

Testimony, recovery and plausible deniability: a response to Peet

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

Forthcoming in *Episteme*

Abstract: According to telling based views of testimony (TBVs), B has reason to believe that p when A tells B that p because A thereby takes public responsibility for B's subsequent belief that p. Andrew Peet presents a new argument against TBVs. He argues that insofar as A uses context-sensitive expressions to express p, A doesn't take public responsibility for B's belief that p. Since context-sensitivity is widespread, the kind of reason TBVs say we have to believe what we're told, is not widespread. Peet doesn't identify any problem with his own argument though he does attempt to limit its sceptical potential by identifying special contexts in which TBVs stand a chance of success. A more general defence of TBVs can be provided by showing Peet's argument to be unsound. I argue that Peet's argument is unsound because it requires us to wrongly suppose that speakers do far less labour than their audiences in context-sensitive linguistic communication. I aim to show why—in the context of the epistemology of testimony and the philosophy of language—it's important to recognize the labour that speakers can do, and so can be held responsible for not doing, in episodes of context-sensitive linguistic communication.

1 Telling based views of testimony

According to Telling Based Views of Testimony (TBVs), an audience has a certain kind of reason to believe that p when she is told that p. The speech act of telling someone that p is, on such views, a public act of taking responsibility for a speaker's belief that p, when that belief is formed in response to being so-told. So if the speaker tells someone that p and p turns out to be false, then the speaker becomes blameworthy. According to TBVs, it's this feature of the speech act of telling that provides an audience with a reason to believe that p when she is told that p.

TBVs are supposed to ward off scepticism about whether testimony based beliefs can be well-justified or constitute knowledge. Peet (2015) presents an argument for the conclusion that TBVs fail at this task: because a speaker frequently uses context-sensitive expressions to express the propositions she asserts to others, a speaker hardly ever performs a public act of taking responsibility for her hearer's belief that p when that belief is formed in response to such assertions. Peet then tries to rein in the scope of this conclusion by describing special kinds of context in which context-sensitivity does not have this devastating effect on the kind of reason described by TBVs.

In this paper, I will argue that Peet's argument is unsound. Peet helps himself to a description of context-sensitive linguistic communication which allocates the lion's share of the work in context-sensitive linguistic communication to the audience. The description thus implies that the audience is generally responsible for failures of communication, not the speaker. This description, however, is partial. The speaker does quite a bit of performative work in making her intended content clear and unequivocal to her audience; and she is accountable for this work. Once we recognize this, two things seem to follow. Firstly, Peet's argument includes a false premise and for that reason, the threat posed by linguistic context-sensitivity can be reined in much more thoroughly than Peet suggests. Secondly, research into context-sensitive linguistic communication within the philosophy of language has been shaped by this description in ways that obscure the likely importance to successful context-sensitive linguistic communication of the broader actions pursued by speaking.

2 Peet's argument against TBVs

Peet defines plausible deniability as follows:

Plausible deniability

An agent has plausible deniability about intending to communicate a proposition *p* with an utterance *u* of a sentence *s* if that agent is able to tell a story (with at least some degree of believability) about their attitudes and expectations at the same time of utterance such that a reasonable agent with those attitudes and expectations could utter *s* with no intention to communicate *p*. (Peet, 2015, p. 31)

Notice that this story does not have to be true.

Peet believes that if a speaker has plausible deniability about intending to communicate that *p* then she can deny responsibility for her audience's belief that *p* when the speaker has asserted that *p*. She can convincingly say that she wasn't taking responsibility for beliefs with that content because *p* wasn't the content she intended to communicate. Moreover, if she *can* deny responsibility for the audience's belief that *p*—regardless of whether or not she in fact *does* deny it—then (according to Peet) she wasn't publicly taking responsibility for the audience's belief that *p*. Consequently, the kind of reason that an audience has to believe that *p* when she is told that *p*, according to TBVs, is absent.

Even if Peet is right up to this point, there need be no problem for TBVs. If plausible deniability is something that sometimes happens, but far from always, then a TBV adherent can simply concede that sometimes audiences don't have a reason to believe what has been communicated to them. Plausible deniability will be a problem for TBVs only if plausible deniability is widespread. In that case, the kind of reason audiences have to believe that *p* when *p* is communicated to them generally doesn't exist. However, setting aside the

exceptions that I'll introduce at the end of this section, Peet believes that if the speaker uses a sentence to communicate that *p* to her audience which includes a context-sensitive expression, then, the speaker will have plausible deniability. Thus, given that context-sensitive expressions are widespread, it follows that plausible deniability is too and thus, in turn, that in a large range of cases, the kind of reason an audience has to believe that *p* when she is told that *p*, according to TBVs, is absent.

Why think that context-sensitivity provides speakers with plausible deniability? An expression is context-sensitive if and only if that expression can have different contents in different contexts. For example, in one context, one can use the adjective "tall" in such a way that five-year-old John falls within its extension whereas in another context one can use the adjective "tall" in such a way that John does not fall within its extension, even though we speak of John in the same condition in both contexts. Notice that context-sensitivity doesn't require a change in the meaning of a word. The meaning of "tall" doesn't change in the way, for example, that the meaning of the word "girl" has changed throughout the word's history: it used to mean young person, it now means young female person. Context-sensitivity is not meaning change. Context-sensitivity is the under-determination of content by context-invariant meaning.

Context-sensitivity, thus understood, has implications for linguistic communication using context-sensitive expressions. If a sentence includes context-sensitive expressions, then there isn't just one proposition that corresponds to it. The sentence can be used to express an open-ended range of different propositions. Hence, an audience cannot read-off the proposition that a speaker intends to tell an audience just by examining the sentence. The audience will have to rely upon additional means. That means will probably consist of what Peet calls "knowledge context":

Knowledge context

the audience's representation of factors such as mutual knowledge, past utterances, Gricean norms, the mutual goals of the conversation, and any other information generally relevant to interpretation other than basic knowledge of the syntax and invariant semantic content of the utterance. (Peet, 2015, p. 31)

However, some of the representations in the audience's knowledge context are liable to be false (cf. (Peet, 2015, p. 31)). This possibility allows the speaker to plausibly deny that the proposition she intended to tell the audience is the proposition the audience thought the speaker intended to tell the audience. She can plausibly deny this by exercising her ability to tell a story about her own attitudes and expectations such that a reasonable person with such attitudes and expectations would not have intended to communicate the proposition the audience took her to communicate. Here is one of Peet's two examples of such a story:

Scalar implicature

We are planning a group trip to a theme park and deciding how many cars to take. I wish to cause logistical problems because I hate fun, so I say 'Matt has three kids' knowing that he has five. On the day of the trip Matt arrives with his five children and we don't have enough space in the car. You challenge me for saying that Matt had only three children. However, I might attempt to maintain that I didn't mean he had *only* three children, I meant that he had *at least* three children, so we would need at least three additional seats. The plausibility of such a story will depend on the way I said 'Matt has three kids' and the immediate preceding utterances, however we rarely recall such minute details of the conversational context, so it would not be difficult to construct a plausible story on which I intended to communicate only that Matt has at least three children. (Peet, 2015, pp. 32–33)

If this example is representative of what can be done when the speaker makes use of a context-sensitive expression, then, with Peet, it seems that we can generalize:

[I]f recovery of an asserted content requires extensive appeal to knowledge context then the speaker will often be able to claim that the audience recovered the wrong proposition, thus disclaiming responsibility for the audience's belief. (Peet, 2015, p. 33)

In this passage, Peet's 'thus' indicates that there is a second conditional in his reasoning that has the consequent of the first conditional as its antecedent. We know, also, from what else he says, that he thinks the antecedent of the first conditional is true. So, making this all explicit, Peet's argument can be laid out as follows:

Peet's Anti-TBV Argument

(1) If recovery of an asserted content requires extensive appeal to knowledge context then the speaker will often have plausible deniability.

(2) If the speaker will often have plausible deniability then she will often be able to disclaim responsibility for the audience's belief.

(3) Recovery of an asserted content often requires extensive appeal to knowledge context.

Therefore:

(4) The speaker is often able to disclaim responsibility for the audience's belief.

If we accept (4) then we have to accept that TBVs are unfit for purpose. If a speaker is able to disclaim responsibility for her audience's belief then she didn't perform a public act of taking responsibility for her audience's belief. So if the speaker is often able to disclaim responsibility for the audience's belief, then the kind of telling which, according to TBVs, gives audiences a reason to believe what they're told, typically doesn't happen.

I said earlier that Peet tries to rein in the scope of the conclusion of his argument. He (2015, p. 48) describes a set of resources which make it harder for 'a speaker to claim mismatch without absurdity.' He thinks that these resources are absent in certain discourses, but not others. He (2015, p. 49) lists as possible discourses in which they are absent: religious discourse, public political discourse, and ethical discourse outside of academic settings. However, Peet does not think that these resources are generally effective in the discourses where they are available. For Peet also distinguishes between one-off attempts to escape responsibility for telling someone that p when p is false by claiming that one never intended to communicate that p, and, repeat offenders, who repeatedly make such such claims. According to Peet (2015, pp. 48–49), the resources which make it harder for a speaker to claim mismatch without absurdity are generally effective only against repeat offenders. Repeat offenders will make their own claims of mismatch increasingly implausible in discourses where the resources Peet describes are available.

For the sake of argument, let's grant that all of this is true. Even so, Peet only manages to provide a way to resist the conclusion of his argument as it applies to repeat offenders speaking in a subset of discourses. Beyond such exceptions, Peet's conclusion remains in place: i.e. the kind of reason we have for believing what we're told—according to TBVs—is absent. I think this is an unnecessarily partial response to the challenge he raises against TBVs. We can find a much more general response to his challenge by showing that one of the premises in his argument is false.

3 Applying the argument to Fricker's and the assurance theorists' TBVs

Peet targets two TBVs in particular: Elizabeth Fricker's (2006, 2012) and assurance views of testimony (Peet lists: (Hinchman, 2005; Moran, 2005; Ross, 1986)). I think Peet wrongly presents his anti-TBV argument as having application, in identical ways, to these two TBVs. Whereas he explains in some detail why Fricker's TBV is vulnerable to plausible deniability, he assumes without argument, that parallel considerations apply to the assurance TBV. Thus when stating the problem that plausible deniability poses for the assurance theorists, Peet writes:

As we observed when discussing Fricker's view, the speaker's public commitment to defend the audience's belief is precisely the type of commitment the deniability problem undermines... speakers don't undertake the assurance theorist's required type of commitment in making [their] utterances (Peet, 2015, pp. 38–39)

In this section, I will argue that the considerations which Peet puts forward in favour of thinking that a speaker's plausible deniability undermines her ability to make Fricker's commitments do not favour thinking that a speaker's plausible deniability undermines her ability to make the commitments described by the assurance theorists.

Being TBVs, according to both Fricker's TBV and the assurance theorists' TBV, the audience of a speaker has a reason to believe that p when the speaker tells her that p because the speaker publicly takes responsibility for the truth of p in telling the audience that p. However, according to Fricker's TBV, the audience has this reason because the speaker will face social sanction if she tells the audience that p when p is not true. It's the threat of social sanction that keeps the speaker in check and which justifies an audience in believing the speaker. For Fricker (2012, pp. 62–63), it's because a speaker would not easily have asserted that p unless p were true, that the audience is justified in believing what she is told. This talk of “ease” is based on what a rational speaker—who wants to avoid social sanction—will probably end up doing. But according to the assurance theorists' TBV, the reason the audience has to believe that p when she is told that p is not that it would be hard for the speaker to avoid social sanction for asserting that p when p is not true. It is rather that the speech act of telling changes the deontic status of the speaker: it places upon her an obligation and it's the fact that the speaker has this obligation that gives the audience a non-evidential reason for believing what she was told. On this, Moran (at least) is rather explicit:

The speaker confers a right of complaint, should what the speaker is told turn out to be false. (Moran, 2005, p. 21)

[I]n testimony in particular the kind of reason for belief that is presented is one that functions in part by binding speaker and audience together, and altering the normative relationship between them. (*ibid*, p.22)

Whether the speaker finds it easy or hard to escape sanction for breaking her obligation is neither here nor there as regards the existence of the kind of reason described by the assurance theorists.

Consequently, plausible deniability threatens the very existence of the reason described by Fricker but it doesn't threaten the very existence of the reason described by the assurance theorists—or if it does, it does not do so in the same way. Given that the speaker has plausible deniability, the threat of social sanction can be avoided by the speaker when she asserts that p when p is not true. Thus there's nothing that makes it hard for her to assert that p when p is not true. So the reason Fricker describes for believing that p when a speaker asserts that p, doesn't exist when the speaker has plausible deniability. But one can give someone a right of complaint even though it's possible (and even if it's easy) for one to convince others that one didn't do so. For example, it's not uncommon for us to commit to doing something or being somewhere by doing things that we could later plausibly say were sarcasm or a joke. But that fact in itself doesn't prove that, on all those occasions, we weren't really making the commitment at all: if we actually went on to convincingly claim that we were joking, when we weren't really joking, then we'd be convincingly lying about what we did, not proving that we didn't make the commitment in the first place. But then, if telling someone that p is to give

someone a right of complaint if p turns out to be false, the fact that one could convince others that one didn't give the audience a right to complain if p in particular turns out to be false, doesn't mean that one didn't do so. Thus plausible deniability doesn't imply that the assurance theorists' reason doesn't exist.

Although Peet wrongly takes it as given that the considerations he sets out against Fricker's TBV also straightforwardly apply to the assurance theorists' TBV, there are amendments that Peet could make to his argument to try and show that the assumption that speakers have plausible deniability implies that assurance theorists fail to avoid scepticism about testimonial beliefs. The most straightforward way to do this is to suppose that if a speaker can easily convince her audience that she didn't commit to the truth of p then she didn't commit to the truth of p:

Commitment Connection (CC)

In order for the speaker to make a commitment to her audience that p is true, by telling the audience that p, it cannot be easy for the speaker to subsequently convince her audience that she didn't intend to tell her audience that p.

If the speaker has plausible deniability, then it is easy for the speaker to subsequently convince her audience that she didn't intend to tell her audience that p. Hence, given CC, the speaker didn't make a commitment to the truth of p to her audience.

But as we have already seen, CC is not plausible. Generally speaking: making a commitment to someone doesn't require that it be very difficult for one to subsequently provide misleading evidence to the effect that one didn't make the commitment. So the connection CC describes between the making of commitments and the ability to convince someone that one didn't make the commitment is too strong. However, if we weaken this connection so that it is not true that if a speaker has plausible deniability then she cannot make commitments to others, then a speaker's possession of plausible deniability is consistent with her being able to make commitments to others. In which case, plausible deniability will pose no direct problem for assurance theorists: their kind of reason (a commitment) doesn't disappear because of plausible deniability.

Alternatively, one might try, not to connect the existence of the audience's reason for believing what is asserted to the speaker's possession of plausible deniability, but instead, to connect the audience's *access* to the audience's reason for believing what is asserted to the speaker's possession of plausible deniability:

Access Connection (AC)

When a speaker asserts that p to an audience, her audience knows that the speaker asserts that p only if the speaker does not have plausible deniability for p.

Although plausible deniability doesn't imply the non-existence of the reason that TBVs say audiences have to believe what they are told, given AC, plausible deniability does imply that audiences typically don't know what's been asserted to them. Given that an audience must know what has been asserted in order to come to have knowledge with the content of what is asserted, AC brings with it the sceptical conclusion that TBV theorists are trying to avoid.¹

But AC is also implausible. To have plausible deniability is not to actually provide, but rather, to be able to provide misleading evidence that leads an audience to give up her belief that the speaker really intended to communicate that *p*. However, generally speaking, if one forms a belief that *p* on the basis of evidence *e*, the mere possibility that someone could provide one with misleading counter-evidence that leads one to give up one's belief that *p* doesn't suffice to show that the evidence upon which one based one's belief that *p* didn't suffice to justify one's belief that *p*—perhaps even to make it knowledge. For example, I could see Mary in the supermarket and form the belief that Mary is in town. Although it doesn't actually happen, it's nonetheless possible that someone subsequently tells me the falsehood that Mary's definitely in Mexico right now. In many such cases, I am likely to abandon my belief. But that doesn't mean that, in actual fact (when there's no misleading counter-evidence), my original reason for thinking that Mary is in town doesn't really provide me with reason sufficient to make my belief that she is in town knowledge. If it did, we would have a problem that applies to a great many of our empirical beliefs, and not just to beliefs based upon the kind of reason described by assurance theorists. Thus, AC is implausible: it implies a very general kind of scepticism regardless of what assurance theorists have to say about knowledge by testimony.

Things are different with Fricker's TBV precisely because she simply defines the act of telling in such a way that the act doesn't happen if the speaker has plausible deniability. So we don't need any further premise (controversial or otherwise) that connects the kind of reason she says we have to believe what we're told to the absence of plausible deniability in order for plausible deniability to pose a problem for that particular TBV.

Perhaps Peet can supplement his argument in some other way. But as things stand, he explicitly provides an argument that links plausible deniability to only Fricker's TBV. So it's upon this TBV that we'll focus in what follows.

4 Responsibility for Mismatches

Let's say that there is a *mismatch* when the proposition that the speaker intends to communicate by uttering a given sentence diverges from the proposition that the audience thinks the speaker intends to communicate. When a mismatch arises, we have available to us three different ways that we can describe the mismatch. With each of these comes a different

¹ Though others might be tempted by it, Peet himself seems unlikely to pursue this possibility. He (forthcoming) has argued that knowledge of what has been asserted (or at least, what has been said) is not required for knowledge by testimony.

allocation of responsibility for the mismatch. Firstly, we can say that the audience misunderstood the speaker: the audience is responsible for the mismatch. Secondly, we can say that the speaker was being unclear, misleading, obtuse etc.: the audience is faultless, the speaker is at fault for the mismatch. This may be so even if the audience has false representations of the context, provided that the speaker is responsible for the audience having such false representations. Thirdly, we can say that “there was a misunderstanding”: we remain neutral as to whether the speaker or the audience is to blame for the mismatch.

With this distinction in hand, we can define a plausible assumption.

Speaker's Responsibility (SR)

When there is a mismatch, if the speaker is responsible for the mismatch then the speaker is responsible for the audience's belief that p (where p is what the audience reasonably took the speaker to assert), even if the speaker intended to communicate some other proposition q .

For example, suppose that I'm giving you directions to my car, and I say, “it's red,” meaning that the car has a red undercoat (whilst its surface colour is blue), but (as is reasonable) you take me to mean that the car has a red surface colour, then I am to blame for you coming to think that the car has a red surface. I am responsible for your belief and can be criticized if (when) it turns out to be unwarranted and false. This is so even though the belief you form is not the one I intended you to form.

SR implies that premise (2) is false. Peet appears to suppose that if the speaker can convince her audience that she didn't intend to communicate what the audience took her to communicate then the speaker can convince her audience that she is not responsible for the audience's belief. But given SR, this isn't true. For if the mismatch arose because of some incompetence on the part of the speaker, then, given SR, the speaker is still responsible for the audience's belief, even if the speaker didn't intend to communicate its content. It's true that the speaker is not obliged to defend the truth of the audience's belief—so she is not responsible for the belief in this sense. But she is responsible for the fact that her audience has this belief, and is blameworthy for that fact, if it's something she did that resulted in the mismatch. So, if she's responsible for the mismatch, she's still responsible for the audience's belief, even given her plausible deniability. We should reject premise (2).

But if SR is a plausible assumption, why does Peet think that the speaker's plausible deniability implies that the speaker can convince her audience that she is not responsible for the audience's resulting belief? Peet doesn't explicitly address this question. But we can piece together a plausible answer by noting some features of Peet's discussion.

Notice first that one subtle but definite feature of Peet's discussion is that Peet repeatedly

chooses, without argument, to use descriptions of mismatches which attribute responsibility to the audience for the mismatch. Here are three examples. Firstly, when Peet describes linguistic communication using context-sensitive expressions, he says:

In order to assign content to a wide range of context-sensitive utterances audiences must rely heavily on their knowledge of the conversational context. Because of this reliance on context there are many ways the recovery of content can go wrong. As a result, speakers are able to make assertions and yet deny responsibility for the proposition asserted, claiming that the audience made a mistake in resolving the context-sensitivity. (Peet, 2015, p. 29)

Peet supposes that if the recovery of content goes wrong—if there is a mismatch—then *speakers* will be able to deny responsibility for communicated contents because the *audience* made a mistake. Secondly, and similarly, consider again the passage in which Peet summarizes the main moves of his argument against TBVs:

[I]f the recovery of an asserted content requires extensive appeal to knowledge context then the speaker will often be able to claim that the audience recovered the wrong proposition, thus disclaiming responsibility for the audience's belief. (Peet, 2015, p. 33)

The extensive appeal to knowledge context is supposed to open up space for inaccurate representations of the context in the knowledge context. Peet supposes here that if there are inaccurate representations in the knowledge context then the speaker is not responsible for the belief that the audience forms, believing it to be what the speaker asserted. The possibility that the speaker is responsible for the inaccurate representations in the audience's knowledge context, is quietly elided from consideration. Thirdly, when Peet (2015, p. 40) does explicitly consider the possibility that *the speaker* has made a mistake about the context, and on this basis has been misleading, he assumes that the speaker can use this mistake to excuse herself. Peet thus applies different standards to speakers and audiences. When an audience makes a mistake about the knowledge context, the audience is to blame for forming a belief other than that which the speaker intended. But when a speaker makes a mistake about the context, this absolves the speaker of responsibility for the audience's belief. Peet seems to be assuming that if there is a mismatch then it's the audience's fault even if it's the speaker who caused the mismatch.

If Peet were right that speakers are generally not responsible for mismatches—*audiences* are—then SR would be false. So our reason for rejecting premise (2) would fall away. Is Peet right to assume what he does about mismatches? He makes no comment about his choice of description. But one clue as to why he describes mismatches in this way can be found in the way that he tends to describe episodes of context-sensitive linguistic communication. Consider once again, the opening passage of Peet's paper:

In order to assign content to a wide range of context-sensitive utterances audiences must rely heavily on their knowledge of the conversational context. Because of this reliance on context there are many ways the recovery of content can go wrong. As a result, speakers are able to make assertions and yet deny responsibility for the proposition asserted, claiming that the audience made a mistake in resolving the context-sensitivity. (Peet, 2015, p. 29)

Here Peet focuses entirely upon the role played by the audience in context-sensitive linguistic communication. The speaker's role is merely to produce an utterance. The rest of the work takes the form of the "recovery" of content—and this work is done exclusively by the audience. Thus construed, in context-sensitive linguistic communication, aside from her utterance, the speaker is a passive objective of interpretation—a hermeneutical puzzle. The audience is the active participant—the interpreter. Peet portrays the speaker and audience in these passive and active roles in episodes of context-sensitive linguistic communication repeatedly (ibid, p.29, p.30, p.31, p.32, p.33, p.35, p.43, p.44, p.45, p.47). If one thinks of context-sensitive linguistic communication in these terms then one will think—given the plausible principle that one cannot be held responsible when one couldn't have done otherwise—that the speaker just cannot be held responsible for mismatches. SR will appear irrelevant to whether the speaker or the audience is responsible for the belief the audience forms in response to apparently being told that *p*, because the speaker cannot be responsible for mismatches, given the passive role speakers play in episodes of context-sensitive linguistic communication. Thus, premise (2) of Peet's argument will not be under threat.

In construing the distribution of work in context-sensitive linguistic communication in this way, Peet is (apparently) following those whose work he builds upon—amongst whom are influential voices in the study of context-sensitive linguistic communication. He cites Fricker (2012) (from whom Peet borrows the thrust of his argument), Carston (2002) and Sperber and Wilson (1986) (who are also cited by Fricker). Fricker, Carston and Sperber and Wilson all describe context-sensitive linguistic communication in a similar, and often the same, way as Peet. For example, on page 1 of Sperber and Wilson's classic text on context-sensitivity, they orient their readers to the intended subject matter thus:

Communication is a process involving two information-processing devices. One device modifies the physical environment of the other. As a result, the second device constructs representations similar to representations already stored in the first device. Oral communication, for instance, is a modification by the speaker of the hearer's acoustic environment, as a result of which the hearer entertains thoughts similar to the speaker's own. The study of communication raises two major questions: first, what is communicated, and second, how is communication achieved? (Sperber & Wilson, 1986, p. 1)

In this passage, the speaker is reduced to being a speaker—talking (i.e. modifying the hearer's acoustic environment) is all she does. It's then unsurprising that the rest of their book focuses entirely upon the matter of how the audience of a speaker can recover the speaker's intended proposition without any systematic attention paid to how the speaker can make her intention available to her audience. Similarly, in Carston's equally influential book on linguistic context-sensitivity, from the very start she takes Sperber and Wilson's research enterprise as her own:

Sperber and Wilson (1986a/1995b) set out to give an account of how hearers (interpreters) reach the intended interpretation of an utterance quickly and with relative certainty, despite the fact that there are, in principle, indefinitely many possible interpretations, all compatible with the linguistically encoded content. The question that has to be addressed is how it is that hearers do not get into even lengthier chains of inference, searching for more and more possible interpretations and comparing them, in a bid to find the 'right' one or the 'best' one. (Carston, 2002, p. 2)

As with Sperber and Wilson, Carston's book subsequently focuses upon the work done by the audience ('how hearers (interpreters) reach the intended interpretation of an utterance quickly and with relative certainty') without any great concern paid to how a speaker could do work that makes her intended content clear to her audience. Where Carston (*ibid*, p.7) does consider the work of the speaker, it is limited to her choice of words.

Despite the venerable (apparent) origin of Peet's construal of context-sensitive linguistic communication, should we accept it? In the next section, I argue that we should not. The speaker can do much more than talk when making apparent the content she intends to communicate. Thus, there will be occasions when she is responsible for the false representations in the audience's knowledge context, not the audience herself. So SR does have application and thus premise (2) of Peet's argument is false.

5 What speakers can do

I proceed in three stages. In the first two, I describe a pair of general features of actions done in the presence of others. In the third, I explain how these two general features together provide a speaker with a way of making her intention clear to her audience which does not reduce to uttering a sentence and meaning something by it.

5.1 Actions in public places

Recognition of our intentional actions seems to come naturally. Elizabeth Anscombe was right when she wrote:

I am sitting in a chair writing, and anyone grown to the age of reason in the same world would know this as soon as he saw me, and in general it would be his first

account of what I was doing; if this were something he arrived at with difficulty, and what he knew straight off were precisely how I was affecting the acoustic properties of the room (to me a very recondite piece of information), then communication between us would be rather severely impaired. (Anscombe, 1957, p. 8)

But the recognition of action can be facilitated and hindered by the person doing it. In his *The Presentation of Self in Everyday Life* the sociologist Erving Goffman seeks to describe the work of making clear what it is that we are doing, when we are in the presence of others: what he (1974, p. 26) calls the 'dramaturgical problems of presenting the activity before others.' Mostly using Goffman's own observations, I will describe seven ways in which this can be done.

Firstly, there are props and settings we can use which will make more obvious what it is that we are doing. For example, if you go for a jog in a suit then you are likely to provoke the question: what is she doing? This contrasts with what is likely to happen if you go for a jog in running-clothes. The same question is much less likely to be raised because it's much more immediately apparent that you are not running from a thief, running after someone, late for a meeting, etc. Similarly, if a window cleaner climbs a ladder to inspect a window before cleaning it, if he has the right over-alls, a parked van with an appropriate logo or description, and his cleaning tools, then it will be all the more obvious what he is doing than if these props were absent. Agents have control over what props they use in the doing of activity and in what setting they attempt to do it. These can be used to make it clearer what is being done.

Secondly, one can infuse one's activity with 'signs which dramatically highlight and portray confirmatory facts that might otherwise remain unapparent or obscure.' (Goffman, 1974, p.40). For example, a man in a suit, who crawls on all fours, in the grass, may, when noticed by another, begin muttering to himself, "where *is* that cuff link?" in order to make it clearer what it is that he is doing. When a baseball umpire has to make a confident, informed judgement, he may emphasise the speed and confidence with which he gives it in order to ensure that impression. A person who is waiting for another in a place where it might not at all be obvious that he is waiting for someone, may glance at his watch more regularly and look around more obviously than he otherwise might, in order 'to give a visible familiar shape to his intention.' (Goffman, 1963, pp. 78–79)

Thirdly, the agent may hide distracting elements of the activity which would put the activity into an ambiguous or misleading light. A wife who tells her husband that she's going home may downplay her tiredness so that it doesn't get misinterpreted as irritation at something done by her husband during the evening. A young philosopher disappointed at her latest rejection, may hide her disappointment so that it doesn't ruin and cast doubt upon the genuine congratulations she gives to her colleague.

Fourthly, distracting mistakes may be practised against to ensure their not happening. A radio

guest may practice his answers to foreseeable questions so as to avoid creating a misleading impression that he cannot correct, if it arises on air. A person going on a date, a job interview, an outing with unknown acquaintances may envisage possible conversations so as to avoid “putting her foot in it” in a way that is not recoverable once the date, interview or outing has commenced.

Fifth, within the confines of conversation, there are various things one can do to make it clearer what it is that one is doing, at a given point in the conversation. Here are three examples. Firstly, a speaker can establish or find a place in the conversation that doesn't obscure what she is doing. Turns in a conversation commonly establish particular tasks for the person selected as the next speaker. For example, if someone asks a question, then the task is to answer the question. As the conversation analyst Emmanuel Schegloff writes:

What is critical here is that the action which some talk is doing can be grounded in its *position*, not just its *composition* – not just the words that compose it, but its placement after a question. Talk after a question invites hearing for how it could be answering... (Schegloff, 2007, pp. 20–21)

The same applies to similar cases. If someone provides an assessment and directs her speech at you then you are given the task of agreeing or disagreeing with the assessment (Pomerantz, 1984). If someone announces news to you then a receipt and an assessment of the news is expected (Terasaki, 2004). If someone makes an offer, then it's to be accepted or rejected. If someone offers an invitation then that sets up the task of accepting or rejecting it. Farewells are to be met with farewells and greetings with greetings. If you try to perform an action at the wrong point in a conversation (e.g. acceptance of an offer just after an assessment has been made) then it will be unclear (or certainly, less clear) what action you are trying to perform. Secondly, speakers can prepare for an action in a conversation over the course of several turns. For example, announcements may be preceded by a series of preliminary questions that make it clearer that an announcement is in the offing: e.g. “did you hear the terrible news?” (Terasaki, 2004). Similarly, if someone wants to be understood to be telling a story, then she will provide a small preface that helps to ensure that others don't cut her off before she finishes (Sacks, 1974). Thirdly, speakers can arrange or make clear relevant aspects of the wider circumstances that bear upon what action they are performing at a given point in the conversation. For example, whether an utterance (regardless of its form) will get heard as a question or a statement depends to a great degree on the known relative epistemic positions of the speaker and her recipient (Heritage, 2012). If you and I both know that you were at the party and I was not, and consequently that you were in a position to know who was there, and I am not, then if I say, “John was there,” it's much more likely to get heard as a question than if our epistemic positions were reversed—in which case it would sound like a statement. One can make one's conversational action—statement or question—clearer by making it clear that there is this asymmetry to our epistemic positions.

So this is the first feature of actions done in the presence of others that I want to register: they can be done in ways that make it more or less obvious what is being done. Although anyone grown to the age of reason may indeed be able to tell by looking what Anscombe was doing, that doesn't mean that when in the presence of others, there are not various ways that Anscombe could make it more or less obvious what she was doing by changing precisely how she was doing it. If she screamed whilst writing, if every 10 seconds she jumped twice and sat down again, if she wrote one word per page, if she carved the words on her table (and didn't write them on her paper), then to any onlookers it would become more obscure than it otherwise would have been what she was doing.

5.2 Instrumental normative constraints

The second feature of actions done in the presence of others that I want to register is as follows. Actions can be related as means to ends: one Φ s by Ψ ing. For example, one can be sailing to the isle of Wight by moving the rudder, one can write a paper by typing on the keyboard, and one can sabotage the wedding by making the announcement. Moreover there are various manners in which one can do actions. One can move the rudder, type on the keyboard, or make the announcement in various different ways. For instance, one can regulate the way one moves the rudder in accordance with the currents of the ocean. One can type on the keyboard so as to generate simple well-formed sentences. The ends one pursues by a given means place instrumental normative constraints on the manner in which one pursues the means: given that one is Φ ing by Ψ ing, and given the circumstances, unless one Ψ s in a given manner, one's Ψ ing doesn't amount to, or contribute towards, one's Φ ing. So insofar as one is Φ ing by Ψ ing, one ought to Ψ in a particular manner. For example, unless one moves the rudder in a particular way, one won't be sailing to the isle of Wight by moving the rudder. Unless, one types in a given way, one won't be writing a paper by typing.

5.3 Speaker's work

If we put together the fact that agents can make their actions more or less clear to others with the fact that one's broader actions place normative instrumental constraints on one's narrower intentions, then we have a method for publicizing some of one's intentions. By making it clear what activity one is engaged in, one can make apparent to others what one's "narrower" intentions ought to be, and hence, given that one is a rational agent, what one's "narrower" intentions in fact are. For example, if others know that you're sailing to the isle of Wight by moving the rudder, and they know enough about the circumstances that you're in to see what that requires of moving the rudder, then they can know how you ought to be moving the rudder, and thus, given that you are a rational agent, they can know how you intend to move the rudder too. In this way, publicizing one's broader projects (using the techniques sketched in section 5.1) is a way of giving access to some of one's (other) intentions. This suffices to show that the depiction of a speaker as ultimately passive in context-sensitive linguistic communication misrepresents the speaker's role in context-sensitive linguistic communication.

Consider, for example, quantifier domain restriction: the phenomenon wherein the set of objects over which a quantifier ranges is shifted by changing the context in which the quantifier is used. Consider a sentence that Peet uses in one of his examples:

(5) Every beer is in the fridge.

Consider the following two contexts in which (5) may get used, and notice how what it is that the speaker is doing by or in uttering (5) changes what it would be most reasonable to take the speaker to intend as the set of objects over which the quantifier phrase “every beer” ranges.

The Wedding Party

Abe and Bob are packing up the bar for the night. Upstairs in the bar, there is a store for extra merchandise: crisps, beers, wines, etc. There's no need for these items to be downstairs and it's safer to leave them upstairs away from customers' wandering hands. But Abe and Bob had ordered-in an extra batch of beer for this Friday because they were hosting a wedding party and wanted to make sure they had enough. Unfortunately, their fridge wasn't big enough to store all this beer at once. So, many of the beers were stored out in the yard, round the back of the bar, where they would be kept cool, protected from customers' wandering hands and where they'd be more accessible than the beers upstairs. But now that the party is over, and they've sold many of the beers from the fridge, there's room enough to put all the beers from the yard into the fridge. This they are doing. At some point in that task, Abe wants to know whether he can go home. Abe asks Bob, “are we done?” Bob replies, “yep, every beer is in the fridge, let's go.”

Here, given that all the relevant information has been sketched, it is most reasonable to take the set over which the quantifier phrase “every beer” ranges to be the set of beers from the yard. It would be unreasonable, given the activity that Bob is pursuing in uttering (5), for Bob to intend to use the quantifier phrase “every beer” so that it ranges over a set of objects that includes the beers upstairs. This is because what Bob is trying to do by uttering (5) is apparent enough: he's trying to provide an answer to Abe's question, and to do that, it would be unreasonable to use “every beer” so that it ranged over the beers upstairs—such would make the utterance of (5) unsuited to establishing whether the two are indeed done for the night. Provided that the broader activity that Bob is pursuing by uttering (5) is clear enough to Abe, and to anyone else who may be interested in how he intends the quantifier phrase, it'll be clear enough how Bob ought to intend the quantifier phrase, and thus, if he is rational, how he is using that quantifier phrase.

If we change the activity that Bob is engaged in when he utters (5) then we can change what it

would be reasonable for him to intend by the quantifier phrase in (5). Consider:

Heatwave

It's the middle of summer. There's a heatwave. The air-conditioning is broken. If any beers in Abe and Bob's bar are left in the heat for too long, their flavours will change in ways that the customers will not like, and they'll be unable to sell the beers—leading to a loss of earnings (which they cannot afford). So this week they don't order so many beers. This way, they can ensure that they have enough space in the fridge to store all the beers in the bar—upstairs and downstairs—in the fridge. There are no beers, this time, outside in the yard. During the process of shifting beers to the fridge, Abe wants to know if they've finished the task at hand. He says to Bob, “are we done?” and Bob replies “yep, every beer is in the fridge.”

This time, given what Bob is doing in uttering (5), it is most reasonable to interpret the quantifier phrase as ranging over all the beers in the bar. For given Bob's apparent aim in using the sentence—viz. to indicate whether the activity they are pursuing is complete—he satisfies that aim only by so using the quantifier phrase.

Thus one can help to make a given intention with which the quantifier phrase is used the one that seems most reasonable, by making apparent factors that would make that intention instrumentally reasonable to have, and by ensuring that distracting factors which might lower confidence in the reasonableness of that intention are hidden from view or not present at all.

In various ways, Bob might have intended something else by the quantifier phrase, in either case. This would be because he was actually pursuing some activity other than that which we take him—given the description of the situation provided—to be pursuing. In *Heatwave*, for example, he might have under-estimated how badly the bar needed to save money, and for that reason, thought that only the upstairs beers need be moved to the fridge, and not all. But he might not have made this apparent in an earlier conversation wherein they had decided to engage upon the task of moving the beers. If indeed, the financial situation is that bad, and Bob had made a mistake about this, then the responsibility for a resulting mismatch would fall to Bob for being unclear and not to Abe for misunderstanding Bob.

We can conclude then that there's work that the speaker does in episodes of context-sensitive linguistic communication, besides uttering words, and hence, that if the speaker fails to do this work well, then mismatches that result from such will be the speaker's fault and not the audience's. It seems then that there's something wrong after all with Peet's exclusive focus on the work done by audiences in context-sensitive linguistic communication. Consequently, we should reject premise (2) of Peet's argument: a speaker's possession of plausible deniability does not imply that the speaker can escape social sanction for her audience's resulting belief.

6 Predictably communicating

Peet (*ibid*, pp. 39-40) discusses and rejects a response to his argument which is similar to the response that I have provided. It's worth explaining how the response I have provided is different from the response Peet considers and why it is not subject to Peet's reasons for rejecting the position he discusses. Consider the following thesis:

Predictably False (PF)

Often, if a speaker makes an assertion which predictably results in a false belief then the speaker will be held responsible for the false belief.

Given PF, we might expect that speakers will not easily predictably communicate that *p*, if *p* were false—for, given PF, they will face social sanction if they do so. If so, then Fricker's reason for believing what one is told might well be found for propositions that one would predictably believe when a speaker makes a given assertion. Thus, the problem that plausible deniability poses for Fricker's TBV could be bypassed. Peet grants PF but he denies that it has this consequence. He gives two reasons for the denial. Firstly, he (2015, p.39) thinks that speakers 'are generally far more careless about what they might predictably communicate than what they clearly and openly communicate.' If speakers are sufficiently careless about what they predictably communicate, then, even if speakers are held responsible for what they predictably communicate, that doesn't seem to stop them from easily predictably communicating false propositions. But then there is no reason of the kind described by Fricker for believing what someone predictably communicates. Secondly, Peet (2015, p.40) denies that PF, even if true, significantly reduces the plausible deniability possessed by a speaker for what she predictably communicates. If so, then, Peet assumes, she will be able to avoid responsibility for predictably communicating false propositions. Hence, once again, we shouldn't expect an audience to have Fricker's kind of reason for believing what the speaker predictably communicates.

Neither of Peet's reasons for rejecting the PF-based view apply to the objection that I'm raising against Peet's argument. Consider first whether Peet's first reason for rejecting the PF-based view has application to the objection I have presented. The objection I have presented is that premise (2) of Peet's argument is false because SR is non-vacuously true. Let's suppose that Peet is right that we don't generally care whether the propositions we predictably communicate are true. That's entirely consistent with it being true that one does care about being blamed for the following: an audience comes to believe that *p*, when one didn't intend to communicate that *p* but the audience nonetheless believed that one intended to communicate that *p*, and, moreover, one is responsible for the discrepancy between one's intention and the belief that resulted. So Peet's first reason for rejection doesn't apply to the objection I've raised. Consider now whether Peet's second reason for rejecting the PF-based view has application to the objection I have presented. The objection I presented was not an

attempt to show that there's any less plausible deniability than Peet thinks there is. It was rather an attempt to show that plausible deniability doesn't imply freedom from social sanction. But the fact that we have plausible deniability for propositions we predictably communicate doesn't provide us with any reason for thinking that plausible deniability implies that the speaker is free of the social sanction to which I have tried to draw attention. But then Peet's second reason for rejecting the PF-view has no application to the objection I have raised against Peet's argument.

Although there may appear to be a similarity between the PF-view and the objection I have raised against Peet's argument, the objections Peet raises against the PF-view do not work against the objection I have raised.

7 Peet's about-turn

Consider Peet's two illustrations of how plausible deniability can allow a speaker to escape social sanction for asserting that p when p is not true. I've already stated Scalar implicature (cf. section 2). The second example is as follows:

Quantifier domain restriction

It is the start of the new year and we have organised a party for the new graduate students. We have a variety of beers on offer, but there are some special craft lagers I want for myself (even though they were brought for the guests). I have stored most of the beer in the fridge, but I have put the craft lagers outside. Sally, one of the new students, arrives and asks where the beer is, so I tell her 'every beer is in the fridge'. Later on you find the craft lagers outside and ask me why I told Sally that every beer was in the fridge. In response to this challenge I might attempt to construct a story along the following lines: I had heard that Sally was a vegan, and I am aware that craft lagers often contain animal products. So when I said 'every beer is in the fridge' I didn't mean every beer we had purchased for the party, I meant every beer which was safe for Sally, as a vegan, to drink. (Peet, 2015, p. 33)

In both of Peet's examples, there's good reason to suspect that even if Peet's story is believed, he (as the speaker in each example) will, in each case, be blamed for the mismatch that is implied by the story he tells. In Scalar implicature, even if Peet managed to convince his fellow conversants that he didn't mean to communicate what they thought he did, insofar as their interpretation of what he intended was reasonable at the time, it seems likely that Peet would be blamed for the mismatch. Put it this way: Peet could assert that p only if p is true, or alternatively, he could be reckless in asserting p , assuming that he can dodge whatever sanction he'll face for asserting that p if p is not true, by telling a story of mismatch like the one in Scalar implicature. But surely the first option is preferable: there's a significant risk that his audience will blame him for the mismatch, if he claims that there's been one. So if he is concerned about social sanction, he's better off asserting that p only if p is true. Similarly, in

Quantifier domain restriction, Sally's interpretation of Peet's utterance seems perfectly reasonable, and Peet's intended content (even though it might have been plausibly something he genuinely meant to communicate) is at least a little delinquent given that Sally had been given no reason to think, in the context, that Peet knew she was vegan (assuming that Peet is not mistaken about this), or that he would be so presumptuous as to try deciding on her behalf what she should and should not be consuming. Peet takes a significant risk of being blamed for the mismatch that his story would imply happened, if he tells such a story. So again, if he cares about social sanction, he's safer from it if he aims to make an assertion only if he's confident of its truth.

Peet seems to concede some of this when he presents the examples:

For now it suffices to note that on one-off occasions we would usually let such matters slide and speaker would usually get away with making such conversational manoeuvres. This is not to say that they wouldn't be criticised for being unclear, simply that they would not be held to what they communicated. (*ibid*, p.32)

When Peet says 'they would not be held to what they communicated,' he means that they would not be expected to defend its truth. He is nonetheless willing to concede that they are liable to criticism for being unclear. But at this point of his paper, Peet supposes that this kind of social sanction (i.e. 'for being unclear' and causing a mismatch) is no help to the TBV theorist against the threat of plausible deniability (see also (*ibid*, p.38)). And yet, later on in his paper, when attempting to rein in the damage plausible deniability does to TBVs, Peet seems to change his mind, and allow that this kind of social sanction is of some help to the TBV theorist:

If there is a clear common use for a term (or common default contextual value), and the audience reasonably assigns such a value only to be met by the mismatch move later on, then the audience is able to maintain that the speaker should have been more explicit about their intention, maintaining that they are partly responsible for the resultant belief. (*ibid*, p.48)

As on p.32, the kind of responsibility for the belief that Peet thinks the speaker has is not an obligation to defend its truth, but rather a blameworthiness for its existence. But here, unlike on p.32, Peet supposes that this kind of responsibility *does* serve to curtail the damage plausible deniability does to TBVs. Admittedly, Peet supposes that being held responsible for a mismatch implies absence of plausible deniability—a false supposition, given the way that Peet defines “plausible deniability.” Nonetheless, if Peet had conceded on p.32 what he concedes on p.48, then the problem he attempts to pose for TBVs wouldn't get started: his examples may be examples in which the speaker has plausible deniability, but that doesn't mean they are examples in which the speaker is not to blame for the supposed mismatch.

Such blame will be an incentive for a rational speaker not to rely upon their possession of plausible deniability when deciding whether to assert something the truth of which she is uncertain. Premise (2) of his argument would have been obviously false when stated.

Moreover, notice that if we reject premise (2) of Peet's argument then we have a much more general response to the problem that plausible deniability poses for TBVs than Peet himself thinks is available. Peet proposes that, mostly, it's only if a speaker repeatedly makes denials in a certain subset of discourses that she will betray that she's lying and thereby reduce the plausibility of these denials. One-off denials—on any discourse—will remain plausible, and so given that plausible deniability suffices for the avoidance of social sanction, Fricker's reason will cease to exist on most occasions in which a speaker doesn't make denials too often. But if we reject premise (2) of Peet's argument on the ground that plausible deniability doesn't suffice to ensure that one is not to blame for a purported mismatch, then we can see how a speaker will be incentivised to avoid making even a one-off (plausible) denial if she cannot ensure that she won't be to blame for the mismatch she will have implied has happened. Whenever it is easier for her to tailor her asserting behaviour to what she can justify than it is for her to be epistemically reckless in her asserting behaviour while remaining confident that she can exercise her capacity for plausible deniability in a way that leaves her blameless for the purported mismatch, she'll be incentivised to keep her asserting behaviour in epistemic check, for fear of social sanction. Thus we can see that Peet's argument fails to show that Fricker's reasons vanish in contexts in which even a one-off plausible denial is available.

8 Avoiding responsibility for mismatches

An improved version of Peet's argument would explicitly acknowledge the need to ensure that speakers are not responsible for mismatches in order for speakers to be able to confidently avoid social sanction for asserting something false, by exercising their capacity to make plausible denials. Such an argument would look like this:

Improved Anti-TBV Argument

- (1) If recovery of an asserted content requires extensive appeal to knowledge context then the speaker will often have plausible deniability.
- (2) If the speaker will often have plausible deniability and, if, often when she does so, she won't make herself blameworthy for mismatches then she will often be able to disclaim responsibility for the audience's belief.
- (3) Recovery of an asserted content often requires extensive appeal to knowledge context.
- (4) Often, when a speaker has plausible deniability, she won't make herself blameworthy for mismatches.

Therefore:

- (5) The speaker is often able to disclaim responsibility for the audience's belief.

The question now is, why should we accept premise (4)? I will not make an earnest attempt to show that premise (4) is true or false. But I will explain why the various considerations which Peet puts forward in favour of the ease with which a speaker can construct a story that would allow her to make a plausible denial do not support premise (4).

Firstly, Peet argues that plausible deniability is common by providing support for the claims that context-sensitivity is widespread and that context-sensitivity requires that audiences rely upon knowledge context. However, the fact that the audience relies upon knowledge context when recovering the speaker's intended content doesn't by itself tell us who is responsible for false representations within knowledge context. If I am representing you, and you are pretending to be a taxi driver, when in fact you have no such credentials, but I think you're a taxi driver, then I have a false representation of you. But I have that false representation because of your efforts at deceiving me, and not because I made a mistake in the way I formed that representation of you. You are to blame for my false representation. Not me.

Secondly, when Peet (2015, p. 39) wants to stress just how easy it is to produce a story that would imply that one didn't intend to communicate what one's audience took one to communicate, he alludes to the fact that what the speaker will have intended will depend upon her mental states and the speaker is epistemically privileged with respect to those states, as compared with the audience. This privilege means that the speaker is more authoritative than the audience on factors that bear upon what the speaker intended. Thus, commonly, the audience will not be able to effectively challenge the speaker's claims about her own intentions. But again, none of this implies that it's easy for the speaker to show that she has no blame for any (purported) resulting mismatch. It doesn't tell us whether the audience was unreasonable in her interpretation of the speaker or instead the speaker failed to make the relevant state of her mind available to her audience.

Thirdly, as we have seen, Peet's illustrations of plausible denials are not illustrations of ways in which the speaker can easily avoid blame for mismatch—in both examples, the speaker takes a significant risk (in putting forward his denial) that he will be blamed for the implied mismatch.

So, the considerations that support the contention that plausible deniability is easy to come by do not provide support for premise (4) of the improved argument. It might be possible to find a way to defend premise (4). The trick would be to find a strategy or procedure for constructing stories that (confidently) leave the speaker blameless, but which don't simply reject the speaker's competence at communicating what she intends—since such a rejection (as we have seen) would be inconsistent with other premises of the argument and would lead to a general scepticism about testimonial knowledge, rather than a problem specifically for TBVs. But, as I said, I won't pursue such a strategy here.

9 Studying linguistic context-sensitive communication

We have seen that influential voices in the philosophical study of context-sensitive linguistic communication have a tendency to think of speakers as speakers only. Such a tendency automatically skews our sense of what speakers and their audiences achieve when context-sensitive linguistic communication goes well. The speaker can modify her words but besides that she can do very little to make plain what content she intends to communicate by producing them: one is left imagining someone who tries to carry out Wittgenstein's experiment (say "it's cold here" and mean "it's hot here") merely by playing with the stress on the words "it's cold here." The audience is then left with the hermeneutical challenge of "inferring" what the speaker's hidden intention might be. Those who seek to describe and explain context-sensitive linguistic communication are then supposed to describe the strategies available to the audience for surmounting this hermeneutical challenge.

But a speaker is not just a speaker. A speaker is a baker giving instructions in the kitchen, or taking orders at the counter. A speaker is not reduced in her efforts at publicizing her intentions merely to choosing between different words and the stress to place upon them. For she can openly embed her speech within broader activities that she dramatizes for her audience in a broader context (one that includes conversational turns, social events, roles etc.). These will place instrumental normative constraints on the contents she could rationally intend by her words when she utters them. The task of the audience in context-sensitive linguistic communication can be significantly facilitated if the speaker dramatizes those actions of hers which bear upon the intentions she has for her words. Given this possibility, one wonders why philosophers of language who have an interest in context-sensitive linguistic communication have generally ignored the research that has been done on these contexts, and the actions done within them. Although informing one's accounts of context-sensitive linguistic communication with psychological research is considered *de rigueur*, informing one's accounts with conversational analytic studies of conversational actions done by talking, and with detailed ethnomethodological studies of the non-conversational actions done by talking, as well as the contexts in which all this happens, are thought to be permissibly substituted with J.L. Austin's or Charles Travis' imagined scenarios. But perhaps this is a mistake. If speakers are incentivized not to exercise their capacity to produce plausible denials by a filigree of obligations to publicize their intentions when speaking, then it is a mistake to think that context-sensitive linguistic communication is achieved prior to or independently of the doing of other things; and if it is, then we cannot study the one without studying the other.

Acknowledgements

This research was funded by the Estonia Research Council through grants PUT632 and IUT20-5. I would like to thank those in attendance of the Summer Work in Progress Seminar at the Department of Philosophy, University of Tartu, for very helpful feedback on an earlier

draft of this paper. I would also like to thank the anonymous referees of this journal for extensive and very helpful comments from which this paper has significantly benefited.

References

- Anscombe, E. (1957). *Intention*. Cambridge, MA: Harvard University Press.
- Carston, R. (2002). *Thoughts and Utterances: The Pragmatics of Explicit Communication*. Oxford: Wiley.
- Fricker, E. (2006). Second-hand Knowledge. *Philosophy and Phenomenological Research*, 73(3), 592–618.
- Fricker, E. (2012). Stating and Insinuating. *Proceedings of the Aristotelian Society Supplementary Volume*, 86(1), 61–94.
- Goffman, E. (1963). *Behavior in Public Places: notes on the social organization of gatherings*. Toronto: Collier-Macmillan.
- Goffman, E. (1974). *The Presentation of Self in Everyday Life*. Hamondsworth: Penguin Books.
- Heritage, J. (2012). Epistemics in Action: Action Formation and Territories of Knowledge. *Research on Language and Social Interaction*, 45, 1–29.
- Hinchman, E. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research*, 70(3), 562–587.
- Moran, R. (2005). Getting told and being believed. *Philosophers' Imprint*, 5(5), 1–29.
- Peet, A. (forthcoming). Testimonial Knowledge without Knowledge of What is Said. *Pacific Philosophical Quarterly*.
- Peet, A. (2015). Testimony, Pragmatics, and Plausible Deniability. *Episteme*, 12(1), 29–51.
- Pomerantz, A. (1984). Agreeing and Disagreeing with Assessments: Some features of Preferred/Dispreferred Turn Shapes. In *Structures of Social Action: Studies in Conversation Analysis* (pp. 57–101). Cambridge: Cambridge University Press.
- Ross, A. (1986). Why do we believe what we are told? *Ratio*, 1, 69–88.
- Sacks, H. (1974). An Analysis of the Course of a Joke's Telling in Conversation. In *Explorations in the Ethnography of Speaking* (pp. 337–353). Cambridge: Cambridge University Press.

- Schegloff, E. (2007). *Sequence Organization in Interaction: a primer in conversation analysis* (Vol. 1). Cambridge: Cambridge University Press.
- Sperber, D., & Wilson, D. (1986). *Relevance: Communication and Cognition*. Oxford: Blackwell.
- Terasaki, A. (2004). Pre-Announcement Sequences in Conversation. In *Conversation Analysis: Studies from the First Generation* (pp. 174–223). Amsterdam and Philadelphia: John Benjamins.