across as disingenuous.

interlocutor recognises that another interlocutor is presupposing something they were not previously presupposing, they can render the context non-defective by quietly adding the relevant proposition to their set of presuppositions. This is a common feature of ordinary conversations. If you didn't know that I had a child, and I stated 'I am sleep deprived because my son wakes up and cries every hour or so', then you will likely add to your set of presuppositions the proposition that I have a son. You will recognise that I was presupposing this, and in order to render the context non-defective you will add this proposition to your set of presuppositions.

However, accommodation is not an all or nothing affair. It can be partial or fleeting, and when this is the case it can fail to render the context non-defective. Imagine that you have a friend who is a known bull-shitter. They will often make up stories or exaggerate their exploits in order to make themselves appear more interesting than they actually are. You find this annoying, but you are willing to let it slide because they have other redeeming virtues. Suppose they state, "when I climbed El Capitan I used a kilo of chalk." You do not believe that they ever climbed El Capitan, however you do not want to call attention to or explicitly reject this presupposition. Doing so would risk generating open conflict, and you do not wish to jeopardize your friendship. You could just accommodate the presupposition, and act, for the rest of the conversation, as if your friend did climb El Capitan. However, there is also a range of intermediate options.

Firstly, of course, one could adopt the passive aggressive strategy of making clear one's rejection of the presupposition by making statements which presuppose or imply (but don't explicitly state) that it is false. You might, for example, ask about your friend's progress on a beginner level climb you know they have been struggling with. This serves to make it clear that you reject their presupposition, but it does so without generating open conflict. More generally, one can make statements which insinuate the falsity of the presupposition in such a way that one retains plausible deniability about the insinuation. This way, if explicitly challenged, one can deny that one rejects speaker's presupposition (see Camp (2018) for discussion of insinuation and plausible deniability). This strategy does not involve accommodating the presupposition in any substantive sense, it simply allows one to avoid explicit rejection. Whether or not this strategy produces a non-defective context will depend on whether one's rejection of the presupposition is recognised and accommodated by one's interlocutor.

However, the passive aggressive strategy still carries risks. There are better options if one wishes to avoid conflict. For example, one might simply change the subject. Doing so would not render the context non-defective, for your friend would continue to presuppose something you do not. They will continue to presuppose that they have climbed El Capitan, whilst you may, unbeknownst to them, be disposed to object to the presupposition were it raised again. However, changing the subject might render the context "close enough" to being non-defective, for it will remove the conversational relevance of the problematic proposition.

Similarly, one might accommodate the presupposition to a certain extent, but place a limit on how accommodating one is willing to be. For example, you may adopt a disposition to not object, or make one's rejection of the presupposition known unless put on the spot, and forced to make a statement which carries with it the relevant presupposition. For example, you may let your friend continue with their tall tales until they ask what you were doing whilst they were on El Capitan, at which point you might make your rejection of their presupposition apparent. This would be to act, in some respects, as if the presupposition is true, but it would not amount to full accommodation. One is not disposed to act as if the proposition is true for the sake of conversation. Such partial accommodation will only partially mend the context. But it will be enough to allow for smooth communication, and it will be enough to avoid conflict.

The above examples are mundane and relatively inconsequential. However, the strategy of partial accommodation, or of attempting to avoid the subject, can be problematic. As Stalnaker hints at the end of the quoted passage, we can communicate things with our silence. It is unlikely that we will

communicate much when we fail to call out our lying friend (other than, perhaps, that we don't disbelieve them). However, sometimes there will be an expectation that, if p were false, we would object to the presupposition or statement of p . This will particularly be the case when somebody presupposes a proposition they know to be controversial or inflammatory. If we fail to object, then this can communicate that we share the speaker's belief in relevant presupposition (loosely following Swanson (2017) I'll call this "omissive implicature"). Consider the following case from Maitra (2012):

An Arab woman is on a subway car crowded with people. An older white man walks up to her, and says "Fucking terrorist, go home. We don't need your kind here." He continues speaking in this manner to the woman, who doesn't respond. He speaks loudly enough that everyone else in the subway car hears his words clearly. All other conversations cease. Many of the passengers turn to look at the speaker, but no one interferes
Maitra, 2012, p. 101.

The victim's silence is unlikely to communicate much. Nobody would expect her to object, as it is likely that she will feel highly intimidated. The silence of the other passengers will communicate somewhat more. It is true that, as Swanson (2017) notes, there will usually be a number of possible explanations for the passengers' silence. One explanation might be that they share the speaker's beliefs. But an alternative explanation may be that they too feel threatened into silence, or that they consider the speaker so far below them that they are not worth challenging. Thus, to anybody for whom these possible explanations are salient, the passengers' silence will not communicate much.

However, not all of these explanations will be salient to the speaker. The speaker is unlikely to see themselves as being threatening (at least, not toward the non-Arab passengers), and they are unlikely to see themselves as being unworthy of the other passengers' engagement. Among the most salient hypothesis for the speaker will be that at least a reasonable number of the other passengers share their beliefs, or at least respect their right to state their hateful views. This will have a number of detrimental effects. Firstly it will render the context defective. It may be that, if put on the spot, any of the passengers would explicitly reject the speaker's racist beliefs. In this sense, they are not presupposing that the speaker's beliefs are acceptable, nor are they presupposing the authority of the speaker to make such statements (contra Langton (2018)). After all, they are not disposed to act as if these statements are either true or acceptable, nor that the speaker has the authority to make such statements. Thus, there will be a failure of alignment in presuppositions between the speaker and audience.

Secondly, and more interestingly, the audience's omissive implicature may strengthen the speaker's belief that their racist views are acceptable or correct, and that they have the authority to publicly state them. That we can, in general, rationally strengthen our belief in p by asserting that p is illustrated by the following example from Goldberg (2016):

SCIENCE TEAM: Research Team X specializes in a certain subfield of physical chemistry. The members of X have assembled excellent evidence for their hypothesis H, where (they recognize that) this evidence would be sufficient to convince even the most demanding physical chemist of the acceptability of H. Consequently, they go to a prestigious national physical chemistry conference and announce their findings, concluding that H is acceptable on these grounds. Now, the members of X acknowledge that in such a large audience there will be other people who have different background knowledge, different evidence, and different expertise. Given (their knowledge of) the prestige of the conference and the excellence of the audience, as well as the norms of conference behaviour, they reasonably think to themselves that if there had been a flaw in their experimentation or an alternative explanation for their data someone in the audience would likely have known of this and would have spoken. So, seeing no one raise any such objection, and observing the audience's favourable reaction to their announced findings, they increase their confidence in the thought that there are no relevant pieces of counter-evidence, alternative explanations for their data, or flaws in their experimental set-up.

Like the scientists, the angry racist is performing their act in front of a large and diverse body of peers. If they were wrong, they may reason, somebody would have objected. In this way the audience's silence is damaging to the speaker. It provides them with misleading evidence for the acceptability of their views.

Perhaps this is not such a plausible hypothesis when applied to the angry racist on the train. Perhaps, in reality, any such individual would have sufficient self-awareness to realize that the other passengers are simply intimidated by them, or do not want to escalate the situation. But more mundane versions of this type of scenario are worryingly common. I imagine many readers will have had the experience of being forced to listen to a taxi driver outline their unhinged views on immigration, or of trying to politely escape a political discussion with an intoxicated bigot at a bar. We may on occasion object in such situations. But many of us will not. Rather, we will try to change the subject, or politely remain silent unless explicitly challenged. In doing so we risk communicating our agreement, and we harm the speaker. The more the speaker's statements go unchallenged, the stronger the misleading evidence that their statements are correct. A one-off incident may not have a noticeable effect. However, if a speaker finds themselves in such situations frequently, and they are regularly met with a non-confrontational silence, this could cumulatively lead to significant epistemic damage.

This section has ended on something of a digression. We have moved from consideration of whether or not defective contexts are typically self-correcting, to a discussion of the harms which can be brought about by failure to fix the context (even when the context is close enough to being non-defective for communication to proceed smoothly). However, the examples are nonetheless illuminating. They illustrate the manner in which Stalnaker's argument rests on a problematic simplification about conversation: the claim that, in conversation, communication is the aim of the game. Clearly communication will typically be one of our primary aims in conversation. However, it is one amongst many. Conversation also serves as a means of bonding, of generating familiarity, as forming a basis for cooperation, or as an opportunity for self-expression. Our desire to, for example, avoid conflict, or allow the speaker to feel that we are on good terms, may override our interest in communicating smoothly. The result will often be defective contexts which fail to self-correct.

It may be objected that Stalnaker's argument, charitably understood, is more limited in scope. That is, perhaps we should understand Stalnaker to simply be claiming that in cooperative contexts where communication *is* the aim of the game, defective contexts will typically be self-correcting (at least, where this is possible). This is a far weaker and less interesting claim. After all, very few (if any) of our conversational exchanges will be purely communicative in aim. So there will be little motivation for the claim that we can ignore defective contexts in our linguistic theorizing.² Furthermore, it might be objected that the situations discussed so far are somehow atypical, and that Stalnaker's claim should be taken only to pertain to typical, or normal contexts. In the following section I will argue that even typical contexts, where communication is one of the primary aims, can be problematically defective without self-correcting.

3. Opaquely Defective Contexts

In the previous section we considered what might be called transparently defective contexts: contexts in which at least one of the conversational participants realizes that the context is defective, but is either unable or unwilling to correct it. However, contexts can also be opaquely defective. That is, defects can

² To be clear, Stalnaker himself never claims that we can largely ignore defective contexts in our theorising. This is simply something that follows fairly naturally from the stronger reading of his argument. So this more charitable construal of his argument may well be more accurate. This is not of great consequence here. I take the stronger version of the argument to be *prima facie* compelling regardless of whether or not it is endorsed by Stalnaker. And it is the stronger version of the argument that might tempt us to sideline defective contexts in our theorizing. So it is the stronger version of the argument that is my focus.

arise without anybody realizing it. Here I will focus on one way in which opaquely defective contexts might arise: through the potentially approximate and inexact nature of typical linguistic communication.

It seems quite plausible that communication is, fairly often at least, approximate and imprecise. Minor and seemingly inconsequential misunderstandings are not out of the ordinary. Much of our speech is loose, and many of us regularly communicate with concepts we only partially understand. As Yalcin (2014) puts a related point, it may be that, typically, “coordination on items of content is a highly approximate, more-or-less affair, with perfect coordination on content not being especially important, and rarely or never happening” (Yalcin, 2014, 24). If this is right, then defective contexts will abound: in typical communicative exchanges, one interlocutor’s representation of the common ground will merely approximate the other interlocutors’ representations of the common ground. After all, the proposition the audience adds to their representation of the common ground (i.e., the propositions they immediately entertain and automatically attribute to the speaker through their sub-personal systems of comprehension) will merely approximate the proposition intended by the speaker (i.e. the proposition the speaker will think of themselves as asserting and which they will, as a result, add to their representation of the common ground). However, if the interlocutor’s representations of the common ground approximate each other closely enough, these differences will usually go unnoticed.

There are three questions which are natural to ask when presented with this argument:

1. Do we have good reason to believe that communication is typically imprecise in these ways?
2. If communication is typically imprecise in this way, wouldn’t interlocutors take this imprecision into account whilst updating their representations of the common ground, thus rendering the resulting contexts non-defective?
3. Surely, if defective contexts commonly arise from the imprecision typical of linguistic communication, the resultant contexts will still be “close enough” to being non-defective? This seems especially obvious if, as you say, the interlocutors’ representations of the common ground will typically approximate one another closely enough for any differences to go unnoticed.

In the following subsections I will address each of these questions in turn.

3.1. The Imprecision of Communication

I won’t attempt to establish beyond doubt that linguistic communication is typically imprecise and approximate in the manner discussed above. To do so would take us beyond the scope of this chapter. Instead, I will merely aim to illustrate the plausibility of the hypothesis through a discussion of the prevalence of conceptual variation: the idea that different interlocutors will typically assign slightly different concepts to the same words. There are other potential sources of approximation and imprecision in communication, such as vagueness, loose talk, and context sensitivity. However, I have discussed many of these issues at length elsewhere, so I will refrain from doing so here.³

I take it to be uncontroversial that different interlocutors will often understand the same word in slightly different ways. They may have slightly different patterns of association, their conceptions may center around slightly different prototypes, they may be disposed to make slightly different inferences with the same concept, etc. For example, my prototype for “cat” is likely modeled on my own cat, a relatively small, furry animal, fairly passive in behavior (except when provoked by mischievous hands), and very affectionate. These will be among the features which make up my ‘cat’ prototype. Likewise, my familiarity with this cat affects the associations and inferences I will most immediately and automatically make when somebody uses the word “cat.” Your prototype for “cat” will likely be similar. That is, it will list many of the same features. However, for some readers at least, there is likely

³ See Peet (2015, 2016, 2018, and forthcoming). See Davies (forthcoming a, forthcoming b) for responses to some of these considerations.

to be a small degree of divergence. At the very least, these features may be weighted differently. The same will likely be true of the typical associations and inferences automatically drawn when interpreting utterances using the word 'cat'.

If factors such as those discussed above determine what a word means in a speaker's idiolect, then, often, the same word will mean different things to different interlocutors. However, this latter claim is controversial. Externalists, for example, reject the claim that meaning is determined (solely) by such factors, instead holding that the meaning a word has for a speaker will be determined in large part by their environment. That is, the meaning 'cat' has when I use it (in thought or talk) does not just depend on the various associations and inferences which affect my use of the term. For externalists the meaning of 'cat' depends on the use of the term 'cat' in the wider linguistic community I implicitly defer to when using the term. And these patterns of use will, in turn, be determined by the various patterns of inference and association which underlie and guide the use of 'cat' by members of this community. As a result, as long as all the interlocutors are implicitly deferring to the same linguistic community, the meaning of 'cat' will be the same for each interlocutor. Thus, when a speaker says "my cat is sleeping" the audience, by implicitly deferring to the same linguistic community as the speaker, will be able to update their representation of the common ground with precisely the same proposition as that intended by the speaker.

I am doubtful that this will solve the problem. There are two reasons for this. Firstly, the externalist response depends on the assumption that interlocutors will typically defer to the same linguistic communities. But it is not clear that this assumption is correct. An individual act of deference will always be directed at some linguistic community. There will be some linguistic community it picks out. There must be some story to tell about which linguistic community is picked out by a given act of deference. One natural story to tell here is that the relevant linguistic community will be determined by the speaker's conception of the community most relevant to their present use of the relevant term. But this leaves space for small idiosyncrasies to emerge in different interlocutor's conceptions of the communities to which they are deferring. Such idiosyncrasies seem inevitable, much as they do with our prototypical conceptions of, say, cats. The result is that interlocutors will typically defer to slightly different (although largely overlapping) linguistic communities in typical situations. Tiny variations in the overall patterns of use in each community will determine tiny differences in meaning assigned by each interlocutor. So, when one speaker says "my cat is sleeping" the proposition each interlocutor updates their representation of the common ground with will be slightly different (but extremely similar). Of course, it might be objected that the community deferred to is not determined by the audience's conception of the relevant community. But then we're owed a story about what *does* determine the relevant community. It cannot be deference again, for then we would be led to a regress. Alternatively, it could be that some groups are reference magnets for states of implicit deference. It is not clear to me how such reference magnetism would work. However, it may be possible to tell a story here. So I do not wish to put too much weight on this consideration. My point is simply that it is an open (and pressing) question for externalists exactly what determines which precise community is deferred to in a given instance of thought or talk. And whether or not the externalist is able to hold that interlocutors typically coordinate precisely in communication will depend on how they respond to this challenge.⁴

My second worry regarding the externalist response is that even if externalists are able to guarantee that there is some proposition (the proposition determined by use of the relevant sentence in the relevant community) which is always added to each interlocutor's representation of the common ground, there is nothing to stop additional propositions determined by each interlocutor's idiosyncratic uses being

⁴ This challenge becomes especially pressing when we consider items of language other than, say, natural kind terms. For natural kind terms the relevant community may be experts regarding the subject matter at hand. Such communities may have epistemic properties which make them reference magnets for states of deference. It is less clear what community would be reference magnetic for uses of words like "bottle" or "chair", and how such reference magnetism would actually work. Thanks to Sandy Goldberg and Joey Pollock for discussion here.

added to their individual representations of the common ground, or shaping their understanding of the common ground in problematic ways.

To see this, consider Tyler Burge's classic example of the patient who thinks she has arthritis: Betsy believes she has arthritis in her thigh. She goes to the doctor, and tells her doctor as much, only to be told that arthritis is an inflammation of the joints, meaning that one cannot get arthritis in one's thigh. Betsy takes this on-board, and ceases to believe she has arthritis in her thigh. Prior to her visit to the doctor Betsy had an idiosyncratic understanding of the concept "arthritis." Even though she deferred in her use of the term in everyday thought and talk, the idiosyncrasies of her prior understanding would have affected her representation of the common ground when she communicated with the term. For instance, if somebody told her, "I'm pain-free other than the arthritis in my lower extremities," Betsy would, through her implicit deference to her linguistic community, update her representation of the common ground with the proposition that the speaker is only experiencing arthritis based pain in their lower extremities (cf. Burge (1979)). However, whilst the speaker would take themselves to be eliminating from the common ground worlds in which they have pain in their thigh, for Betsy these situations would still seem to be live possibilities. In a sense, Betsy's representation of the common ground will be inconsistent, although not in a way that is transparent to her.⁵ The arthritis case involves a fairly major idiosyncrasy in understanding. However, smaller idiosyncrasies in understanding (such as those arising from slightly different associations and inferences underlying use of the term 'cat') are common. And I see no reason to doubt that they would influence our representations of the common ground in a similar way (but to a much smaller degree).

Indeed, it appears that such differences in understanding are often manipulated by strategic speakers. If a term is understood by one segment of the population as carrying a particular association, or as giving rise to particular inferences, then a speaker can use this term in order to communicate something to that segment of the population that may be missed by others. This allows the speaker to maintain plausible deniability regarding the hidden message. Such terms are referred to as 'dogwhistles' or 'code words' (see Mendelberg (2001), White (2007), Stanley (2015), Khoo (2017), and Saul (2018)). For example, suppose a politician wishes to target voters with underlying racial prejudices, but does not want to be seen as overtly racist. They might highlight the fact that their opponent's proposed welfare policies disproportionately target 'inner city families', hoping that their intended audience (racially resentful white voters) will understand them as communicating that black families will benefit disproportionately from their opponent's policies. There is much to be said about code words, but I will not discuss them further here since they have a chapter of their own in this volume. The important point for our purposes is that it is unclear how they could have their powerful effects if everybody understood them in the same way. That is, if everybody understood code words in the same way it is not clear that speakers could use them as a means to maintain plausible deniability.

Thus, I conclude that the thesis of conceptual variation is plausible, even on externalist views of meaning. Furthermore, it seems clear that if the thesis of conceptual variation is correct, linguistic communication will typically be approximate (to varying degrees).

3.2. Epistemic Responsibility when Communicating Approximately

At this point it is natural to object that if communication is typically approximate in the way I have suggested, then surely interlocutors will take this into account when updating their representations of the common ground. If I know that your use of the word 'cat' is likely to differ slightly from my own, but I don't know how, then it would be epistemically irresponsible of me to update my representation of the common ground in accordance with my own specific use of the word 'cat.' Rather, I should update the context in a more coarse-grained way. I should eliminate fewer possibilities from my

⁵ It might be held that the context in this situation is non-defective, since each interlocutor's representation of the common ground is identical, but their conceptions of (or understandings of) the common ground differ. Strictly speaking this might be correct. However, there is a clear sense in which such contexts are still defective. Coordination on content is not the be-all and end-all of communication. Mutual understanding matters as well (Pollock (2015)).

conception of the common ground than I would otherwise. In doing so I am able to prevent the context from becoming defective.

There are a number of problems with this objection:

1. The above objection assumes that language users are standardly aware of the fact that we typically use the same words in ever so slightly idiosyncratic ways. However, it is far from clear that this is the case. Ultimately this is an empirical question. But I am highly doubtful that most unreflective language users ever think about this, and even more doubtful that it effects their linguistic practice.

2. Although forming a more coarse-grained belief may reduce the risk of forming a false testimonial belief, it is unlikely to render one's context non-defective.⁶ In order for the context to be rendered non-defective the speaker and hearer must each adjust their representations of the common ground in precisely the same way. Suppose we have two interlocutors, Ashima and Kai. Ashima asserts the sentence S . For Ashima, S expresses the proposition p^1 . For Kai it expresses the proposition p^2 . But they are each aware of the possibility of divergence in meaning. So neither of them updates their representation of the common ground with the precise proposition they associate with S . Rather, they each adjust with a more coarse-grained proposition (i.e. a proposition which eliminates fewer worlds from the set of open possibilities). In order for this to work, they must update the context with exactly the same proposition, which means they must each make slightly different deviations from their standard use of the term. It is hard to see how, without knowing the precise idiosyncrasies of one-another's use of S , they could reliably converge on the same proposition with which to update their representations of the common ground.⁷

3. Even if it was possible for interlocutors to adjust their updates to the common ground in such a way that they converge on precisely the same meanings, this would be highly impractical. Communication moves quickly. Interpretation is largely unreflective. If communication required that we reflect precisely on the ways in which our uses of various bits of language may differ, and adjust our representations of the common ground accordingly, then communication would be massively cognitively demanding. This would be highly impractical. As long as the interlocutors coordinate approximately, it will not be worth the extra expenditure of cognitive resources to secure precise coordination. Thus, even if such adjustments were possible in principle, it is unlikely that they typically occur in practice.

For these reasons, I find the epistemic responsibility concern unpersuasive. However, my third response to this objection lends further force to the final potential worry I highlighted in the opening of §3. Surely, to the extent that communication is typically approximate, the contexts which result are close enough to being non-defective for it not to matter. After all, if they were not close enough to being non-defective it would be worthwhile to expend the additional cognitive resources necessary to render them close enough to being non-defective. I will consider this objection in the final section.

3.3. "Close Enough" Communication

In the previous sections I have suggested that typical communicative contexts will be defective; different interlocutors will typically have slightly different representations of the common ground.

⁶ For discussion of the "coarse-grained proposition" approach to testimonial belief in cases of inexact communication arising from context sensitivity see Peet (2016) and Davies (forthcoming b).

⁷ It might be thought that vagueness can salvage the situation. That is, Kai and Ashima each update their representations of the common ground in a vague way, and thus avoid updating their respective representations of the common ground with different propositions. This does not solve the problem, for we are still left with the question of how they are to update their representations of the common ground in the same vague way. Indeed, on some views of vagueness (such as many valued approaches) introducing vagueness may make coordination even more miraculous (MacFarlane (2016)).

However, Stalnaker never claimed that typical communicative contexts would be non-defective. Rather, he claimed that they would typically either be non-defective, or close enough to being non-defective for it not to matter. It is natural to feel that the considerations I have adduced at best suggest that contexts will rarely be completely non-defective, but will nonetheless be close enough to being non-defective for it not to matter.

This line of thought is further buttressed by two considerations I appealed to in the previous sections. Firstly, to the extent that typical interlocutors' representations of the common ground differ, they will usually approximate one another well enough for any differences to go unnoticed. Secondly, the cognitive capacities which underlie our ability to communicate, and the social and epistemic norms that govern and guide our communicative practices, are adaptive. They have developed to allow us to achieve our communicative ends. If typical communicative contexts were not close enough to being non-defective for it not to matter, then our norms and capacities would surely adapt in such a way as to ensure that we do typically coordinate well enough for any divergences in meaning to be inconsequential (call this 'the adaptiveness argument').

There is a great deal to be said for this line of reasoning. However, it forces us to ask what it really amounts to for a context to be close enough to being non-defective for it not to matter. This question is not easily answered. One might of course claim that a context is close enough to being non-defective if the interlocutors' representations of the common ground are close enough to allow for successful communication. This was hinted at in §2.2 when I noted that Stalnaker's remarks on defective contexts may best be taken as concerning contexts where communication is the name of the game. But this just pushes the problem back a step: what is successful communication? If successful communication requires precise coordination on meaning or content, then contexts typically will not be sufficiently non-defective for it not to matter.

We can make progress by considering the core idea behind the adaptiveness argument: linguistic communication has many functions. There are many important things it allows us to do. The various mechanisms and norms which underlie our communicative practices have evolved and adapted to allow us to do these things smoothly and successfully. That is, they have evolved to allow us to successfully achieve our communicative ends. So, communicative contexts will typically be close enough to being non-defective for us to be able to successfully achieve our communicative ends.

The problem we now run into is that linguistic communication has many different functions. And these apparent functions will typically require different levels (and potentially different forms) of coordination on meaning and content.⁸ For example, in communicating with others we attempt to coordinate our actions and plans so that we can work together to achieve ends of mutual interest, we acquire true beliefs (and even knowledge) from others, we express our feelings and emotions to others, we establish power relations and shape others' cognitive worlds in ways which give us an advantage over them, we agree and disagree, we consent, and we issue commands. Some of these functions will no doubt be more central to the way our norms and linguistic abilities have developed than others. With respect to those functions which are more peripheral to the pressures and circumstances which gave rise to our linguistic norms, it is an open question whether typical communicative contexts really are sufficiently close to being non-defective for these functions to be reliably achieved.

For many of these functions to be reliably achieved it may not even be required that audiences typically form *true* beliefs via communication. As long as the beliefs they form are close enough to the truth to have the same behavioral consequences in a wide enough range of circumstances it is not clear that there would be sufficient pressure to force further adaptation. Furthermore, even if the acquisition of true beliefs is a central function of linguistic communication (as I am inclined to think it is), it is not clear that our communicative practices would be sufficiently attuned to reliably support the acquisition

⁸ For two recent forms of pluralism about communicative success see Abreu Zavaleta (2018), and Pollock (Forthcoming).

of *knowledge*. Truth and knowledge have the same behavioral consequences. Yet it might be thought that the reliable acquisition of knowledge would place considerably higher demands on the systems of norms and cognitive capacities which underlie our communicative practices than the acquisition of true belief. For example, knowledge has a modal component: it requires that we get it right not only in the actual world, but also all nearby worlds. If our communicative practices developed in such a way that for instance, we acquired true beliefs 90% of the time, approximately true beliefs 9% of the time, and dangerously false beliefs 1% of the time, they would have high statistical reliability. However, this would not necessarily secure modal reliability. It would be consistent with such a situation that when we form true beliefs there will usually be a small set of nearby worlds in which we acquire only approximately true beliefs. If this was the case, then typical communicative contexts would be sufficiently non-defective to support the reliable spread of true belief, but not knowledge.⁹

Of course, this is purely speculative. But it serves to illustrate that just because communicative contexts are typically close enough to being non-defective to support the central functions of linguistic communication, we should not conclude that we can safely ignore defective contexts in our philosophical and linguistic theorizing. Indeed, I hope that throughout this paper it has become clearer to the reader just how important it is to be cognizant in our theorizing of the various different ways in which communicative contexts can be rendered defective, and the important impact such defects can have. Defective contexts are ubiquitous. They are not typically self-correcting. And they have significant potential to undermine our communicative aims in various ways.

Conclusion

This article has merely scratched the surface, and I have raised more questions than I have answered. I hope that future research on defective contexts will answer some of the questions I have raised. However, the scope of this paper has been very narrow. There are many issues for future work on defective contexts to contend with which have not been broached here. I would suggest the following as promising additional avenues for future research in this area:

1. In this article I have focused on defective contexts defined as those in which the presuppositions of the interlocutors do not coincide. But there are sophisticated alternatives to presuppositional approaches to context. For example, linguists such as Cragie Roberts (2004, 2012), and philosophers such as Schoubye and Stokke (2016) see contexts as being structured by the questions under discussion. Such approaches to context have potential to shed light on contextual deficiencies missed by focusing only on contexts as sets of presuppositions.
2. I have primarily focused on defects which negatively affect our ability to spread information. But, as mentioned at various points above, conversation has many different aims. It is conceivable that contexts in which information transfer occurs smoothly may nonetheless be defective along other dimensions. For example, a discourse could be framed in terms of concepts, which are offensive or otherwise harmful to some of the interlocutors. We need to have a better grasp of the various aims of communication in order to fully understand the various ways in which contexts, and communication more generally, can be defective.
3. I have primarily focused on face to face communication between individuals. However, our means of communication are rapidly expanding. And with new forms of communication come new ways for communication to go wrong. We must be cognizant of the ways in which developing communication technologies (such as new forms of social media) constitute defective platforms for reliable (or otherwise unproblematic) communication.

⁹ How pressing this worry is will depend on how demanding one takes the conditions on knowledge-yielding communication to be. In Peet (Forthcoming) I argue that the conditions for knowledge-yielding communication are actually fairly undemanding. However, many previous authors have taken the conditions on knowledge-yielding communication to be significantly more demanding (for example, Goldberg (2007), and Heck (1995)). If they are correct, then this worry may be pressing.

Papers Cited

- Abreu Zavaleta, Martín. (2018). *Semantic Variance*. PhD Thesis. New York University.
- Anderson, L. (2017). Hermeneutical Impasses. *Philosophical Topics*, 45 (2), 1-19.
- Burge, T. (1979). Individualism and the Mental. *Midwest Studies in Philosophy*, 4, 73-121.
- Camp, E. (2018). Insinuation, Common Ground, and Conversational Record. In D. Fogal, D. Harris, and M. Moss (Eds.) *New Work in Speech Acts*. Oxford: Oxford University Press.
- Davies, A. (Forthcoming a). Testimony, Recovery, and Plausible Deniability: A Response to Peet. *Episteme*
- Davies, A. (Forthcoming b). Testimonial Knowledge and Context Sensitivity: A New Diagnosis of the Threat. *Acta Analytica*.
- Fricke, M. (2007). *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Goldberg, S. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification*. New York: Cambridge University Press.
- Goldberg, S. (2016). Can Asserting that p Improve the Speaker's Epistemic Position (And is that a Good Thing). *Australasian Journal of Philosophy*, 95 (1), 157-170.
- Heck, R. (1995). The Sense of Communication. *Mind*, 104 (413), 79-106.
- Khoo, J. 2017. Code Words in Political Discourse. *Philosophical Topics* 45 (2), 33-64.
- Langton, R. (2018). The Authority of Hate Speech. In J. Gardner, L. Green, and B. Leiter (Eds.) *Oxford Studies in Philosophy of Law, Volume 3*. Oxford: Oxford University Press.
- Lee, H. (1960). *To Kill a Mockingbird*. New York: J.B. Lippincott & Co.
- MacFarlane, J. (2016). Vagueness as Indecision. *Proceedings of the Aristotelian Society*, 90, 255-283.
- Maitra, I. (2012). Subordinating Speech. In I. Maitra and M.K. McGowan (Eds.) *Speech and Harm: Controversies over Free Speech* (pp. 94-121). Oxford: Oxford University Press.
- Mendelberg, T. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton, NJ: Princeton University Press.
- Mills, C. (2007). White Ignorance. In S. Sullivan and N. Tuana (eds.) *Race and Epistemologies of Ignorance* (pp. 11-39), New York: SUNY Press.
- Peet, A. (2015). Testimony, Pragmatics, and Plausible Deniability. *Episteme* 12 (1), 29-51.
- Peet, A. (2016). Testimony and the Epistemic Uncertainty of Interpretation. *Philosophical Studies*, 173 (2), 395-416.
- Peet, A. (2017). Epistemic Injustice in Utterance Interpretation. *Synthese*, 194 (9), 3421-3443.
- Peet, A. (2018). Testimonial Knowledge without Knowledge of What is Said. *Pacific Philosophical Quarterly* 99 (1), 65-81.
- Peet, A. (Forthcoming). Knowledge-Yielding Communication. *Philosophical Studies*.
- Pohlhaus Jr, G. (2012). Relational Knowing and Epistemic Injustice: Toward a Theory of Wilful Hermeneutical Ignorance. *Hypatia*, 27 (4), 715-735.
- Pollock, J. (2015). Social Externalism and the Problem of Communication. *Philosophical Studies*, 172 (12), 3229-3251.
- Pollock, J. (Forthcoming). Context and Communicative Success. In T. Ciecierski and P. Grabarczyk (eds.), *The Architecture of Context and Context-Sensitivity*. Studies in Linguistics and Philosophy, Springer.
- Roberts, C. 2004. Context in Dynamic Interpretation. In L.R. Horn, and G. Ward (eds.) *The Handbook of Pragmatics*. Oxford: Blackwell Publishing. 197-220.
- Roberts, C. 2012. Information Structure in Discourse: Toward an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics* 49 (5). 1-69.

- Saul, J. 2018. Dogwhistles, Political Manipulation, and Philosophy of Language. In D. Fogal, D. Harris, and M. Moss (eds.) *New Work on Speech Acts*. NY: Oxford University Press. 360-383.
- Schoubye, A. Stokke, A. 2016. What is Said? *Noûs* 50 (4). 759-793.
- Stalnaker, R. (1973). Presuppositions. *Journal of Philosophical Logic*, 4, 447-457.
- Stalnaker, R. (1978). Assertion. *Syntax and Semantics*, 9, 315-332.
- Stalnaker, R. (2002). Common Ground. *Linguistics and Philosophy*, 25 (5-6), 701-721.
- Stanley, Jason. (2015). *How Propaganda Works*. Princeton, NJ: Princeton University Press.
- Swanson, E. (2017). Omissive Implicature. *Philosophical Topics*, 45 (2), 117-137.
- White, I. 2007. When Race Matters and When It Doesn't: Racial Group Differences in Response to Racial Cues. *American Political Science Review* 101(2), 339-54.
- Yalcin, S. (2014). Semantics and Metasemantics in the Context of Generative Grammar. In A. Burgess and B. Sherman (eds.) *Metasemantics: New Essays on the Foundations of Meaning* (pp. 17-55). Oxford: Oxford University Press.