

Knowledge and Implicatures*

Michael Blome-Tillmann

Abstract

In recent work on the semantics of ‘knowledge’-attributions, a variety of accounts have been proposed that aim to explain the data about speaker intuitions in familiar cases such as DeRose’s *Bank Case* or Cohen’s *Airport Case* by means of pragmatic mechanisms, notably Gricean implicatures. This paper argues that all pragmatic explanations of the data regarding ‘knowledge’-attributions are unsuccessful and concludes that in explaining those data we have to resort to accounts that (a) take those data at their semantic face value (*Epistemic Contextualism*, *Subject-Sensitive Invariantism* or *Epistemic Relativism*), or (b) reject them on psychological grounds (Williamson’s *Moderate Insensitive Invariantism*). To establish this conclusion, the paper relies solely upon widely accepted assumptions about pragmatic theory, broadly construed, and on the Stalnakerian insight that linguistic communication takes place against the backdrop of a set of mutually accepted propositions: a conversation’s common ground.

1. Introduction

The recent literature on the semantics of ‘knowledge’-attributions has seen a remarkable proliferation of views according to which the contextual variations in our willingness to ascribe ‘knowledge’ are to be accounted for by means of pragmatic theories. Patrick Rysiew, for instance, has forcefully argued (2001, 2007) that the truth-conditions of ‘knowledge’-attributions are invariant with the context of ascription while the conditions under which ‘knowledge’-sentences are felicitously assertible may vary with the ascriber’s context. Other authors—most notably Peter Unger (1975, 1984), Jonathan Schaffer (2004b), and Jessica Brown (2006)¹—have also argued for what I shall call *Pragmatic Invariantism* (henceforth ‘PI’). According to the pragmatic invariantist, some assertions of ‘knowledge’-sentences give rise to Gricean implicatures that can help account for our intuitions in familiar cases such as DeRose’s *Bank Case* or Cohen’s *Airport Case*.² Thus, what varies with the conversational context is, according to PI, not the proposition se-

* Thanks to Brian Ball, Mikkel Gerken, Jonathan Schaffer, Robert J. Stephens, Tim Williamson and two anonymous referees for this journal.

¹ For further versions of PI see (Davis 2004, 2007), (Hazlett 2009), (Pritchard 2010) and (Black 2005).

² See DeRose and Cohen. I discuss the bank case below in §2.

manically expressed by a particular ‘knowledge’-sentence, but rather the proposition pragmatically imparted or implicated by an utterance of that ‘knowledge’-sentence: context affects assertibility-conditions, not truth-conditions.³

Where precisely are we to locate these pragmatic views in the debate about the semantics of ‘knows’? These theories, given their common aim to account for the data about ‘knowledge’-attributions by means of pragmatic mechanisms, are the main rivals to other established theories in the field, such as *Subject-Sensitive Invariantism* (henceforth ‘SSI’), *Epistemic Contextualism* (‘EC’), *Epistemic Relativism* (‘ER’), and *Moderate Insensitive Invariantism* (‘MII’). Moreover, it ought to be noted that the defenders of pragmatic accounts usually claim, just like the defenders of EC and MII, that their theories offer an attractive resolution of sceptical paradoxes.⁴ Thus, if these pragmatic accounts can withstand closer scrutiny, then they present an interesting approach to the semantics of ‘knowledge’-attributions, not only from a linguistic, but also from an epistemological point of view. Yet, even though the aforementioned pragmatic views have attracted some critical attention in the literature, there is as of now no comprehensive study of their overall prospects. This paper fills that gap in the literature by providing a comprehensive and critical evaluation of all versions of PI currently on the market.⁵

The paper is structured as follows. In Section 2, I briefly introduce the data at the heart of this debate—namely, DeRose’s *Bank Cases*—and explain how EC, SSI and MII aim to account for these data (I shall ignore epistemic relativism for the sake of brevity here). Sections 3 and 4 then critically discuss two different variants of PI: *Sceptical Pragmatic Invariantism* (SPI) and *Moderate Pragmatic Invariantism* (MPI). In Section 5, I develop an argument against MPI that builds on Stalnaker’s (1998; 1978) account of assertion and the notion of the common ground in a conversation. Section 6 discusses a version of PI according to which ‘knowledge’-ascriptions carry *conventional* implicatures, while Section 7 discusses the view that ‘knowledge’-ascriptions trigger what Bach (1994) calls *implicitures*. Section 8 finally sums up the discussion and concludes that the data from the *Bank Cases* cannot be accounted for by means of any type of recognized Gricean or Neo-Gricean pragmatic mechanism.

³ (DeRose 1999) refers to pragmatic invariantist strategies as ‘Warranted Assertibility Manoeuvres’ or simply ‘WAMs’. I prefer, in this paper, the label ‘pragmatic invariantism’.

⁴ To my knowledge, the defenders of SSI and ER do not claim to be able to resolve sceptical paradoxes.

⁵ Critical assessments of PI can be found in (DeRose 1999, 2002, 2009: 117-124), (Halliday 2005), (Leite 2005), and (MacFarlane 2005: , §3.1.1). For responses to the objections in these works see (Brown 2006), (Rysiew 2005) and §5 of this paper. For further critical comments concerning PI, see (Cohen 1999: 60, 2000: 137-138), (Hawthorne 2004: 115-118), and (Stanley 2005: 15).

2. What is Pragmatic Invariantism?

To begin with, let us take a brief look at Epistemic Contextualist, Subject-Sensitive Invariantist, and Moderate Insensitive Invariantist approaches to knowledge and the semantics of ‘knows’. Note, however, that I do not mean to endorse any of these views in this paper. The point of presenting EC, SSI, and MII in this section is merely to provide the conceptual framework needed for the discussion of PI that follows.

To begin with consider David Lewis’s (1996) contextualist definition of the satisfaction of ‘knows’:

(L_{EC}) x satisfies ‘knows p ’ in $C \leftrightarrow x$ ’s evidence eliminates every $\neg p$ -world that is relevant in C .⁶

On Lewis’s view, one’s evidence doesn’t have to eliminate all counterpossibilities to p for one to satisfy ‘knows p ’ in a conversational context C , but only those that are, in a sense further specified by Lewis, *relevant* in C : according to (L_{EC}), the predicate ‘knows’ is a context-sensitive expression, an expression that can change its content with the ascriber’s context. In addition to (L_{EC}), Lewis offers a set of rules specifying precisely which possibilities count as relevant in a given conversational context and it is those rules that are meant to determine how the content of ‘knows’ is influenced by particular contextual factors. Now, the details of Lewis’s account of relevance do not have to concern us further here. What I am interested in at this point is merely the general framework Lewis offers for modelling the semantics of ‘knows’: once this framework is in place, we are in a much better position to characterise the different versions of PI that I want to focus on below.

Before turning to PI, however, let us briefly consider SSI, another view about knowledge and the semantics of ‘knowledge’-attributions that has attracted much attention recently. As defined by its main advocates, (Hawthorne 2004) and (Stanley 2005), SSI claims that whether a subject x knows p partly depends on the subject’s rather than the ascriber’s context.⁷ To illustrate this view further, we can characterise it along lines similar to Lewis’s account. In fact, Lewis’s definition (L_{EC}) can fairly easily be turned

⁶ Cp. (Lewis 1996). Note that Lewis doesn’t use the phrase of a ‘relevant counterpossibility’, but rather that of a counterpossibility that isn’t “properly ignored” in a given conversational context. The difference between Lewis’s and my formulation is merely terminological: a counterpossibility w is relevant in C iff it isn’t properly ignored in C . In this paper I shall use ‘relevant’ and ‘irrelevant’ for reasons of brevity and terminological uniformity. Lewisian accounts have been popular recently. See, for instance, (Ichikawa 2011a, 2011b), for interesting discussion.

⁷ Stanley dubs his version of SSI ‘Interest-Relative Invariantism’, but these details are irrelevant here. The main ideas underlying SSI first occurred in (Fantl and McGrath 2002), but see also (Fantl and McGrath 2009).

into an SSI-account of knowledge—viz., by replacing the phrase ‘in C ’ in (L_{EC}) , which establishes reference to the ascriber’s context, with the phrase ‘in x ’s context’, which establishes reference to the subject’s context. Here is (L_{SSI}) :

(L_{SSI}) x satisfies ‘knows p ’ in $C \leftrightarrow x$ ’s evidence eliminates every $\neg p$ -world that is relevant in x ’s context.

Once we add Lewis’s rules of relevance to (L_{SSI}) we attain a fully functional version of SSI.⁸

Finally, let us consider the linguistic data that EC and SSI are in competition with PI to account for. Here is the *Bank Case* as presented by (Stanley 2005: 3-4):

Low Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Realizing that it isn’t very important that their paychecks are deposited right away, Hannah says, ‘I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning.’

High Stakes. Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, ‘I guess you’re right. I don’t know that the bank will be open tomorrow.’⁹

Our intuitions concerning *Low Stakes* are that Hannah speaks truly when she self-ascribes ‘knowledge’. In *High Stakes*, however, our intuitions are reversed: in *High Stakes* our intuitions are that Hannah speaks truly when denying that she ‘knows that the bank will be open on Saturday’ (henceforth ‘knows O ’).¹⁰

Since Hannah is both subject and ascriber in the above cases, (L_{EC}) -based EC and (L_{SSI}) -based SSI can account for our intuitions along the same lines—namely, by claiming that there are $\neg O$ -worlds uneliminated by Hannah’s evidence that are irrelevant in *Low Stakes* yet relevant in *High Stakes*—the $\neg O$ -worlds at issue being $\neg O$ -worlds in which the bank has changed its hours recently. The reason why these $\neg O$ -worlds are irrelevant in *Low Stakes* but not so in *High Stakes* has, of course, to do with the fact that

⁸ Note that neither Hawthorne nor Stanley explicate SSI along the lines of (L_{SSI}) .

⁹ For the original *Bank Case* see (DeRose 1992: 913).

¹⁰ As is standard procedure in the literature on ‘knowledge’-attributions, I assume that competent speakers’ in general have these intuitions and that they represent data that any account of the semantics of ‘knows’ has to account for.

the stakes for Hannah are low in *Low Stakes* but high in *High Stakes*.¹¹ Thus, both EC and SSI have a *prima facie* powerful explanation of the data generated by *Low Stakes* and *High Stakes*.

How does MII account for the above data? Firstly, note that MII, being both insensitive and invariantist, rejects EC and SSI. It is, however, also moderate—that is, MII claims that speakers in ordinary contexts do in fact know most of the propositions that we intuitively take them to know. To illustrate the view, we can again model MII along the lines of Lewis’s (L_{EC}):

(L_{MII}) x satisfies ‘knows p ’ in $C \leftrightarrow x$ ’s evidence eliminates every $\neg p$ -world that is epistemically relevant.

According to (L_{MII})-based MII, the set of $\neg p$ -worlds that x needs to eliminate in order to know p is invariant across contexts: it changes neither with the subject’s nor with the ascriber’s practical interests or conversational goals. Now, one might wonder how MII will account for the data from *High Stakes*. Not many defenders of MII have addressed this issue. The most detailed approaches, however, can be found in (Gerken 2011, 2012; Nagel 2008, 2010; Williamson 2005a). According to Williamson, for instance, our intuitions in *High Stakes* are deceptive and the “effect of psychological bias” (ibid., p. 234). Hannah speaks falsely in *High Stakes* and we mistakenly judge that she speaks truly because we assign excessive weight to the counterpossibility that the bank has changed its hours recently.¹² Thus, MII has an attractive, simple, and *prima facie* rather powerful explanation of the *Bank Case* data.

Finally, let us turn towards PI’s explanation of the data. According to PI, the proposition semantically expressed by Hannah’s utterance in *High Stakes* is the negation of the proposition semantically expressed by her utterance in *Low Stakes*: PI is an invariantist account of the semantics of ‘knows’. Given that PI is invariantist, however, its defenders owe us an account of the datum that Hannah’s utterances in the *Bank Cases* seem to have the same truth-value, despite expressing incompatible propositions. As is presumably obvious by now, the response the defenders of PI give to this challenge makes essen-

¹¹ According to some theorists (see, for instance, Schaffer 2006) the possibility that the bank has changed its hours recently is relevant in *High Stakes* but irrelevant in *Low Stakes* not because of what is at stake but rather because Sarah, in *High Stakes*, has made that possibility salient by mentioning it, which she has not done in *Low Stakes*. (Schaffer and Knobe 2012) provide further evidence to this effect (see also (DeRose 2011) for interesting discussion). I shall leave aside these details about what exactly determines the relevance of possibilities at a context in this paper and focus on the fact that the data from the bank cases are to be accounted for—independently of what the exact source of our intuitions is.

¹² See (Nagel 2008, 2010) for discussion of Williamson’s view.

tial use of pragmatic concepts such as the notion of a conversational implicature: on the most standard versions of PI at least one of Hannah's utterances in the *Bank Cases* triggers an implicature that plays a crucial role in explaining our intuitions.^{13,14}

Before going into more detail, however, it is worthwhile to provide some structure for the following sections by distinguishing between different versions of PI. Firstly, consider the distinction between what I shall call *Moderate Pragmatic Invariantism* (MPI) and *Sceptical Pragmatic Invariantism* (SPI).¹⁵ According to MPI, Hannah's utterance in *Low Stakes* and most of our everyday 'knowledge'-attributions semantically express truths, while according to SPI our everyday 'knowledge'-attributions semantically express falsehoods: semantically speaking, SPI is sceptical where MPI is moderate. Secondly, note that we can distinguish between versions of PI according to the type of implicature triggered by 'knowledge'-attributions. In particular, we can distinguish between *conversational* and *conventional* versions of the view, depending on whether the relevant implicatures are meant to be conversational or conventional. In what follows, I shall initially restrict my attention to the conversational versions of SPI (Section 3) and MPI (Section 4). Towards the end of the paper, in Section 6, I shall turn to the corresponding conventional versions of those views, while Section 7 is devoted to the view that the phenomenon at issue is what Bach (1994) calls a *conversational implicature*.¹⁶

3. Sceptical Pragmatic Invariantism

One of the most systematic and rigorous accounts of PI in the literature—namely, the one defended by Jonathan Schaffer (2004b)—is a version of *Sceptical Pragmatic Invariantism* (SPI). The central idea underlying SPI goes back to (Unger 1975), but it ought to be emphasised here that both Schaffer and Unger have, in more recent writings, surrendered

¹³ It is sometimes argued that PI is theoretically especially economical, because it explains the data from the bank cases by means of pragmatic mechanisms that are in place already and independently motivated (see, for instance, (Hazlett 2009: 609), (Schaffer 2004b)). Such arguments, however, mistakenly assume that the theoretical mechanisms employed by the opponents of PI are *ad hoc* and not independently motivated. It is far from clear whether this is the case.

¹⁴ For further, alternative views on the semantics of 'knows' see the relativist approach defended by (MacFarlane 2005) and the speech act theoretic account proposed by (Turri 2010, 2011).

¹⁵ I have borrowed the terms 'sceptical' and 'moderate invariantism' from (Hawthorne 2004: 113).

¹⁶ Further explanations of the *Bank Case* data have been proposed in the literature. (Bach 2005), for instance, proposes that Hannah does not 'know' in *High Stakes* because she ceased to believe the proposition at issue. However, Bach's view is controversial, because it cannot readily account for familiar data from *High Ascriber/Low Subject* cases—that is, cases in which the subject is in a low stakes situation and clearly has an outright belief that the bank will be open tomorrow, while the ascriber is in a high stakes situation and so can seemingly truthfully assert 'Hannah doesn't know that the bank will be open tomorrow'. While Bach's view is interesting and deserves an in-depth discussion, restrictions of space do not allow me to go into more detail here.

SPI in favour of other positions.¹⁷ However, Schaffer's earlier work on PI is extremely interesting and rich in detail and therefore warrants closer examination.

According to (Schaffer 2004b), the semantics of 'knows' is given by the following principle:

(L_{SPI}) x satisfies 'knows p ' in $C \leftrightarrow x$'s evidence eliminates every $\neg p$ -world.¹⁸

Given this principle, the predicate 'knows' does not change its content with context and the truth-conditions of 'knowledge'-ascriptions are invariant: (L_{SPI}) gives an invariantist account of the semantics of 'knows'. Moreover, note that, given our Lewisian notion of evidence, (L_{SPI}) entails radical scepticism about the external world.¹⁹ According to (L_{SPI}) we do not know that we have hands, for knowing that we have hands would require our evidence to eliminate *every* world in which we do not have hands, including worlds in which we are handless brains in vats. By definition, however, worlds in which we are handless brains in vats resist elimination by our evidence: we have the same Lewisian evidence in those worlds as we do in actuality. Thus, (L_{SPI}) entails that the sentence 'I know O ' as uttered by Hannah in *Low Stakes*, as well as most of our everyday 'knowledge'-attributions, semantically expresses a falsehood: SPI is not only an invariantist but also a sceptical view. If it really is the case that Hannah's utterance in *Low Stakes* expresses a falsehood, however, the question arises as to how SPI accounts for our intuition that she speaks truly?

SPI's response is, of course, that Hannah's utterance in *Low Stakes* triggers a true implicature. To be more precise, SPI has it that, for all conversational contexts C , utterances of ' x knows p ' semantically express (12) while implicating (13), at least in quotidian or low-stakes contexts:²⁰

- (1) x 's evidence eliminates every $\neg p$ -world.
- (2) x 's evidence eliminates every $\neg p$ -world that is relevant in C .

Thus, the proposition that Lewis claims to be the semantic content of a 'knowledge'-attribution in a given context C is what the defender of SPI claims to be its pragmatically

¹⁷ Schaffer now defends contrastivism, a version of epistemic contextualism that is not subject to some of the major objections to standard contextualism. See (Schaffer 2004a, 2005).

¹⁸ According to Schaffer (2004b), it is the subject who eliminates a counterpossibility, not her evidence. I shall ignore this difference and assume that a subject can eliminate a possibility iff her evidence eliminates it.

¹⁹ See (Lewis 1996: 224) for his notion of evidence. For Lewis, our evidence consists in the totality of our experiences and memory states narrowly individuated.

²⁰ Schaffer (2004b) does not discuss the issue whether the implicature at issue is triggered in all conversational contexts or only in non-sceptical ones.

conveyed content in *C*: what Lewis puts into the semantics, SPI puts into the pragmatics. Moreover, note that, according to SPI, Hannah semantically expresses a falsehood when uttering ‘I know *O*’ in *Low Stakes*, but conversationally implicates a truth—namely, the proposition that her evidence can eliminate all $\neg O$ -worlds that are relevant in *Low Stakes*. At first glance, SPI seems to offer an attractive explanation of the data from both *Low Stakes* and *High Stakes*.

But let us turn towards a more critical evaluation of SPI. Firstly, note that any pragmatic explanation of our apparently semantic intuitions must give an explanation from general conversational principles, such as Grice’s conversational maxims.²¹ Here is the *Generality Constraint*:

Generality Constraint (GC):

Postulate conversational implicatures only if the implicature’s presence can be explained from general conversational principles.

All parties involved consider GC uncontroversial and I shall therefore refrain from discussing it further. But does SPI conform to this methodological constraint? On Schaffer’s view, SPI’s explanation of the bank case data is given by means of entirely general, independently motivated pragmatic principles. To be more precise, Schaffer (2004b) argues that the implicatures postulated by SPI are triggered by violation of Grice’s *First Maxim of Quality*. Here is Grice’s formulation of that maxim:

*Quality*₁: Do not say what you believe to be false.

Given *Quality*₁, Schaffer explains that ‘knowledge’-attributions such as Hannah’s in *Low Stakes* are comparable to other violations of *Quality*₁—namely, to cases of exaggeration or hyperbole. On the view at issue, Hannah’s utterance of ‘I know *O*’ in *Low Stakes* is analogous to utterances of sentences such as ‘I’ve been waiting for ages’ or ‘I’ve heard that a million times’: the speaker conveys a true and informative proposition by exaggerating and overstating what is actually the case. SPI has a *prima facie* plausible and appealing explanation of the bank case data that clearly adheres to GC: *Quality*₁ is a pragmatic maxim that is independently motivated and a familiar and indispensable tool in the Gricean toolbox.

There are, however, convincing arguments against the view just explicated. As MacFarlane (2005: 206) has pointed out, competent speakers are in fact in a position to

²¹ Cp. (DeRose 2002) and (Stanley 2005: 15) for this point.

easily detect cases of exaggeration and hyperbole as cases of non-literal speech. Consider (3):

(3) I've been waiting for ages.

Competent speakers of English can easily be made aware of the fact that (3) semantically expresses a falsehood on any occasion of ordinary use. In fact, competent speakers are quick to realise that utterances of (3) convey that the speaker has been waiting for what counts as a long time in the relevant context rather than that the speaker has literally been waiting for ages: speakers immediately identify (3) as a trope.²²

To illustrate this phenomenon further, consider the following dialogue:

(4) *Exaggeration:*

A: I've been waiting forever!

B: That's absurd; no human being could ever wait that long.

A: Oh c'mon—I was just exaggerating!

As the felicity of this dialogue demonstrates, competent speakers recognise that A's first utterance is an instance of non-literal speech: if standard cases of hyperbole such as (3) were non-transparent to competent speakers, then they would not consider A's second remark natural and felicitous. Rather, they would expect A to reinforce her earlier remark by repeating her previous utterance. Furthermore, it is worthwhile emphasising that the hyperbole in (3) and (4) is not unusual in this respect: other cases of hyperbole give rise to the same phenomena. Consider, for instance, the phrases under (5):

(5) *apologise a thousand times, cry a flood of tears, spend tons of money, be as old as the hills*

Even though the phrases in (5) are used non-literally on every ordinary occasion of use, they are nevertheless hyperbolic in an obvious way: competent speakers quickly identify them as figurative when asked whether they are speaking literally—dialogues analogous to (4) can easily be construed for the phrases in (5). Consequently, competent speakers detect without difficulty the non-literality of even highly formulaic and idiomatic instances of hyperbole.²³

With respect to 'knowledge'-ascriptions, however, the situation is rather different. If our everyday 'knowledge'-attributions are in fact hyperbolic, then competent speakers are unaware of their non-literality. To see this consider the following dialogue:

²² See also (MacFarlane 2005: , §3.1.1) on this point.

²³ It is thus not quite right that “[h]ighly formulaic tropes are particularly non-obvious” or that “[o]ur linguistic intuitions provide evidence for acceptability, and do not discriminate between semantic and pragmatic sources” (Schaffer 2004b). See also (Rysiew 2001: 496) for similar claims.

- (6) *Exaggeration_K*:
 H: I know that the bank will be open tomorrow.
 S: #That's absurd; no human being could ever know that.
 H: #Oh c'mon—I was just exaggerating!

Certainly, both S's response and H's second utterance in (6) are conversationally defective in ways in which A's and B's utterances in (4) are not defective. S's response in (6), for instance, seems puzzling and unmotivated: it is not at all clear that S is addressing the alleged literal content of H's first utterance. And H's admission to have been speaking hyperbolically seems equally out of line: there is no clear sense in which H's first utterance can be construed as being hyperbolic. If SPI were correct, however, we should expect these utterances to be conversationally unobjectionable, as seen in (4). Consequently, there is convincing evidence that 'knowledge'-attributions are not instances of hyperbole.

Leaving aside the view that everyday 'knowledge'-attributions are hyperbolic, let us consider the view that they are nevertheless systematic violations of Grice's maxim of *Quality*₁. Could ordinary 'knowledge'-attributions be assimilated to other types of violation of *Quality*₁, such as cases of metaphor and irony?²⁴ As the following examples demonstrate, similar problems arise for these views. Consider the case of irony first:

- (7) *Irony*:
 A: John is an atheist.
 B: Yes, and so is the pope.
 A: That's absurd; the pope is clearly *not* an atheist.
 B: Oh, c'mon—I was just being ironic!
- (8) *Irony_K*:
 H: I know that the bank will be open tomorrow.
 S: #That's absurd; no human being could ever know that.
 H: #Oh, c'mon—I was just being ironic!

As (7) shows, competent speakers are, in standard cases, in a position to detect irony. From (8), however, it is obvious that similar data are not observed with respect to 'knowledge'-attributions. The view that 'knowledge'-ascriptions are instances of ironic speech is hopelessly implausible.

How about the view that utterances such as Hannah's in *Low Stakes* are instances of metaphorical speech? Consider the following discourses:

- (9) *Metaphor*:
 A: She's made of stone, this girl.
 B: That's absurd; human beings aren't made of stone.
 A: Oh, c'mon—I was speaking metaphorically!

²⁴ See (Grice 1989).

- (10) *Metaphor_K*:
 H: I know that the bank will be open tomorrow.
 S: #That's absurd; no human being could ever know that.
 H: #Oh, c'mon—I was speaking metaphorically!

Again, the examples strongly suggest that ordinary 'knowledge'-ascriptions are not cases of metaphorical speech. Summing up, the claim that Hannah's utterance in *Low Stakes* and our everyday 'knowledge'-attributions more generally are systematic violations of *Quality*₁, and thus similar to cases of hyperbole, irony or metaphor, has exceedingly implausible consequences.

Independently of Schaffer's version of SPI, Wayne Davis (2007) has suggested another approach on which our ordinary 'knowledge'-ascriptions are violations of *Quality*₁. Davis's approach is, on the face of it, more promising, but I shall argue that it suffers from the same crucial defect as all versions of SPI. What exactly is Davis's approach? According to Davis (2007: 395) our utterances of the form 'x knows p' "are used loosely to implicate [that x] is close enough to knowing p for contextually indicated purposes" or that "the condition for asserting 'p' and using it in practical reasoning are satisfied." Davis's view, however, is troubled by recalcitrant data very similar to the previous versions of SPI. Consider firstly typical cases of loose use such as (11), (12), and (13):

- (11) *Loose Use_H* (context: looking at a map of Europe):
 H: France is hexagonal.
 S: Well, strictly speaking that's not true.
 H: Oh c'mon—I was just speaking loosely!
- (12) *Loose Use₃* (context: the clock shows 3.03pm):
 H: It's 3pm.
 S: Well, strictly speaking that's not true.
 H: Oh c'mon—I was just speaking loosely!
- (13) *Loose Use_W* (context: Wayne lives in Springfield, a suburb of Washington):
 H: Wayne lives in Washington.
 S: Well, strictly speaking that's not true.
 H: Oh c'mon—I was just speaking loosely!

As with the previous examples, we observe the by now familiar datum that recognized cases of violation of *Quality*₁ are easily discovered as cases of non-literal speech and, in this particular case, as instances of loose use. With respect to 'knowledge'-ascriptions, however, this is not the case. Consider (14):

- (14) *Loose Use_K* (context: as in *Low Stakes*):
 H: I know that the bank will be open tomorrow.
 S: Well, strictly speaking that's not true.
 H: #Oh c'mon—I was just speaking loosely!

As (14) demonstrates, Hannah’s utterance of ‘I know that the bank will be open tomorrow’ in *Low Stakes* is not an instance of loose use. For if it were, then Hannah should be in a position to become aware of the fact that she had been speaking loosely and her response of ‘I was just speaking loosely’ should be felicitous.²⁵

At this point, the defender of SPI might be tempted to move away even further from Schaffer’s original version of SPI and maintain that the data from *Low Stakes* can be explained by recourse to one of the remaining Gricean conversational maxims. According to Grice (1989: 26-27), there are four general types of conversational maxims—namely, what he calls *Quantity*, *Quality*, *Relation* and *Manner*:

- Quantity*₁: Make your contribution as informative as is required.
- Quantity*₂: Do not make your contribution more informative than is required.
- Quality*₂: Do not say that for which you lack adequate evidence.
- Relation*: Be relevant.
- Manner*: Be perspicuous.²⁶

I take it to be fairly obvious that none of these maxims can do the work required by the defenders of SPI. Let me be brief: there are simply no convincing reasons to accept the view that Hannah’s utterance in *Low Stakes* is, when taken literally, overly informative or uninformative, evidentially unfounded, conversationally irrelevant or overly imprecise, prolix or convoluted.²⁷ To the contrary, Hannah’s utterance in *Low Stakes* is a clear case

²⁵ It might be objected at this point that the data are not as clear and straightforward as suggested here. Consider the following case:

- (i) *Loose Use_P* (context: looking at a patchwork piece on a bedspread)
 - H: That piece is hexagonal.
 - S: Well, strictly speaking that’s not true.
 - H: ?Oh c’mon—I was just speaking loosely!

Given (i), it might seem that, unlike in the cases presented in the main text, the looseness of use is sometimes not immediately obvious. If this was correct, Davis could, in response to my examples, take the view that Hannah’s ‘knowledge’-claim in *Low Stakes* is also an example of loose use that is not obvious, but that can be revealed by giving some further information: say that Hannah cannot rule out the possibility that the bank has changed its hours recently. Once such additional information is added, it might seem quite natural for H in (14) to reply that she was just speaking loosely. Now, I am not entirely certain whether H’s second utterance in (i) is in fact as infelicitous as H’s second utterance in (14). In fact, I should mention that I have a rather strong tendency towards assessing (i) along the same lines as (11)-(13)—that is, as entirely felicitous. However, I realize that intuitions may vary concerning (i), and I shall therefore, in the paragraphs to follow, supplement the above argument against Davis’s view with a theoretical objection to all version of SPI (including Davis’s loose use account) that is independent of the above data.

²⁶ Grice considers the Maxim of Manner a “supermaxim”, that has the maxims ‘Avoid obscurity of expression’, ‘Avoid ambiguity’, ‘Be brief’ and ‘Be orderly’ as submaxims. There is no need to discuss these maxims individually here, as they are epitomised in the maxim of manner as presented above.

²⁷ In fact, there is a strong argument against the view that *Relation* can do the work required by SPI. Note that the alleged semantic content of ‘I know *O*’ in *Low Stakes*—namely, the proposition

of unproblematic and felicitous literal speech and this is precisely how we ought to interpret the utterance.

As a final datum in support of this view, note that S's utterance in (6), (8) and (10) makes no conversational sense—independently of H's response. To make this point more obvious, consider (15), which is a slightly altered and abridged version of the discourses previously discussed:

(15) *Semantic Inaccessibility in Low Stakes*:

H: I know that the bank will be open tomorrow.

S: #That's false; no human being could ever know that.

As the infelicity of S's comment in (15) demonstrates, what SPI claims to be the literal semantic content of Hannah's 'knowledge'-attributions in *Low Stakes* is unavailable to competent speakers; independently of which maxim the defender of SPI claims to be violated. Thus, even though some versions of SPI may have a *prima facie* plausible response to the generality constraint, they fare rather badly with respect to the actual data of allegedly analogous instances of implicature.

Before moving on, however, let me present a more principled argument against SPI that does not rely on any linguistic data. To begin with, remember that according to SPI we never use the expression 'know' with its literal meaning in everyday contexts.²⁸ Rather, according to SPI, we use the expression 'know' in everyday contexts with its non-literal meaning—whether this is a hyperbolic, metaphorical or contextually loosened meaning. According to SPI, 'know' thus exemplifies what Bach (1987: 79ff) calls "standardized non-literality": an expression that is standardly used non-literally is only rarely—or sometimes even never—used with its literal meaning. To get a better grasp on the phenomenon, recall the expressions in (5), repeated here for convenience:

(5) *apologize a thousand times, cry a flood of tears, spend tons of money, be as old as the hills*

It is fairly obvious that each of these phrases is standardly used non-literally: when we utter, in everyday contexts, the sentence 'Jane has apologized a thousand times' we usually do not mean that Jane has apologized a thousand times. Rather, we mean, say, that she

that Hannah's evidence eliminates every $\neg O$ -world—entails the allegedly implicated content—namely, the proposition that Hannah's evidence eliminates every conversationally salient $\neg O$ -world: the conversationally salient $\neg O$ -worlds are a proper subset of the $\neg O$ -worlds. Thus, if the latter proposition is relevant, how could the former, logically stronger proposition be irrelevant?

²⁸ This is even so when we utter negative 'knowledge'-ascriptions: if we were to use 'know' literally in negative 'knowledge'-ascriptions, then those negative 'knowledge'-attributions should appear trivial and uninformative to us.

has apologized more often than was expected. Similar considerations apply for the remaining phrases.

One might wonder how it is possible that the phrases in (5) have the literal meanings they have despite being standardly used to convey different contents. Note that this is not a trivial question, for if we adopt the eminently plausible view that the linguistic meaning of an expression is determined by its use in a speech community—or by the conventions governing its use—then there seems to be no room for a difference between literal meaning and standard speaker meaning. However, note that it is not really surprising that there is such a difference in the above cases. The phrases in (5) have a literal meaning distinct from what they are standardly used to convey in virtue of being composed out of simpler expressions: their literal meanings are functions of their syntax and the literal meanings of their ultimate constituents. This idea also allows us to give an explanation of why ordinary speakers recognize the literal meaning of the phrases in (5) straight away: they grasp it by means of standard compositional semantic interpretation.

Let us return to SPI and the semantics of ‘know’. How can we explain the alleged fact that ‘know’ is standardly used non-literally? Note that the above explanation that accounts for the standardized non-literality of the examples in (5) cannot be employed here, for ‘know’ is obviously not syntactically complex and thus decomposable in a way that leaves room for a difference between literal meaning and standard speaker meaning. If this is so, however, it follows that the standard or ordinary use of ‘know’—or the conventions governing that ordinary use—must be taken to determine its literal meaning. For otherwise, what should we take to determine the linguistic meaning of ‘know’, if not the use of competent speakers of English? If we accept the view that standard use determines standard (or literal) meaning, then it seems that SPI’s claim that our standard uses are non-literal must be false.

The sceptical invariantist might at this point be tempted to argue that her own expert usage of the word ‘know’ has a privileged status over ordinary speakers’ uninformed usage of the word—similarly to how we might be tempted to think that the geometer’s expert use of ‘hexagonal’ has a privileged status over our everyday loose use. But it is no doubt implausible and excessive to ascribe such linguistic authority to the sceptical invariantist, given the controversial status of SPI in the literature: why should we not instead defer to the ‘expert judgments’ of epistemic contextualists, subject-sensitive invariantists,

or moderate insensitive invariantists?²⁹ It thus seems that the defenders of SPI have no plausible story as to how the predicate ‘know’ has acquired the literal meaning that they claim it has and their claim that the semantics of ‘know’ is not in fact anchored in the speech behaviour or the ordinary usage of ‘know’ by competent speakers of English becomes increasingly *ad hoc* and implausible. I shall thus leave behind SPI and its various pragmatic incarnations, for the view is, as we have seen, not only problematic from an empirical perspective but rather also dubious from a more theoretical point of view.

4. Moderate Pragmatic Invariantism

Let us turn to *Moderate Pragmatic Invariantism* (MPI), the view that our everyday ‘knowledge’-attributions semantically express truths. MPI has been rather popular in the recent literature—different versions of the view have, amongst others, been proposed by Brown (2006), Rysiew (2001, 2007), Pritchard (2010), Hazlett (2009) and Black (2005). To begin with, consider the following principle, which I shall use as a model for the semantics MPI assigns to ‘knowledge’-attributions:

(L_{MPI}) *x* satisfies ‘knows *p*’ in *C* ↔ *x*’s evidence eliminates every ¬*p*-world that is *epistemically relevant*.³⁰

In addition to this principle, the defenders of MPI explain that what qualifies as an *epistemically relevant* alternative does not vary with the conversational context. However, they admit that there is a different notion of a *con conversationally salient alternative* that is context-sensitive. Thus, according to the view at issue, ‘knowledge’-attributions semantically express (16) but—at least in high-stakes situations—con conversationally implicate (17):

(16) *x*’s evidence eliminates every ¬*p*-world that is *epistemically relevant*.

(17) *x*’s evidence eliminates every ¬*p*-world that is *con conversationally salient* in *C*.³¹

Since, according to MPI, a world can be conversationally salient in a context without being epistemically relevant, the content semantically expressed by ‘knowledge’-attributions can, in high-stakes contexts, be considerably easier to satisfy than the content conversationally implicated.

²⁹ Note also that, since Schaffer has given up SPI, only Wayne Davis is left defending a version of SPI. I take it to be implausible that Wayne Davis is the lonesome expert determining all by himself the literal meaning of ‘know’ in English.

³⁰ Rysiew (2001, 2007), Brown (2006) and Hazlett (2009: 605) accept (L_{MPI}) or a close version of it. Pritchard (2010) and Black (2005) do not formulate their views in terms of the elimination of counterpossibilities, but (L_{MPI}) is nevertheless firmly within the spirit of their views.

³¹ Rysiew (2007: 488) uses the terms “relevant” and “salient”, while in his (2007: 637) he discusses possibilities that are “considered” and “worth taking seriously”.

How does this view account for the *Bank Case* data? Consider *Low Stakes* first. According to MPI, the epistemically relevant and the conversationally salient possibilities coincide in *Low Stakes*. Thus, in *Low Stakes*, Hannah’s utterance of ‘I know *O*’ both semantically expresses and conversationally implicates a truth. In *High Stakes*, however, an utterance of ‘I know *O*’ is conversationally misleading, for in *High Stakes* that utterance—even though having the same semantic content—conversationally implicates that Hannah’s evidence eliminates every $\neg O$ -world that is conversationally salient in *High Stakes*, including those $\neg O$ -worlds in which the bank has changed its hours recently. Given that Hannah’s evidence does not eliminate those particular $\neg O$ -worlds, the proposition pragmatically imparted by Hannah’s assertion of ‘I know *O*’ in *High Stakes* is false. MPI seems to be able to account for the *High Stakes* data.

However, it is worthwhile taking a closer look. As DeRose (1999: , §11; 2002: 191) and MacFarlane (2005) have pointed out, the datum from *High Stakes* was not that potential utterances of ‘I know *O*’ by Hannah would convey falsehoods—even though this is, of course, a datum that we ought to explain, too. Rather, the datum from *High Stakes* was that Hannah’s *negative* ‘knowledge’-attribution—that is, her utterance of ‘I don’t know *O*’, conveys a truth in *High Stakes*. According to (L_{MPI}), however, that utterance semantically expresses a falsehood—namely, the proposition that Hannah’s evidence does not eliminate all *epistemically* relevant $\neg O$ -worlds. And this datum is not accounted for by claiming that a potential positive ‘knowledge’-attribution would have triggered a false implicature.

To resolve this issue, the defenders of MPI might want to add the following principle to their account:

Converse Implicatures (CI):

If an utterance of a sentence *S* conversationally implicates *p* in *C*, then an utterance of $\neg S$ conversationally implicates $\neg p$ in *C*.

This principle, however, is exceedingly implausible, for virtually all paradigmatic cases of conversational implicature are counterexamples to *Converse Implicature*.³²

³² Brown (2006: 420/425) seems to assume that CI holds at least for those implicatures that are triggered by violations of *Relation*. For illustration, Brown discusses Grice’s example of a man standing besides a car that has obviously broken down asking a passer-by:

- (1) A: Is there a garage nearby?
 B: Yes, there’s one around the corner.

In the envisaged case, B’s utterance conversationally implicates that the nearby garage is open. Brown then considers the following discourse, in which B utters (roughly) the negation of (1B):

What are the moves available to the defender of MPI at this point? Black (2005: 336), Hazlett (2009: 616-619) and Pritchard (2010: 89-90) argue that we should simply reject the data from *High Stakes*. Hazlett considers two different ways to do this—firstly, the idea that we should reject the data from *High Stakes* because Hannah “has made a philosophical mistake” (Hazlett 2009: 619) when uttering ‘I don’t know *O*’ and, secondly, the idea that Hannah’s utterance in *High Stakes* is felicitous only because it involves unusual stress on ‘know’. Both of these approaches are *ad hoc* and unmotivated. Firstly, it is simply not correct that Hannah’s utterance in *High Stakes* is only felicitous if it contains a special stress on ‘know’. Secondly, the explanation that Hannah made a mistake and is simply wrong in *High Stakes* is crucially incomplete: it needs to be supplemented with a story about why Hannah—and competent speakers more generally—are prone to making the mistake at issue. Once such a story is in place, however, we will have an explanation of the *High Stakes* data—including the infelicity of potential positive ascriptions in *High Stakes*—that is independent of the postulation of conversational implicatures: rejecting the data from *High Stakes* renders superfluous a pragmatic explanation of those data, and Hazlett’s version of MPI ultimately collapses into an account resembling Williamson’s MII.

A similar approach to the problem of accounting for the data from *High Stakes* can be found in (Black 2005: 336) and (Pritchard 2010: 89-90): both Black and Pritchard do not seem to share the intuition that Hannah’s utterance in *High Stakes* conveys a truth and they therefore reject the data. This strategy is subject to the same objection as Hazlett’s: Black and Pritchard need an explanation for why the remaining theorists in the field—including the remaining defenders of PI—have a rather strong but mistaken intuition that Hannah’s utterance in *High Stakes* conveys a truth even though, on their view, it

-
- (2) A: Is there a garage nearby?
 B: No, there’s no garage nearby.

Brown claims that B’s utterance in (2) conversationally implicates (or “pragmatically conveys”)—in line with *Converse Implicatures*—that there is no open garage nearby. Now, I agree that B’s utterance in (2) conveys that there is no open garage nearby. But this is so because (2B) semantically entails—and thus semantically conveys—that proposition. The fact that the phenomenon at issue is not a conversational implicature can be demonstrated further by the cancellability test: ‘There’s no garage nearby; but there’s an open garage nearby’ is contradictory in all contexts, because it semantically expresses a contradiction. Thus, B’s utterance in (2) does not conversationally implicate that there is no open garage nearby and Grice’s garage example does not lend support to *Converse Implicatures*. To the contrary, it is, in fact, a counterexample. Now, while this problematic consequence can be circumvented by claiming that the actual implicature in the garage case is the proposition that A can get fuel around the corner, which is not semantically entailed by (2B), it is nevertheless the case that other paradigmatic cases of Relation implicatures (such as Grice’s famous letter writer) are counterexamples to (CI).

conveys a falsehood. But again, once such an explanation is in place, there are no data left to be explained by means of conversational implicatures: rejecting as mistaken the intuition that Hannah conveys a truth in *High Stakes* renders MPI's appeal to conversational implicatures superfluous.³³

The only detailed and custom-made response to the data from *High Stakes* that is wholly in the spirit of MPI can be found in (Rysiew 2001, 2007) and (Brown 2006). For the sake of brevity I shall, in what follows, focus on Rysiew's version of the view, but the following considerations apply equally to the account proposed by Brown (2006). Rysiew argues not only that Hannah's potential utterance of 'I know *O*' in *High Stakes* violates *Relation*; he also argues that Hannah's actual utterance in *High Stakes* does so. Consider the following passage from (Rysiew 2001: 491):

Witness, then, how natural it would be for [Sarah in *High Stakes*] to reason as follows: '[Hannah] has just said 'I [...] don't know [*O*].' And she has said this in response to my raising a doubt as to whether she can really be so sure—after all, banks, as I've just reminded her, do change their hours. Presumably (on the assumption that she's conforming to [the Cooperative Principle], she wouldn't have said what she has unless she thought that there were possibilities incompatible with the bank's being open tomorrow—specifically, that it has recently changed its hours—that she could not rule out. For to say, 'I don't know...' is to indicate that one isn't sure. On the other hand, if this isn't what [Hannah] meant in saying that she didn't know the bank would be open, then I'm not quite sure what she did mean. On the assumption, then, that her conversational contribution is to the point and has been made in light of what I've just said, that must be what she intends to communicate—viz., that she can't rule out the possibility that the bank has just recently changed [its] hours.'³⁴

Rysiew has a *prima facie* very plausible explanation of the data in *High Stakes*.

However, let us take a closer look at the idea that Hannah's utterance in *High Stakes* triggers a *Relation* implicature. Remember from Section 3 that any account that postulates a conversational implicature must respect the *Generality Constraint* (GC). Now, it has been argued before that MPI does not have an explanation of the pragmatic mechanisms at issue that is derived from fully general conversational principles such as the Gricean maxims. DeRose (2009), for instance, objects to MPI on these grounds. To

³³ For critical discussion of the data in the bank cases based on empirical studies see (May et al. forthcoming) and (Feltz and Zarpentine forthcoming). I take it that there are good reasons to meet the results of these studies with scepticism; however, this article is not the place to for a detailed discussion. Moreover, note that the studies at issue claim to have shown that ordinary speakers have the intuition that Hannah speaks falsely in *High Stakes*. Thus, if we took those data seriously, no pragmatic mechanisms would be needed in the first place: Pragmatic Invariantism would be unmotivated. For further interesting discussion and experimental research on the *Bank Cases* see (Schaffer and Knobe 2012). For critical discussion of the mentioned studies see (DeRose 2011).

³⁴ (Rysiew 2001); I have adjusted the personal pronouns to fit Stanley's case. (Brown 2006: 421 ff.) agrees with Rysiew's explanation but pairs it with a slightly different account of how context determines what counts as relevant.

see what he has in mind, we first need to introduce some terminology. On the background of his contextualist approach, DeRose uses ‘knows_M’ and ‘knows_H’ respectively to denote the epistemic states Hannah is in when satisfying ‘knows’ in *Low Stakes* and in *High Stakes* respectively. Thus, Hannah in both cases knows_M but does not know_H that the bank will be open tomorrow. DeRose then proceeds and paraphrases Rysiew’s view as the view that the sentence Hannah utters in *High Stakes* semantically expresses, in *High Stakes*, the false proposition that she does not know_M *O*, while the utterance pragmatically conveys the true proposition that Hannah does not know_H *O*. Thus, according to DeRose’s reconstruction of Rysiew’s view, Hannah conveys a truth by uttering a sentence that semantically expresses a falsehood.

DeRose then objects to this strategy by pointing out that what is, according to Rysiew’s view, semantically expressed by Hannah’s utterance is far from irrelevant in *High Stakes*. DeRose:

We will grant that the *issue* or the question of whether or not the subject knows_H is the most relevant issue to the purposes at hand (and is therefore more relevant than is the question of whether he knows_M): Whether or not he should go into the bank to brave the lines [...] seems to stand or fall on the issue of whether he knows_H. But that the speaker doesn’t even know_M that the bank will be open on Saturday—which according to Rysiew’s account is what the speakers asserted—would of course settle (in the negative) the salient question or issue of whether [Hannah] knows_H. So, in saying that [she] doesn’t ‘know’ that the bank is open on Saturday, on Rysiew’s account, the speaker asserts what would be an extremely relevant thought. (DeRose 2009: 122).

Clearly, DeRose here assumes that the alleged semantic content of Hannah’s utterance entails the allegedly pragmatically imparted content. In other words, he assumes that Hannah’s not knowing_M *O* entails that she does not know_H *O*. And this is correct, for how could one know_H *O* if one does not also know_M *O*, given that one’s knowing_H requires one to eliminate all those possibilities that knowing_M requires one to eliminate, plus a few extra $\neg O$ -possibilities?

However, note that, in response to DeRose’s objection, Rysiew can point out that he is not committed to DeRose’s paraphrase of his view. In fact, Rysiew does not formulate his view in terms of knowing_M and knowing_H. Rather, he claims that the semantic content of ‘I don’t know *O*’ is the proposition expressed by (18), while the pragmatically imparted content is the proposition expressed by (19):

- (18) Hannah’s evidence does not eliminate every $\neg p$ -world that is *epistemically relevant*.
- (19) Hannah’s evidence does not eliminate every $\neg p$ -world that is *conversationally salient* in *C*.

Surely, Rysiew can argue that the proposition expressed by (18) can be relevant in a given context without the proposition expressed by (19) being relevant in that context: after all, (18) does *not* entail (19), and Rysiew can consistently claim that, in *High Stakes*, Hannah’s evidence eliminates all conversationally salient $\neg O$ -worlds, but not all epistemically relevant ones. For illustration, consider the possibility that the bank has closed down permanently three months ago. That possibility is epistemically relevant in *High Stakes* but it is not conversationally salient. As a consequence, we can make perfect sense of Rysiew’s view that (19) is conversationally relevant in *High Stakes* even though (18) is not (e.g., because it is taken for granted in *High Stakes*)—a view that would be rather difficult to make sense of if, as DeRose assumes, (18) were to entail (19).³⁵ Thus, MPI does not fail for the reasons given by DeRose. However, as we shall see shortly, there are other important problems with the view.

5. Quality and Common Ground

To see what is problematic about MPI, we need to take a closer look at how exactly different types of conversational implicature are triggered. To begin with, note that *Quality*₁ implicatures are usually triggered in contexts in which the utterance’s semantic content conflicts with what Stalnaker (2002: 716) calls the *common ground*. In addition, it is worth noting that Stalnaker’s account of assertion provides an attractive explanation for why certain violations of *Quality*₁ and, in particular, the aforementioned cases of hyperbole, metaphor, and irony from Section 3 trigger implicatures. An assertion can, according to Stalnaker, be thought of as a proposal to add the assertion’s content to the common ground.³⁶ On this view an assertion of p is successful just in case, after the assertion, all members of one’s conversation accept p , believe that all accept p , believe that all believe that all accept p , etc.

³⁵ More generally, note that it is not very promising to object to Rysiew’s view by claiming that the proposition expressed by (16)—the alleged semantic content of Hannah’s utterance in *High Stakes*—is in fact conversationally relevant: Rysiew can respond to such arguments by pointing out that Hannah and Sarah are interested in whether the bank has changed its hours recently and thus in whether Hannah’s evidence eliminates the relevant $\neg O$ -worlds in which it has. This issue is unaffected by the fact that Hannah’s evidence does not eliminate all epistemically relevant $\neg O$ -worlds, given that those $\neg O$ -worlds in which the bank has changed its hours are not epistemically relevant. To make the epistemic fact relevant would require an additional assumption, to the effect that knowledge (or the elimination of epistemically relevant worlds) is the norm of practical reasoning. But this is clearly a principle that Rysiew would reject, given that he divorces conversational and practical facts from epistemic facts.

³⁶ Stalnaker formulates his view in terms of reducing the context set. Since the context set is the complement of the common ground the two formulations are equivalent.

What happens, however, if an assertion's semantic content is incompatible with the common ground? Clearly, in such cases the semantic content of the assertion cannot be added to the common ground without leading to an obvious contradiction. Now, given Stalnaker's account of assertion, it is precisely the impossibility of adding the semantic content of the assertion to the common ground that causes the audience to look for a conversational implicature.³⁷ Thus, on the view at hand, the audience's calculation of an implicature begins with the recognition that the assertion's semantic content is not what the speaker intended to communicate, for that content cannot, after all, be readily added to the common ground.

Next, note that utterances that trigger an implicature *because* their semantic content cannot be readily added to the common ground are, as demonstrated by the examples in Section 3, rather easily detectable as instances of non-literal speech. We can thus infer the following general principle, which I shall call the *Detectability Principle* (DP):

- (DP) If an utterance u in a context C that semantically expresses p in C carries an implicature i in C , and if $\neg p$ is part of the common ground in C , then u is easily detectable as a case of non-literal speech.³⁸

Given the classic picture of Gricean pragmatics and Stalnaker's account of assertion, (DP) is rather unsurprising: competent speakers can, on Stalnaker's account, become easily aware of what is and what is not part of the common ground at a context. Thus, given some rather minimal logical capacities, competent speakers are, accordingly, in a position to become aware of whether an utterance's literal semantic content is incompatible with their context's common ground or not. However, it ought to be emphasised that, independently of this theoretical motivation, (DP) receives its strongest and most crucial support from data such as those from Section 3—that is, from examples such as (3)-(5), (7), (9), and (11).³⁹

³⁷ If no conversational implicature is forthcoming, speakers will, of course, consider revising the common ground.

³⁸ As the reader will be aware by now, I use the phrase 'utterance u semantically expresses p in C ' as shorthand for ' u is an utterance of a sentence S that semantically expresses p in C '. Moreover, I should like to emphasise that I do not claim here that all conversational implicatures whatsoever are easily detectable.

³⁹ Note that detectability in the sense at issue here does not entail that hearers always discover an implicature right after hearing the utterance carrying the implicature. Rather, an implicature is, on the present use of the notion, 'detectable' just in case we can construe a felicitous dialogue of the type discussed on pp. 9-11. Thus, a *Quality*₁-implicature is detectable just in case it allows for the construction of a felicitous dialogue in which (a) the hearer does not discover the implicature and interprets the initial utterance literally and (b) the speaker then clarifies her initial intentions by admitting that she hasn't spoken literally but exaggerated, spoken ironically, etc.

Once (DP) is in place, the *Moderate Pragmatic Invariantist* faces a dilemma. To see what I have in mind, note that the proposition that is, according to MPI, semantically expressed by ‘Hannah knows *O*’ either is or is not part of the common ground in *High Stakes*. Firstly, assume that it is part of the common ground—that is, in *High Stakes* Hannah and Sarah accept, for the purposes of the conversation, that Hannah knows *O* and they believe that they accept, for the purposes of the conversation, that Hannah knows *O*, etc. This should be the expected scenario, given MPI’s semantics of ‘knows’ and the plausible assumption that Hannah and Sarah are not mistaken about either the semantics of ‘knows’ or the truth-value of the proposition that Hannah knows *O*.⁴⁰ It is important to note at this point, however, that if the proposition that Hannah knows *O* is part of the common ground, then Hannah’s utterance of ‘I don’t know *O*’ should, according to (DP) be easily detectable as an instance of conversational implicature: Hannah utters a sentence whose semantic content is incompatible with the context’s common ground. Thus, if Hannah’s utterance is supposed to trigger an implicature, then that implicature must, by virtue of (DP), be detectable without difficulty.

As can be demonstrated easily, however, the hypothesised implicature and the alleged violation of *Quality*₁ are not detectable by competent speakers. Consider the following dialogue:

(20) *High Stakes MPI*:

S: But banks do sometimes change their hours.

H: You’re right, I don’t know that the bank will be open tomorrow.

S: #That’s absurd; of course you know that.

H: #Oh c’mon—I was just exaggerating/being ironic/speaking metaphorically/speaking loosely!

As the infelicity of S’s utterance in (20) demonstrates, the literal semantic content of Hannah’s utterance in *High Stakes* is not in conflict with the common ground in *High Stakes*: if it were part of the common ground and thus uncontroversial in *High Stakes* that Hannah knows *O*, then S’s utterance in (20) should be felicitous and unproblematic—similarly to the responses given in the cases discussed in Section 3, such as (4), (7), (9), and (11). Moreover, as the infelicity of H’s reaction to S’s utterance in (20) shows, Hannah’s utterance is not an instance of hyperbole, irony, metaphor, or loose talk. As a consequence, Hannah’s utterance in *High Stakes* cannot be successfully assimilated to any of

⁴⁰ Note again that the defenders of PI need to avoid the claim that Hannah and Sarah are mistaken about the semantics of ‘knows’, for if they do not, they owe us an explanation of why Hannah and Sarah—and thus competent speakers more generally—are wrong about *High Stakes*. But that explanation, if successful, would render superfluous a pragmatic explanation of the data. See below for an elaboration on this point.

the received cases of utterances that are, at the semantic level, in violation of *Quality*₁. The moral to draw from (20) is, accordingly, that the proposition that Hannah knows *O* is, contrary to our above assumption, not part of the common ground in *High Stakes*.⁴¹

Let us thus turn to the second horn of the dilemma, according to which it is not part of the common ground in *High Stakes* that Hannah knows *O*. To see what is problematic about this view, note that if, in *High Stakes*, it is not part of the common ground that Hannah knows *O*, then this means that Hannah and Sarah do not accept, for the purposes of the conversation, that Hannah knows *O*, or do not believe that they accept that proposition, etc. But why would Hannah and Sarah in *High Stakes* not accept or believe that they accept that proposition, given that it is supposedly true? Why would Hannah and Sarah make this mistake?

The only plausible explanation that comes to mind is that, just before Hannah's last utterance in *High Stakes*, Hannah and Sarah mistakenly believe that the possibility that the bank has changed its hours recently (henceforth '*H*') is epistemically relevant and thus needs to be eliminated by Hannah's evidence in order for her to know *O*. This seems rather plausible, for if Hannah and Sarah were not to believe that *H* is epistemically relevant, then they would, presumably, believe the alleged truth that Hannah knows *O*, and thus accept that proposition for the purposes of the conversation, believe that they accept it, etc. In other words, if Hannah and Sarah were not mistaken about the epistemic relevance of *H*, then it would be part of the common ground that Hannah knows *O*. Moreover, it is worth noting that before Hannah's final utterance in *High Stakes*, Sarah does not know whether Hannah can eliminate the possibility that the bank has changed its hours recently; and it is precisely because of this lack of knowledge that Sarah does not accept or presuppose, for the purposes of the conversation, that Hannah knows *O*: given MPI, Sarah must be mistaken about the epistemic relevance of *H* and thus about the truth-value of the proposition that Hannah knows *O*.

But now the defenders of MPI are in trouble, for they need an explanation of why Hannah and Sarah in *High Stakes* are wrong about the truth-value of the proposition that Hannah knows *O*. Why do Hannah and Sarah falsely believe that Hannah does not know *O*, given that, according to MPI, she clearly does? Surely, it will not be impossible to

⁴¹ A referee for this journal points out that S's utterance in (20) is infelicitous because S herself has just brought up the error-possibility that the bank has changed its hours recently (before A's first utterance). Given that S brought up that error-possibility, it follows directly that she does not presuppose that H knows *O*, for otherwise S's bringing up the mentioned error-possibility would be incompatible with her own presuppositions.

give such an explanation.⁴² What is important here, however, is that once such an explanation is in place, we can utilize it to explain not only Hannah and Sarah's mistake in *High Stakes*, but also our own intuition that Hannah's utterance of 'I don't know *O*' in *High Stakes* conveys a truth despite semantically expressing a falsehood—without any further appeal to implicatures. Consider, for instance, the possible explanation that Hannah and Sarah mistakenly believe that *H* is epistemically relevant because they overestimate its probability: they give increased probabilistic weight to *H* because *H* is psychologically salient to them and because the costs of error are particularly high.⁴³ If we were to accept this line of thought as a viable explanation of why Hannah and Sarah believe that *H* is epistemically relevant, then why should we not also accept it as a viable explanation of our own intuition that Hannah's utterance in *High Stakes* is true? The reason why we intuit that Hannah's utterance conveys a truth in *High Stakes* is, on this view, simply that we overestimate the probability of *H*, and therefore mistakenly believe it to be epistemically relevant. What is a plausible explanation of Hannah's and Sarah's mistake is an equally plausible explanation of our mistake. Thus, on the second horn of our dilemma, MPI again collapses—like Hazlett's, Pritchard's and Black's approaches—into a Williamson-style MII: MPI's pragmatic explanation of the data from *High Stakes* presupposes an antecedently available and independently viable explanation of the data at issue.

Let me sum up the above considerations. I have argued that MPI must claim that the proposition that Hannah knows *O* either is or is not part of the common ground in *High Stakes*. If it is part of the common ground, we should expect competent speakers to detect the fact that Hannah's utterance of 'I don't know *O*' in *High Stakes* is in conflict with the common ground. The fact that we do not see such behaviour casts significant doubt on the view under consideration. If, on the other hand, the defenders of MPI claim that it is not part of the common ground in *High Stakes* that Hannah knows *O*, then they owe us an explanation for why Hannah and Sarah fail to take for granted the alleged truth that Hannah knows *O*. Once such an explanation is in place, the defenders of MPI have rendered superfluous an explanation of our intuition that Hannah speaks truly by means of conversational implicatures. Summing up, all accounts of MPI are either demonstrably

⁴² Note that it will not do to claim that Hannah and Sarah mistakenly believe that Hannah does not know *O*, because utterances expressing the proposition that she does not know *O* would, in *High Stakes*, trigger a false implicature. Whether we believe a proposition is, if implicatures with differing truth-values are involved, independent of whether we accept an utterance expressing it: if you know that John ate all the cookies, you may reject as false an utterance of 'John ate some of the cookies', but you will nevertheless believe that John ate some of the cookies. You do, after all, believe that he ate all of them.

⁴³ Cp. (Gerken 2011, 2012; Nagel 2008, 2010; Williamson 2005a, 2005b).

false (first horn of the dilemma) or explanatorily incoherent (second horn of the dilemma), independently of which Gricean maxim they claim to be violated by Hannah's utterance in *High Stakes*.⁴⁴

6. Conventional Implicature

Let us leave behind the idea that the *Bank Case* data can be accounted for by means of conversational implicatures and turn our attention towards conventional implicatures. Even though the view that the *Bank Case* data can be accounted for by conventional implicatures has not been suggested explicitly in the literature, Rysiew (2007: 658, fn. 616) at times compares 'knowledge'-ascriptions with paradigm examples of conventional implicatures ('but') and both Rysiew (2001: 495, 2007: 639) and Brown (2006: 428) resort to Grice's claim that "some implicatures cannot be 'comfortably' cancelled"—a claim that Grice himself only makes with respect to conventional implicatures.⁴⁵ Now, even though it is, I take it, entirely clear that neither Rysiew nor Brown want to commit to the view that Hannah's utterances in the bank cases carry conventional implicatures, the view itself deserves our attention.

⁴⁴ As a referee for this journal points out, it is worthwhile discussing another issue brought up by Rysiew: as Rysiew (2005: 62) points out, "it is among the sophisticated [pragmatic] invariantist's central claims that our pretheoretic intuitions as to what we're 'saying' are insensitive to the semantic/pragmatic distinction." In response, note firstly that Rysiew's claim does not affect the argument presented in the main text. According to the above argument, either the proposition that Hannah knows *O* is part of the common ground in *High Stakes* or it is not. If it is, the argument goes, MPI violates the independently plausible (DP). If it is not, MPI is explanatorily self-defeating. Thus, Rysiew's comment to the effect that our intuitions are generally insensitive to the semantics/pragmatics distinction leaves the argument in the main text untouched.

Furthermore, note that Rysiew's view itself is too strong and subject to an enormous battery of counterexamples: any uncontroversial and recognized case of conversational implicature is one in which our intuitions distinguish clearly and precisely between speaker meaning and sentence meaning or what is said and what is merely meant or implicated (e.g., Grice's gas station). Thus, pace Rysiew, the intuitions of competent speakers are, in standard cases of conversational implicature, rather obviously sensitive to Grice's distinction and tend to pinpoint in a reliable way the relevant differences in content.

Of course, the defender of MPI might at this point retreat to the considerably weaker claim that our intuitions do *not always* distinguish reliably between what is said and what is conversationally implicated, and then claim that Hannah's utterance in *High Stakes* is precisely such a case in which our intuitions go awry. However, if the defender of MPI were to argue along these lines she would have to provide us with evidence for her view by producing clear and uncontroversial examples of conversational implicatures with respect to which our intuitions are insensitive and then provide further evidence that Hannah's utterance in *High Stakes* functions analogously to those uncontroversial cases. In the absence of such analogous examples or, as we might put it, 'partners in crime', the claim that our intuitions "are insensitive to the semantic/pragmatic distinction" in a way that can be utilized by the pragmatic invariantist remains *ad hoc* and unsubstantiated.

⁴⁵ See (Rysiew 2007: 658, fn. 616) and (Grice 1989: 46) for his remarks on conventional implicatures and cancellability.

To begin with, consider the following two classic examples of conventional implicature:

- (21) Marie is poor, but she's honest.
 - a. Marie is poor and Marie is honest.
 - b. Poor people are not usually honest.
- (22) Even Bart passed the test.
 - a. Bart passed the test.
 - b. Bart was among the least likely to pass the test.⁴⁶

I shall say, in what follows, that (21a) and (22a) express the *truth-conditional contents* of utterances of (21) and (22) respectively, and that (21b) and (22b) express such utterances' *conventionally implicated contents*. Now, what is crucial for our purposes is the familiar fact that sincere utterances of (21) and (22) appear true to competent speakers just in case their truth-conditional contents appear true.⁴⁷ The perceived truth-values of the conventionally implicated b-propositions are entirely irrelevant with respect to the truth-evaluation of utterances of (21) and (22). The truth-conditional contents of the sentence at issue have, in other words, *intuitive primacy* over the conventionally implicated contents.⁴⁸

Once we appreciate this phenomenon of primacy it is fairly obvious that the notion of a conventional implicature is of no use for the defender of PI. Remember that the defender of PI claims that the apparent truth of some 'knowledge'-attributions is accounted for by the truth of their implicatures. In particular, as we have seen in Section 3, the defender of SPI must claim with respect to *Low Stakes* that the truth of Hannah's utterance is explained by the truth of an implicature, which is argued to override the falsity of what is semantically expressed. The same holds for MPI and *High Stakes*: as we have seen in Section 4, the defender of MPI claims that Hannah's utterance in *High Stakes* semantically expresses a falsehood while implicating a truth, and it is again the truth of the implicature that supposedly accounts for our intuition that the utterance at issue conveys a truth. Thus, both SPI and MPI claim that our intuitions track the truth-value of the implicated proposition rather than the truth-value of the literal or truth-conditional content.

⁴⁶ Example (21) is from (Grice 1961: 234) and (22) is borrowed from (Potts 2007).

⁴⁷ Cp. Bach's (1999: 331) definition of the notion of a conventional implicature.

⁴⁸ Moreover, note the truth-conditional content also has primacy over the conventionally implicated content in cases in which the a-content is false while the b-content is true. To see this, consider (22) and assume that its a-content is false while its b-content is true. Clearly, in such a situation we have the intuition that utterances of (22) are false, even though their conventionally implicated content is true.

We can now see that conventional implicatures are of no use to the defender of PI: with respect to conventional implicatures, speaker intuitions track the truth-value of the truth-conditional a-content rather than the conventionally implicated b-content. As Stanley and Szabó (2000: 239-240) put it in a different context, “conventional implicatures [...] tend to add extra information to the proposition expressed, rather than override what is said.” Thus, neither SPI nor MPI can explain the *Bank Case* data by claiming that the relevant utterances carry conventional implicatures. The notion of a conventional implicature is of no use for the defenders of PI.

7. Bachian Implicature

In his latest paper on the topic, Rysiew (2007: 629) suggests yet another view—namely, the view that the phenomenon at issue in *High Stakes* is not that of a conversational implicature, but rather an instance of what Bach (1994) calls “implicature”.⁴⁹ Now, some readers might be somewhat sceptical as to whether the notion of a Bachian implicature picks out a unique phenomenon that cannot be explained more conservatively by a range of other, more received linguistic mechanisms.⁵⁰ The discussion of these issues, however, goes well beyond the scope of this paper. Nevertheless, it ought to be emphasised here that an account of the data from *High Stakes* that relies on the controversial notion of an implicature will *ipso facto* be controversial, too. Thus, the appeal to implicatures undermines to some degree the initial attraction of PI, for the defender of PI can now no longer maintain that her explanation of the relevant data appeals exclusively to the received notions of orthodox Gricean pragmatics.

Besides this downside, there is further trouble for Rysiew’s proposal. To see what I have in mind let us take a brief look at Bach’s notion of an implicature. According to Bach, “[i]n implic-i-ture, one says something but does not mean that; rather, what one means includes an implicit qualification on what one says, something that one could have made explicit but did not.”⁵¹ “Implicatures involve”, as Bach puts it, “an unexpressed qualification on what is said.”⁵² Moreover, there are, according to Bach, two different types of implicature—namely, those involving “expansion” and those involving “comple-

⁴⁹ Rysiew prefers this approach because, as he puts it, “in implicatures, properly so-called, one means what one says but also something else” (Rysiew 2001: , p. 510, fn. 32); also (Rysiew 2007: 643), which is (apparently) not so with Bachian implicatures. However, Rysiew’s claim about conversational implicatures is surely incorrect, as the Gricean treatment of phenomena such as metaphor, irony, hyperbole, and loose use—that is, of violations of *Quality*₁, suggests.

⁵⁰ See, for instance, (Stanley and Szabó 2000) and (Stanley 2002).

⁵¹ (Bach 2001: 251).

⁵² (Bach 2001: 253).

tion”.⁵³ In the case of completion, the sentence uttered does not express a full proposition, but merely a “propositional radical” that needs to be “filled in”,⁵⁴ while expansion is an operation on an already complete proposition that is expressed by the sentence uttered. Since the defenders of both MPI and SPI claim that ‘knowledge’-ascriptions semantically express complete propositions (cp. Sections 3 and 4), they cannot adopt the view that the implicatures at issue are completion implicatures. If they are to be implicatures at all, then they must be so-called expansion implicatures.

Here are a few examples of expansion implicatures, where the bracketed material expresses the allegedly “unarticulated constituents”⁵⁵ of the propositions typically conveyed by utterances of their preceding sentences:⁵⁶

- (23) a. You’re not going to die. (from that cut)
b. I have eaten breakfast. (today)
c. I have eaten caviar. (before)
d. France is hexagonal. (roughly)
e. Dennis had sex and got herpes. (as a result)
f. They fell in love and got married. (in that order)

As Bach observes concerning examples such as the above ones, implicature “expansion involves [...] what might be called ‘lexical’ strengthening, in that what is being communicated could have been made fully explicit by the insertion of additional lexical material.”⁵⁷

Given the characterisation of expansion in terms of lexical strengthening it is far from obvious that Hannah’s utterances in the *Bank Cases* are instances of expansion implicatures. To see this, note that it is rather unclear what the allegedly suppressed lexical material—the unarticulated constituents—of Hannah’s utterance in *High Stakes* or *Low Stakes* could be: according to Bach, expansion implicatures can be made fully explicit by including the “appropriate lexical material [...] in the utterance.”⁵⁸ Thus, if we took Hannah’s utterance of ‘I don’t know *O*’ in *High Stakes* or of ‘I know *O*’ in *Low Stakes* to be a case of expansion, then it should—just as in the examples in (23)—be fairly obvious what the additional but suppressed lexical material of that utterance would consist in. That this

⁵³ (Bach 1994: 126).

⁵⁴ (Bach 1994: 154).

⁵⁵ The phrase ‘unarticulated constituent’ goes back to (Perry 1986), Bach (1994: 127, fn. 124) uses it only in passing.

⁵⁶ For further examples see (Bach 1994: 128).

⁵⁷ (Bach 1994: 134).

⁵⁸ (Bach 1994: 140).

is far from obvious, however, casts significant doubt on the view that the utterances at issue are instances of Bach's expansion implicature.⁵⁹

Summing up, Bachian implicatures describe a phenomenon rather different from the one exhibited by the *Bank Cases* and are accordingly of little use in explaining the relevant data.

8. Conclusion

Pragmatic invariantist explanations of the data from the *Bank Cases* are sometimes claimed to have a rather attractive advantage over other accounts of those data—namely, their explanatory conservatism and simplicity. Once we accept PI, the argument goes, we can account for the data at issue by means of pragmatic mechanisms alone; mechanisms the postulation of which is independently motivated by a vast array of unrelated linguistic data. There are, I take it, two problems with this line of reasoning. Firstly, it is far from obvious that the accounts competing with PI over an explanation of the *Bank Case* data are explanatorily less conservative and simple. Secondly, and more importantly, there is, as we have seen in the course of this paper, very convincing evidence that a satisfactory explanation of the aforementioned data cannot be given in pragmatic terms. As I have argued above, all versions of PI that aim to account for the data in terms of implicatures or similar pragmatic phenomena either violate the *Detectability Principle* or are explanatorily incoherent. The arguments that I have used to establish this conclusion rely exclusively on widely accepted assumptions about the nature of conversational implicatures, the role of the common ground in linguistic communication, and empirically well supported principles.⁶⁰ The alleged theoretical virtues of PI are, as a consequence, only apparent, and in explaining the data from the *Bank Cases* we have to resort to accounts that either take those data at their semantic face value—that is, theories such as *Epistemic Contextualism*, *Subject-Sensitive Invariantism* or *Epistemic Relativism*—or that reject them entirely on psychological or empirical grounds.⁶¹

⁵⁹ Leite (2005: 226) makes a similar point.

⁶⁰ Note also that since my arguments rest only on the *Detectability Principle*, the claim that linguistic meaning is determined by ordinary use, and other platitudes concerning Stalnaker's notion of *common ground*, resorting to alternative pragmatic frameworks such as Sperber and Wilson's (1986) relevance framework or Horn's (2004) neo-Gricean approach will not help the pragmatic invariantist: those frameworks are equally committed to the theoretically neutral *Detectability Principle*, the connection between meaning and use, and the relevant platitudes concerning Stalnaker's notion of the *common ground*.

⁶¹ See, for instance, Williamson's *Moderate Insensitive Invariantism* that accounts for the data by means of psychological considerations. For an interesting discussion and further development of Williamson's view see (Nagel 2008, 2010).

References

- Bach, Kent (1987). Thought and Reference. Oxford, Clarendon.
- Bach, Kent (1994). "Conversational Implicature." Mind & Language **9**(2): 124-162.
- Bach, Kent (1999). "The Myth of Conventional Implicature." Linguistics and Philosophy **22**(4): 367-421.
- Bach, Kent (2001). "Speaking Loosely: Sentence Nonliterality." Midwest Studies in Philosophy **25**(1): 249-263.
- Bach, Kent (2005). The Emperor's New 'Knows'. Contextualism in Philosophy: On Epistemology, Language and Truth. G. Preyer and G. Peter. Oxford, OUP: 51-90.
- Black, Tim (2005). "Classic Invariantism, Relevance, and Warranted Assertability Manœuvres." Philosophical Quarterly **55**: 328-336.
- Brown, Jessica (2006). "Contextualism and Warranted Assertibility Manœuvres." Philosophical Studies **130**: 407-435.
- Cohen, Stewart (1999). "Contextualism, Skepticism, and the Structure of Reasons." Philosophical Perspectives **13**: 57-89.
- Cohen, Stewart (2000). "Replies." Philosophical Issues **10**(1): 132-139.
- Davis, Wayne A. (2004). "Are Knowledge Claims Indexical?" Erkenntnis **61**: 257-281.
- Davis, Wayne A. (2007). "Knowledge Claims and Context: Loose Use." Philosophical Studies **132**: 395-438.
- DeRose, Keith (1992). "Contextualism and Knowledge Attributions." Philosophy and Phenomenological Research **52**: 913-929.
- DeRose, Keith (1999). Contextualism: An Explanation and Defence. The Blackwell Guide to Epistemology. J. Greco and E. Sosa. Oxford, Blackwell: 187-205.
- DeRose, Keith (2002). "Assertion, Knowledge, and Context." The Philosophical Review **111**(2): 167-203.
- DeRose, Keith (2009). The Case for Contextualism: Knowledge, Skepticism, and Context. Vol. 1. New York ; Oxford, Oxford University Press.
- DeRose, Keith (2011). "Contextualism, contrastivism, and X-Phi surveys." Philosophical Studies **156**(1): 81-110.
- Fantl, Jeremy and Matthew McGrath (2002). "Evidence, Pragmatics and Justification." The Philosophical Review **111**(1): 67-94.
- Fantl, Jeremy and Matthew McGrath (2009). Knowledge in an Uncertain World. Oxford ; New York, Oxford University Press.
- Feltz, Adam and Chris Zarpentine (forthcoming). "Do You Know More When it Matters Less?" Philosophical Psychology.
- Gerken, Mikkel (2011). "Epistemic Focal Bias." Australasian Journal of Philosophy: 1-21.
- Gerken, Mikkel (2012). On the Cognitive Bases of Knowledge Ascriptions. Knowledge Ascriptions. J. Brown and M. Gerken. Oxford, Oxford University Press: 140-170.
- Grice, Herbert Paul (1961). "The Causal Theory of Perception." Proceedings of the Aristotelian Society, Supplementary Volume **35**: 121-153.

- Grice, Herbert Paul (1989). Studies in the Way of Words. Cambridge, Mass., Harvard UP.
- Halliday, Daniel (2005). "What Explains Our Intuitions about Knowledge Ascriptions?" Proceedings of the Aristotelian Society **105**(3): 393-402.
- Hawthorne, John (2004). Knowledge and Lotteries. Oxford, OUP.
- Hazlett, Allan (2009). "Knowledge and Conversation." Philosophy and Phenomenological Research **78**(3): 591-620.
- Horn, L. R. (2004). Implicature. The Handbook of Pragmatics. L. R. Horn and G. Ward. Oxford, Blackwell: 3-28.
- Ichikawa, Jonathan (2011a). "Quantifiers and epistemic contextualism." Philosophical Studies **155**(3): 383-398.
- Ichikawa, Jonathan (2011b). "Quantifiers, Knowledge, and Counterfactuals." Philosophy and Phenomenological Research **82**(2): 287-313.
- Leite, Adam (2005). "Some Worries for Would-Be WAMmers." Grazer Philosophische Studien **69**(1): 101-126.
- Lewis, David (1996). "Elusive Knowledge." Australasian Journal of Philosophy **74**: 549-567.
- MacFarlane, John (2005). The Assessment Sensitivity of Knowledge Attributions. Oxford Studies in Epistemology. T. Gendler and J. Hawthorne. Oxford, OUP: 197-233.
- May, Joshua, Walter Sinnott-Armstrong, et al. (forthcoming). "Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study." Review of Philosophy and Psychology.
- Nagel, Jennifer (2008). "Knowledge Ascriptions and the Psychological Consequences of Changing Stakes." Australasian Journal of Philosophy **86**(2): 279-294.
- Nagel, Jennifer (2010). "Knowledge Ascriptions and the Psychological Consequences of Thinking about Error." The Philosophical Quarterly **60**(239): 286-306.
- Perry, John (1986). "Thought without Representation." Proceedings of the Aristotelian Society, Supplementary Volumes **60**(ArticleType: research-article / Full publication date: 1986 / Copyright © 1986 The Aristotelian Society): 137-150.
- Potts, Christopher (2007). "Into the Conventional-Implicature Dimension." Philosophy Compass **4**(2): 665-679.
- Pritchard, Duncan (2010). Contextualism, Skepticism and Warranted Assertibility Manoeuvres. Knowledge and Skepticism. J. Keim-Campbell, M. O'Rourke and H. Silverstein. Cambridge, MA, MIT Press: 85-104.
- Rysiew, Patrick (2001). "The Context-Sensitivity of Knowledge Attributions." Nôus **35**(4): 477-514.
- Rysiew, Patrick (2005). "Contesting Contextualism." Grazer Philosophische Studien **69**: 51-70.
- Rysiew, Patrick (2007). "Speaking of Knowing." Nôus **41**(4): 627-662.
- Schaffer, Jonathan (2004a). "From Contextualism to Contrastivism." Philosophical Studies **119**: 73-103.
- Schaffer, Jonathan (2004b). "Skepticism, Contextualism, and Discrimination." Philosophy and Phenomenological Research **69**: 138-155.

- Schaffer, Jonathan (2005). *Contrastive Knowledge*. Oxford Studies in Epistemology. J. Hawthorne and T. Szabó Gendler. Oxford, OUP. **1**: 235-271.
- Schaffer, Jonathan (2006). "The Irrelevance of the Subject: Against Subject-Sensitive Invariantism." Philosophical Studies **127**: 87-107.
- Schaffer, Jonathan and Joshua Knobe (2012). "Contrastive Knowledge Surveyed." Noûs **46**(4): 675-708.
- Sperber, Dan and Deirdre Wilson (1986). Relevance: Communication and Cognition. Oxford, Blackwell.
- Stalnaker, Robert (1998). "On the Representation of Context." Journal of Logic, Language, and Information **7**: 3-19.
- Stalnaker, Robert (2002). "Common Ground." Linguistics and Philosophy **25**: 701-721.
- Stalnaker, Robert C. (1978). "Assertion." Syntax and Semantics **9**: 315-332.
- Stanley, Jason (2002). "Making it Articulated." Mind & Language **17**(1&2): 149-168.
- Stanley, Jason (2005). Knowledge and Practical Interests. Oxford, OUP.
- Stanley, Jason and Zoltán Gendler Szabó (2000). "On Quantifier Domain Restriction." Mind & Language **15**(2&3): 219-261.
- Turri, John (2010). "Epistemic Invariantism and Speech Act Contextualism." Philosophical Review **119**(1): 77-95.
- Turri, John (2011). "Knowledge Guaranteed." Noûs: no-no.
- Unger, Peter (1975). Ignorance: A Case for Scepticism. Oxford, Clarendon.
- Unger, Peter (1984). Philosophical Relativity. Oxford, Blackwell.
- Williamson, Timothy (2005a). "Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge." The Philosophical Quarterly **55**: 213-235.
- Williamson, Timothy (2005b). Knowledge, Context and the Agent's Point of View. Contextualism in Philosophy: On Epistemology, Language and Truth. G. Preyer and G. Peter. Oxford, OUP: 91-114.