

Knowledge, Pragmatics, and Error

Dirk Kindermann

dirk.kindermann@gmail.com

Penultimate Draft. *Grazer Philosophische Studien* 2016, special issue on *Context-Relativity in Semantics* (ed. C. Gauker)

ABSTRACT. ‘Know-that’, like so many natural language expressions, exhibits patterns of use that provide evidence for its context-sensitivity. A popular family of views – call it *pragmatic invariantism* – attempts to explain the shifty patterns by appeal to a pragmatic thesis: while the semantic meaning of ‘know-that’ is stable across all contexts of use, sentences of the form ‘S knows [doesn’t know] that p’ can be used to communicate a pragmatic content that depends on the context of use. In this paper, I argue that pragmatic invariantism makes inaccurate predictions for a wide range of well-known use data and is committed to attributing systematic pragmatic error to ordinary speakers. But pragmatic error is unprecedented, and it is doubtful that speakers are systematically wrong about what they intend to communicate.

KEYWORDS. Knowledge attributions, invariantism, error, implicature, pragmatics, context-sensitivity

1 Introduction

How should we account for the shifty patterns of use that ‘know-that’ exhibits? It is widely agreed, though details are heavily contested,¹ that there is contextual variability in speakers’ acceptability judgments and other use patterns. This variability isn’t predicted and explained by traditional invariantist assumptions about the meaning of ‘know-that.’ Four major strategies have been pursued to account for this variability:

1. The semantic strategy: Epistemic contextualists² and relativists³ argue that ‘know-that’ is semantically context-sensitive: the truth-value of knowledge sentences – declarative sentences of roughly the form ‘S knows that p’ and ‘S doesn’t know that p’ – depends on epistemologically relevant factors, such as interests, purposes, stakes, and the salience of error possibilities, at the context of use (contextualism) or at the context of assessment (relativism).⁴

¹Current empirical studies of speakers’ use of knowledge claims deliver a mixed picture. See, e.g., the introduction and essays in [Beebe \(2014\)](#) for a useful overview of the current literature on the alleged influence of practical stakes and salience of error on speakers’ judgments.

²See for instance [Blome-Tillmann \(2014\)](#); [Cohen \(1986, 1987, 1999\)](#); [DeRose \(1992, 1995, 2009\)](#); [Ichikawa \(2010\)](#); [Kompa \(2002\)](#); [Lewis \(1996\)](#); [MacFarlane \(2009\)](#); [Neta \(2003\)](#); [Schaffer \(2004a\)](#); [Schaffer & Szabó \(2014\)](#).

³See for instance [MacFarlane \(2005, 2011, 2014\)](#); [Richard \(2004, 2008\)](#).

⁴Expressivists also account for the alleged context-sensitivity of ‘know-that’ in partly semantic terms; cf. [Chrisman \(2007\)](#).

2. The metaphysical strategy: Subject-sensitive or interest-relative invariantists argue that the knowledge relation supervenes on non-traditional factors such as the interests, purposes, and stakes of the knowing subject, and that this accounts for speakers' shifty patterns of use.⁵
3. The error strategy: some invariantists account for the uses their view does not predict by claiming that speakers are simply wrong when they use 'know-that' in this way, and that this error can be explained by appeal to some general and expectable psychological mechanisms.⁶
4. The pragmatic strategy: A popular strategy among traditional invariantists is to argue that while 'know-that' is semantically invariant – it denotes the same knowledge relation across all contexts of use and assessment – uses of knowledge sentences may pragmatically convey more than their literal semantic content; and this pragmatically conveyed content may vary with context, thus accounting for shifty patterns of use.

My concern in this paper is with the pragmatic strategy (though I'll occasionally touch on the other strategies where comparison is helpful.) Let's call *pragmatic invariantism* any view that endorses the following two theses:

- (I) *Invariantism*: 'Know-that' semantically expresses an invariant meaning; its contribution to the content expressed by sentences in which it occurs does not vary with features of the context (of the subject, user/attributor, or assessor) such as interests, purposes, stakes or salient error possibilities. Nor does the truth value of a sentence 'S knows [doesn't know] that p' at a context of use depend on such contextual features.⁷
- (P) *Pragmatic content*: Many uses of knowledge sentences pragmatically convey a content that may be distinct from the semantic content expressed at the context of use.

Pragmatic invariantism is a popular view among invariantist epistemologists who wish to address our variable intuitive judgments regarding knowledge. Many differences aside, the advocated views can be divided in two groups:

- (M) *Moderatism*: Ordinary agents know a fair amount, and many of our ordinary knowledge attributions are literally true, even when the knowing subject's epistemic position is not strong enough for them to rule out sceptical counter-possibilities.

Moderate Pragmatic Invariantism (MPI) is the combination of (M), (P), and (I); its proponents include [Black \(2005\)](#), [Brown \(2005, 2006\)](#), [Hazlett \(2009\)](#), [Locke \(forthcoming\)](#), [Lutz \(2014\)](#), [Pritchard \(2010\)](#), and [Rysiew \(2001, 2005, 2007\)](#).

- (S) *Strict/Sceptical*: Knowledge is hard to attain, and most of our ordinary knowledge attributions are literally false. Knowledge requires the ability to rule out even (most of the) far-fetched counter-possibilities.

⁵See for instance [Fantl & McGrath \(2002, 2009\)](#); [Hawthorne \(2004\)](#); [Stanley \(2005\)](#); [Weatherson \(2005\)](#).

⁶See for instance [Gerken \(2012, 2013\)](#); [Nagel \(2010b,a\)](#); [Williamson \(2005\)](#) on behalf of moderate or anti-sceptical traditional invariantism, and [Hawthorne \(2004\)](#) for an implementation of the strategy for subject-sensitive invariantism.

⁷What I call 'invariantism' is often called classical, traditional, strict, or insensitive invariantism, to contrast it with forms of subject-sensitive or interest-relative invariantism. I do not mean to deny the label 'invariantism' to the latter views, but merely to keep my label simple here.

Strict Pragmatic Invariantism (SPI) comprises theses (S), (P), and (I); its proponents include Davis (2004, 2007, 2015) and Schaffer (2004b).⁸

To keep the discussion focused, the primary target in this paper will be MPI – specifically, the versions defended by Brown, Hazlett, and Rysiew. However, the main objections apply, *mutatis mutandis*, to other versions of MPI and to SPI.

Many proponents of pragmatic invariantism have focused on explaining use data from bank cases and airport cases. However, a look at a broader variety of data shows that MPI makes the wrong predictions in many cases; further pragmatic theses aside, the views are committed to a particularly implausible form of speaker error. Or so I will argue in sections 2 – 6, taking into account patterns of cross-contextual rejection of knowledge claims, of retraction (section 4), and sceptical paradox (section 5). I will then show in section 7 that parallel objections apply to SPI, which suggests that the trouble is with the pragmatic strategy, not just particular implementations.⁹

A note on terminology before we start. I will use ‘knowledge sentences’ to refer to *sentences* of roughly the forms ‘S knows that p’ and ‘S doesn’t know that p’. I will reserve the term ‘knowledge attribution’ for the affirmative ascription of propositional knowledge, expressed by instances of ‘S knows that p’, and I will use ‘knowledge denial’ to talk about denials of knowledge-possession expressed in uses of ‘S doesn’t know that p’. Knowledge attributions and denials are thus speech acts, and I will use ‘knowledge claim’ to cover both.

2 The Case for Moderate Pragmatic Invariantism

Proponents of MPI have focused on basic variability data that initially motivated contextualist semantics of ‘know-that.’ Once more, consider DeRose’s bank case:

Bank Case A. My wife and I are driving home on a Friday afternoon. We plan to stop at the bank on the way home to deposit our paychecks. But as we drive past the bank, we notice that the lines inside are very long, as they often are on Friday afternoons. Although we generally like to deposit our paychecks as soon as possible, it is not especially important in this case that they be deposited right away, so I suggest that we drive straight home and deposit our paychecks on Saturday morning. My wife says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays.” I reply, “No, I know it’ll be open. I was just there two weeks ago on Saturday. It’s open until noon.”

Bank Case B. My wife and I drive past the bank on a Friday afternoon, as in Case A, and notice the long lines. I again suggest that we deposit our paychecks on Saturday morning, explaining that I was at the bank on Saturday morning only two weeks ago and discovered that it was open until noon. But in this case, we have just written a very large and very important check. If our paychecks are not deposited into our checking account before

⁸ Schaffer has not gone on to defend the view. As Dinges (2015b) notes, SPI can be found (with or without endorsement) in Bach (2010), Conee (2005, 52f.), Douven (2007), Fantl & McGrath (2009, 185–94) and Bonjour (2010, 78). The ‘Pluralistic Skepticism’ advocated by Cappelen (2005) is kindred in spirit to SPI. It accepts (S) and (I) and in lieu of (P) claims that in utterances of knowledge sentences, speakers assert indefinitely many propositions, only some of which are semantically expressed (speech act pluralism).

⁹For further objections to pragmatic invariantism, see e.g. Blome-Tillmann (2013), Dimmock & Huvenes (2014), and Dinges (2015b).

Monday morning, the important check we wrote will bounce, leaving us in a very bad situation. And, of course, the bank is not open on Sunday. My wife reminds me of these facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no, I don’t know. I’d better go in and make sure.” (DeRose, 2009, 1-2)¹⁰

The widely reported response to this contrast-pair of cases is acceptance of DeRose’s knowledge attribution (‘I know it’ll be open’) in case A and acceptance of DeRose’s knowledge denial (‘I don’t know [the bank will be open tomorrow]’) in case B.¹¹ Generalizing from bank cases, we can for our purposes state that theories of the meaning of ‘know-that’ need to account for the following linguistic data:

Basic Variability

1. Knowledge Attributions in Low. Many ordinary knowledge attributions to subjects who are known to have true beliefs and to be in decent epistemic positions regarding the known proposition seem acceptable and true, as in case A.¹²
2. Knowledge Denials in HIGH. Many knowledge denials in HIGH, such as DeRose’s in case B, seem acceptable and true.
3. Knowledge Attributions in HIGH. Knowledge attributions in HIGH seem unacceptable and false.¹³

MPI explains the acceptability of knowledge attributions in Low (case A) in terms of the moderate invariantist semantics for ‘know-that’. Since the bank will be open on Saturday and since DeRose in both A and B believes that the bank will be open and is in a decent epistemic position, ‘I know that it’ll be open’ is true and hence acceptable as used in the context of A.¹⁴

The invariantist semantics also predicts that DeRose’s knowledge denial in HIGH semantically express a falsehood in HIGH. After all, DeRose meets the invariant moderate standard required for knowing according to MPI. But this prediction clashes with intuitive judgment. Similarly, MPI predicts that if DeRose were to utter ‘I know it’ll be open’ in HIGH (B), it would semantically express a truth, contrary to the intuitive judgment that it is unacceptable and false.

¹⁰See also Cohen’s airport case (Cohen, 1999, 58)

¹¹Reservations about the empirical generalizability of such judgments aside (see footnote 1), we may still qualify these reported responses on behalf of MPI: Hazlett (2009) Lutz (2014) and Pritchard (2010) cast doubt on the widespread *acceptability* judgment concerning the knowledge denial in case B and only see the need to account for ordinary speakers’ *rejection* of knowledge attributions (‘I know it’ll be open’) in case B. However, even Hazlett (2009) and Lutz (2014) along with Brown (2006) and Rysiew (2001) (but not Pritchard (2010) offer pragmatic accounts of knowledge denials in cases like B, so I will take my target here to be the version of MPI that accounts for acceptable knowledge denials in case B.

¹²I won’t here be concerned with distinguishing potential non-traditional factors at play in the speech context: practical stakes, interests and purposes, and the salience of error possibilities. Moreover, I’ll keep with the simplifying tradition of calling ‘Low’ all contexts in which such factors determine a low, meetable standard for the epistemic position required to count as knowing. In contrast, HIGH will be the label for contexts which determine a very high epistemic standard necessary for the possession of knowledge.

¹³My arguments in this paper won’t depend on differences between 1st-person knowledge attributions (e.g., ‘I know/don’t know that p’) and 2nd- and 3rd-person knowledge attributions (e.g., ‘S_he knows/doesn’t know that p’). So for simplicity, I will set them aside.

¹⁴Proponents of MPI differ in their explication of ‘epistemic position.’ But the details won’t matter to us here. For ease of exposition, I will keep with Rysiew who couches his view in terms of the subject’s ability to rule out counter-possibilities to the believed proposition.

In order to explain the acceptability of knowledge denials and the inacceptability of knowledge attributions in HIGH, MPIers claim that knowledge claims in HIGH pragmatically convey a content distinct from the semantically expressed content.¹⁵ According to Rysiew (2001, 2005, 2007), who will serve as our foil, DeRose’s ‘I don’t know it’ll be open’ as asserted in HIGH pragmatically conveys, inter alia, that DeRose cannot rule out the counter-possibilities to the bank’s being open that are salient at the context of utterance, HIGH. According to Rysiew, *salient* counter-possibilities need not be ‘relevant’ counter-possibilities – those that the subject needs to be able to rule out in order to count as knowing. Since DeRose in fact cannot rule out that the bank recently changed its hours – a salient but not relevant counter-possibility – the pragmatically conveyed content is true. It is this pragmatically conveyed content that acceptability judgments track. Similarly, an assertion of ‘S knows that p’ in HIGH semantically expresses that p, that S believes that p and that S can rule out the relevant counter-possibilities to the proposition that p. In addition, it conveys the pragmatic content that S can rule out the salient counter-possibilities. This content is false. Inacceptability judgments track (at least) this pragmatically conveyed content.¹⁶

3 MPI’s Implicatures

Proponents of MPI agree that knowledge claims in HIGH involve Gricean conversational implicatures.¹⁷ While my objections won’t depend on details of the Gricean framework, two features of the pragmatic implicatures of knowledge claims are worth highlighting, and they’re best illustrated within the Gricean framework.

First, the distinction between *exploitation* implicatures and *observation* implicatures is important.¹⁸ Observation implicatures do not involve any failure, on the speaker’s part, to fulfil the Cooperative Principle. Consider the following cases of observation implicatures.

¹⁵Not every proponent of MPI attempts to explain the acceptability of knowledge denials in HIGH in pragmatic terms. While Brown, Rysiew and Lutz provide implicature accounts, Pritchard (2010, 89) rejects the datum and Hazlett (2009, 615-79) claims that making and accepting knowledge denials in HIGH is ‘a philosophical mistake’ but also offers a Gricean account, according to which knowledge denials are cases of meiosis, understating the agent’s epistemic position.

¹⁶Versions of MPI differ in what implicatures are generated with raised stakes. While Black (2005, 2008) follows Rysiew’s formulations of the implicature, Brown maintains that an assertion of ‘S knows that p’ can ‘pragmatically convey that S is in a very strong epistemic position, that her belief matches the facts across a wide range of worlds, including some so far away they are not normally taken to undermine knowledge’ (Brown, 2006, 428). According to Hazlett, a knowledge attribution may have ‘one or more of the following conversational implicatures: (I) S’s belief that p is infallible (with respect of salient alternatives), i.e. she can rule out all contextually salient possibilities incompatible with p; (II) S can prove or cogently argue that p (in the current context); (III) S is justifiedly certain that p is true’ (Hazlett, 2009, 607). For Pritchard (2010), first-person knowledge attributions trigger the implicature that the speaker is able to back up the assertion (of the embedded sentence) with adequate evidence. According to Lutz, a knowledge attribution, if made in contexts ‘in which a bit of practical reasoning [...] is salient, pragmatically implicates that S is rational to take p as true in that practical reasoning’ (Lutz, 2014, 1728). Similarly, for Locke the implicature is that ‘it is rationally permissible for S to act as if p in all normal choice situations’ (Locke, forthcoming, section 3).

¹⁷Brown (2006) and Rysiew (2007) in places liken knowledge claims to cases of what Kent Bach calls ‘implicatures’ (Bach, 1994). The differences between Gricean conversational implicatures and Bachian implicatures won’t matter to the features of implicatures I discuss here.

¹⁸Cf. Meibauer (2009, 366), Marmaridou (2000, 230f.), and Dinges (2015a) for observation vs exploitation implicatures.

- (1) The chairperson opened the meeting and took her seat.
+> The chairperson first opened the meeting and then took her seat.
- (2) The chairperson took her seat and opened the meeting.
+> The chairperson first took her seat and then opened the meeting.
- (3) The meeting starts at 5pm.
+> The speaker believes that the meeting starts at 5pm and has adequate evidence for her belief.

(1) and (2) pragmatically implicate, but do not literally express, a different order of events. In both cases, the speaker does observe the Gricean Cooperative Principle (CP)¹⁹—in particular, the submaxim ‘Be orderly’ of the maxim of Manner. In (3), it is the speaker’s observation of a submaxim of Quality—‘Do not say what you believe to be false’ (Grice, 1975, 27)—that allows for the inference that she believes that the meeting starts at 5pm and that she has adequate evidence for her belief.²⁰

In contrast, pragmatic implicatures from simple knowledge claims are exploitation implicatures. They come about by exploitation of the CP, and more specifically of the conversational maxim of *Relation*. It is because of the hearer’s assumption that a cooperative speaker will speak to the counter-possibilities presently *salient* in the conversation that she will infer (consciously or subconsciously) that the speaker must have intended to communicate something different from what her sentence literally, semantically expresses—to wit, that DeRose cannot rule out the relevant alternatives (which do not include the bank’s recently having changed its hours). The speaker *flouts* the maxim of Relation at the level of ‘what she said’, at the level of the semantic content expressed, which triggers the hearer’s calculation of the pragmatic implicature.²¹

Because (1)–(3) are examples of observation implicatures, they are inappropriate but often-used types of analogies to the alleged case of pragmatic implicatures of knowledge claims, which are exploitation implicatures. They illustrate the workings of knowledge-implicatures for MPI just as little as ‘here’ and ‘now’ illustrate the alleged semantic indexicality of ‘know-that’ for the contextualist.

Second, the implicatures generated by knowledge attributions in HIGH are additive. Implicatures can usefully be divided into *additive* and *substitutional* implicatures.²² In cases of additive implicatures, the speaker means what she says (what is semantically expressed in the context) and also something else. When Grice’s speaker utters ‘X is meeting a woman’, he intends to communicate that X is meeting a woman *and* that this woman is not X’s partner,

¹⁹The Cooperative Principle says: ‘Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.’ Grice (1975, 26)

²⁰It’s not universally accepted that cases like (3) present cases of conversational implicatures at all. After all, such alleged Quality-implicatures do not seem to be smoothly cancellable. I use (3) here as an example of an observation implicatures because some proponents of MPI, including Lutz, Pritchard, Rysiew, make use of Quality-implicatures. Thanks to Alexander Dinges for drawing my attention to this issue.

²¹Not all MPIers agree on which maxim is involved. Relation figures in Black’s, Brown’s, Hazlett’s, Locke’s, Lutz’s, and Rysiew’s accounts. However, for Pritchard it’s a submaxim of the second maxim of Quality (‘Do not say that for which you lack adequate evidence’, Grice (1975, 27)) whose exploitation triggers the hearer’s calculation of the implicature; and according to Black and Hazlett, Quantity (‘Make your contribution as informative as is required (for the current purposes of exchange)’ and ‘Do not make your contribution more informative than is required’) is involved in addition to Relation, whereas Lutz also sees Quality involved.

²²See, e.g. Meibauer (2009, 374), Meibauer (2014, 162), and (Dinges, 2015a, 56) for the distinction.

mother, sister, platonic friend, ... (Grice, 1975, 37). In substitutional implicatures, in contrast, the speaker doesn't mean what is literally said but only something else in its place. Typical examples of substitutional implicatures are the linguistic tropes of irony, metaphor, hyperbole and meiosis (understatement). With an assertion of 'He's a fine friend', the speaker intends to communicate that the person in question is *not* a good friend; the speaker does not intend to communicate that the person is a fine friend and also not a good friend.

Knowledge attributions in HIGH are additive: It is natural to assume that part of what they communicate is that the subject believes the proposition in question and that it is true. But neither the belief- nor the truth-condition is conveyed with the implicature that S can rule out the salient counter-possibilities.²³ So if they are part of what is communicated, it must be because the literal content is also communicated. That is, both literal content and implicated content are communicated.²⁴

In contrast, knowledge denials in HIGH must be cases of substitutional implicature, if MPI is to explain intuitive judgments of acceptability and truth. The literal content – that it's not the case that [S believes that p, that p, and that S can rule out the relevant counter-possibilities] – is false. S does have knowledge. So it's only the implicated content that S cannot rule out the salient counter-possibilities that is true and thus can be the target of speakers' truth judgments.²⁵

In sum, implicatures generated by knowledge attributions are additive; implicatures from knowledge denials are substitutional; and both are exploitation implicatures (see Table 1). Implicatures are a motley crew. In drawing analogies to 'standard' cases of (Gricean conversational) implicature, we should carefully observe whether those cases really work in analogous ways and thus serve to illustrate what is going on with knowledge claims.

²³'Rule out' may be heard as factive, in which case it is only the belief-condition that fails to get communicated with the implicature.

²⁴Note, however, that judgments of falsity, or unacceptability, of knowledge attributions in HIGH, given in *Basic Variability*, do not by themselves lend sufficient support to the assumption that belief- and truth-condition are part of the communicated content. If judgments track what the speaker in HIGH means to convey, then falsity judgments are predicted on both the assumption of substitutionality and the assumption of additivity. The pragmatically implicated content by itself is false (substitutional implicature), and the conjunction of the true literal content and the false pragmatically implicated content is also false (additive implicature). That said, I think we do have reason to assume that in HIGH, knowledge attributions communicate the belief- and truth-condition. Consider a bank case B', just like B except that it doesn't contain the phrase 'Remaining as confident as I was before that the bank will be open.' A speaker asked to judge whether as assertion by DeRose of 'I know it'll be open' would be acceptable/true in this case, would not be off in responding 'No, he doesn't know, because he doesn't believe that the bank will be open any more.' But if the knowledge attribution only communicates that DeRose can rule out the salient counter-possibilities, then this justification of the speaker's falsity judgment would be missing the point.

²⁵The fact that the belief- (and truth-)conditions do not get conveyed pragmatically on the required substitutionality presents another problematic issue. Given the substitutional implicature, it should be odd to say 'S doesn't know that p because S doesn't even believe that p' in HIGH. It should be odd because all that 'S doesn't know that p' conveys is that S cannot rule out the salient counter-possibilities, which doesn't require S's belief that p. So S's failure to believe that p is no explanation of their not being able to rule out the salient counter-possibilities. The problem can be avoided by including a belief (and truth-condition) in the pragmatically implicated content.

Table 1:

	Observation vs exploitation	Additive vs substitutional
Knowledge Attributions _{HIGH}	Exploitation	Additive
Knowledge Denials _{HIGH}	Exploitation	Substitutional

4 Retraction and Cross-contextual Rejection

Brown, Hazlett, and Rysiew draw attention to the fact their accounts of basic variability data requires the attribution of a modicum of semantics-pragmatics confusion to ordinary speakers.²⁶ For instance, speakers accept knowledge denials in HIGH and are willing to ascribe (literal) truth to them. They are far less likely to hedge and say, ‘Well, strictly speaking, what you said is false, but of course you conveyed something true.’ Call this feature of their linguistic behaviour *SP-Insensitivity*:

SP-Insensitivity

Speakers are insensitive to the distinction between what the knowledge sentences they use *semantically express* (what is said, semantic content) and what *they mean* by using knowledge sentences (pragmatically conveyed content).

The problem with MPI isn’t that it is committed to *SP-Insensitivity*. Competent speakers need not always be able to distinguish between literal truth and pragmatic acceptability (truth of implicated content) and between literal falsity and pragmatic inacceptability. So we shouldn’t expect that speakers’ linguistic sensitivity affords judgments or other linguistic behaviour that tracks the distinction. We can safely assume that *SP-Insensitivity* is compatible with linguistic competence, understood broadly to include pragmatics.

The problem for MPI is that *SP-Insensitivity* isn’t the kind of ignorance that helps explain enough of the data troubling MPI. Consider patterns of cross-contextual acceptance and rejection, in which speakers, who are in HIGH, assess knowledge claims made by other speakers who are in conversational contexts with no challenging error-possibilities being raised and with low practical stakes (Low). Abe, a friend of Esta’s, has no special access to the exam committee but he was told by a reliable source that 90% of the class passed the exam and he knows that Esta is a student in the upper fifth of the class. He says in Low:

²⁶See Brown (2006, 428), Hazlett (2009, 610, 614-15, 616), Rysiew (2001, 496,502-3506-7), (Rysiew, 2007, 640, 643, 648, 660).

(4) Abe_{Low}: I know that Esta passed the exam.

Suppose Esta, whose admission to graduate school depends on the exam result and who has been pondering failure, overhears Abe's utterance of (4). It seems natural for her to say to herself:

(5) Esta_{HIGH}: No way. / That's false. He doesn't know that.

The trouble with (4) and (5) is that the semantic content and conveyed pragmatic content – if there is an implicature – in Abe's Low context in (4) are both true according to MPI. (It's not clear if an implicature is supposed to be triggered in Low contexts. If it is, it won't do much work, however, since the counter-possibilities salient in Low will not outstrip the 'relevant' counter-possibilities.²⁷) So even if Esta is SP-insensitive, this doesn't explain why she judges (4) false and why this falsity-judgments seems acceptable.²⁸

Similar trouble comes from retraction cases.

- (6) a. Abe: (looking at a zebra in a normal zoo) I know that is a zebra.
b. Esta: But can you rule out its being a cleverly painted mule?
c. Abe: I guess I can't rule that out. I was wrong earlier. I didn't know that it's a zebra.

Assuming that the conversation starts off in an ordinary Low context, the literal semantic content of Abe's knowledge attribution in (6a) (and the implicated content, if there is one) is true. So Abe has no reason to retract his knowledge attribution in (6c). Yet his retraction and concomitant assertion of a past-tense, negated knowledge sentence seem felicitous. MPI with *SP-Insensitivity* offers no account of the felicity of retraction and of HIGH-on-LOW, cross-contextual rejection of knowledge attributions.²⁹

Pragmatic explanations of use data which are not predicted by a given semantic account are often treated as better alternatives to accounts that require the attribution of non-trivial error to speakers. But MPI is itself faced with the fact that their semantic and pragmatic accounts combined cannot explain all of the use data. Barring a third option, they either need to fix their pragmatic or semantic accounts, or introduce a more substantial error theory. Giving up

²⁷Rysiew (2001, 495, 496) claims that knowledge claims '*standardly* functions to impart that the subject can rule out all the salient not-p alternatives' (my emphasis) and that the implicature is 'nearly universal', thus suggesting that knowledge claims present cases of generalized conversational implicatures. In the same vein, Lutz (2014, 1736-7) maintains that knowledge implicatures are generalized implicatures. However, Lutz also states that the implicature is generated 'in a context where a bit of practical reasoning (or subsequent action on that reasoning) is salient' (1728), and there is some room for doubt that the salience of practical reasoning is a strong default feature of contexts in which 'know that' is used. Similarly, for Hazlett (2009, 606-7), knowledge implicatures seem to require contexts in which one of a number of questions is salient: 'Does S have the ability to rule out contextually salient possibilities incompatible with p? Is S capable to giving an argument that will convince her interlocutors that p? Is S certain that p is true?'. Finally, Rysiew (2007, 641) makes it clear that in some contexts (Low), it is only the literal semantic content that is communicated. Thus, the question whether MPI's knowledge implicatures are present in Low as well as HIGH receives no clear answer. I will consider both options where it matters.

²⁸Examples similar to (4) and (5) are given in 'High Attributor-Low Subject' cases, e.g. in Stanley (2005, 115).

²⁹Dimmock & Huvenes (2014, §4) draw a very similar conclusion from retraction cases. As they note, judgments of felicitous retraction are reported by authors whose views do not seem to explain such judgments in a straightforward way (e.g., Hawthorne (2004, 163)) and by authors who do not hold such views (e.g. Feldman (2001, 77), MacFarlane (2005, 202f.)). There are some experimental results on retraction cases involving epistemic 'might' (see Knobe & Yalcin (2014) and Marques (forthcoming) for discussion), but to my knowledge there are no experimental findings on the retraction of knowledge claims.

invariantism to account for the shiftability exhibited in retraction and cross-contextual rejection should be only a last resort to the dedicated invariantist. But the pragmatics can't easily be fixed either. A closer look at (4)–(5) and (6) suggests that speakers tend to take their own current situation to matter for interpretation and truth-/acceptability judgments; Esta is influenced by the stakes of her passing the exam and the possibility of failure, Abe in (6c) takes seriously the possibility of looking at a cleverly painted mule. In standard pragmatics, however, interpretation is sensitive to factors of the *context of utterance*—within the Gricean framework, these are the speaker's intentions, above all. So unless MPI gives up on a core feature of pragmatics, moulding the pragmatic account to predict the data isn't an option.³⁰

This leaves a more substantial error hypothesis as the final option. In short, speakers are wrong about *what they mean* with their utterances. The following two theses help MPI explain retraction and cross-contextual rejection data.³¹

Ignorance and Projection Thesis (IPT)

- (a) In specific kinds of contexts, speakers are systematically ignorant of the fact that the meaning conveyed by knowledge claims is sensitive to the *context of utterance*.
- (b) In those contexts, speakers systematically project the relevant features of their own current context onto the context of utterance.

Adducing IPT, MPIers can make room for the cross-contextual rejection in (5) by claiming that Esta in HIGH brings her own stakes to bear on Abe's context of utterance – ignoring the stakes at Abe's context – and wrongly interprets Abe's knowledge attribution as implicating that Abe can rule out the error possibilities Esta is entertaining in her context. Similarly, Abe in (6c) retracts his knowledge attribution made in (6a)'s Low context because he now wrongly takes himself to have implicated earlier that he could rule out the now salient error possibility that the animal is a painted mule. We will see in section 6 why IPT is problematic.³²

³⁰Logical space has room for an invariantist view with a non-Gricean, assessment-sensitive pragmatics that mimics the predictions of relativism on retraction and cross-contextual evaluation. On the semantic side, content relativism may provide inspiration (e.g. Cappelen (2008), Weatherson (2009)); on the pragmatic side, speech act pluralism could be augmented with an assessment-sensitive rule for interpretation (Cappelen & Lepore, 2005). The norm-relativism of Greenough (2011) is yet another pragmatic option to accommodate retraction data.

³¹Retraction and cross-contextual use data have been central in the case against contextualism. It shouldn't come as a surprise here that they also pose a problem for MPI. MPI is designed to mimic the contextualist's semantic content at the pragmatic level. In either case, the content is derived in and from the context of utterance. Hence, both accounts by themselves are ill-equipped to explain data that seems to run counter to the dependence of content on the context of utterance. (Lutz (2014, 1728) points out that his pragmatic account is not modelled after the contextualist's semantics but rather after subject-sensitive invariantism. On his account, an assertion of 'S knows that p' implicates, 'in a context where a bit of practical reasoning (or subsequent action on that reasoning) is salient, [...] that S is rational to take p as true in that practical reasoning.' However, 'rational' is most certainly a context-sensitive word whose content is determined by the context of utterance. So unless Lutz subscribes to an assessment-sensitive semantics of 'rational', hearers aren't free to interpret the pragmatically conveyed content relative to their own current context, where this context is different from the speaker's context of use. Lutz's account shares the use-sensitivity of the relevant parts of the pragmatic implicature with contextualism's semantic use-sensitivity.

³²Reports of cross-contextual disagreement may also pose a problem for MPI (see Montminy (2009a)), if we can assume that incompatibility of literal contents is insufficient for disagreement. I agree with Dimmock & Huvenes (2014, 3247, fn.19) that judgments about cases of cross-contextual disagreement, in which speakers are not in conversation with each other, are somewhat contentious and that the objection from retraction and cross-contextual evaluation presents a stronger case against MPI.

5 Sceptical Paradox

Arguably a significant part of the attraction of moderate invariantism comes from its rejection of scepticism: There is a lot we know, including that we are not envatted brains. Most proponents of MPI argue that the pragmatics adds to this attraction. It helps us put scepticism in its place while explaining its recalcitrant appeal. Rysiew (2001) makes this point most explicitly, when he proposes a resolution of the sceptical argument:³³

- (SA) P1. I don't know that I'm not a BIV (i.e., a bodiless brain in a vat who has been caused to have just those sensory experiences I've had).
P2. If I don't know that I'm not a BIV, then I don't know that I have hands.
-
- C. I don't know that I have hands.

(SA) arguably presents a paradox. While the premises considered by themselves seem intuitively true, the conclusion by itself seems intuitively false. Yet, (SA) clearly appears to be a valid argument. (There are many ways to cast the phenomenology of intuitions, but I'll stick with Schiffer's (1996) presentation in his objection to contextualism. This should be fair to Rysiew, who sets up MPI's solution of (SA) in competition to contextualism's solution.)

Following Schiffer (1996) and DeRose (1995), we require two things from a 'fully satisfactory' resolution of the paradox: first, it must explain why (SA) is not in fact a paradox; second, it must explain why (SA) *appeared* to present us with a paradox.³⁴

MPI, as a form of moderate invariantism, affords a straightforward unmasking of paradox. P1 is literally false: it expresses the semantic content that it's not the case that [I'm not a BIV and I believe that I'm not a BIV and can rule out the relevant alternatives to my not being a BIV]. This is false, *inter alia*, because I *can* rule out the relevant alternatives; these include only the 'normal', 'this-worldly' alternatives, none of which is a world in which I'm an envatted (Rysiew, 2001, 499). This resolves the paradox. Contrary to appearance, P1 isn't true. (Moreover, according to MPI, P2 is literally true and C literally false.)

The pragmatics comes in for MPI in explaining the appearance of paradox. Just like DeRose claims that an assertion of P1 introduces HIGH (because of DeRose's 'rule of sensitivity'³⁵), Rysiew (2001, 499) points out that P1 makes the sceptical hypothesis that I'm a BIV salient. Hence, an assertion of P1 pragmatically implicates that I cannot rule out the salient possibility that I'm a BIV. This content is true. By *SP-Insensitivity*, speakers can't distinguish between what's literally expressed and what is pragmatically meant by an assertion of P1. They take what

³³Hazlett (2009) and Pritchard (2010) also argue that the pragmatics adds to moderate invariantism's rejection of scepticism, while Brown (2006), Lutz (2014) and Locke (forthcoming) do not explicitly argue for MPI's potential to address sceptical worries. As I point out below, there are good reasons to think that if MPI applies to non-sceptical ordinary knowledge claims, it applies to sceptical ones as well.

³⁴The second requirement receives support from Hazlett's epistemological desideratum that the 'puzzling, paradoxical, confusing' nature of scepticism be explained (Hazlett, 2009, 599).

³⁵Where the [arguably known proposition] P involved is to the effect that a sceptical hypothesis [e.g. that one is not a BIV] does not obtain, then this rule dictates that the standards will be raised to a quite high level, for [...] one must be in a stronger epistemic position with respect to a proposition stating that a sceptical hypothesis is false—relative to other, more ordinary, propositions—before a belief in such a proposition can be sensitive.' (DeRose, 1995, 36) Roughly, a subject S's belief that *p* is sensitive just in case [if it wasn't true that *p*, S would not believe that *p*].

is merely pragmatically conveyed at face value and judge P1 true. This, according to Rysiew, is a mistake, albeit an understandable mistake.³⁶

Rysiew's MPI account of sceptical paradox is exclusively concerned with P1, but this is a mistake. What about P2 and C? If P1 makes the possibility of envatment salient, it presumably remains salient throughout the entire argument. P2, I'll assume, is not just literally but also pragmatically true.³⁷ C is literally false yet pragmatically true: While I can rule out the relevant, this-worldly alternatives to my having hands, I cannot rule out the salient BIV possibility in which I do not have hands. But if (SA) in fact presents us with the intuition that C is false, MPI needs to explain why we fail to interpret it as it is meant in the context of salient sceptical hypotheses. Why do we not intuitively judge it to be true, as predicted by MPI's pragmatics, when the BIV hypothesis is salient?

In order to explain why in evaluating C we ignore the salient BIV possibility and judge it by its literal content (or what it would pragmatically implicate in a Low context in which the salient alternatives do not significantly differ from the relevant ones), MPI must again resort to the stronger error attribution of IPT. Speakers are ignorant of the fact that in the context of (SA), C pragmatically conveys a truth; when prompted with the proposition that they have hands, they are taken by their fallibilistic intuitions, which they locally project onto C in the context of (SA). They are wrong about what C pragmatically conveys in (SA).

Some readers may disagree with this way of diagnosing SA's phenomenology of paradox. One might think that in the context of reasoning through (SA), (some) speakers do judge C to be true — that's the source of any appeal scepticism in (SA) may have. If this is what is going on with (some) speakers, then what MPI needs to explain is why those speakers feel a *tension* between accepting C in the context of reasoning through (SA) and rejecting knowledge denials made with the same sentences in mundane Low contexts.³⁸ If they recognized that C has its truth-conveying implicature *in the context of (SA)* and that denials expressed by the same sentence in Low contexts do not have this implicature, they should not feel any tension. If, in contrast, they do not reliably tell HIGH from Low contexts, project the epistemologically pertinent features of their own current context onto any other context and thereby misinterpret what speakers mean by their knowledge claims made in those contexts, then they may deem their acceptance of the claims in those contexts as being in tension with their rejection of homophonic claims made in their current contexts.

I will stick with the first understanding of the phenomenology of paradox, but not much should hang on this.³⁹ Either way, the problem with MPI's resolution of paradox and its account

³⁶I'm not going to be concerned here with worries that pragmatic accounts can only account for (SA) when it is asserted but not when it is merely presented in thought. See, e.g., [Miller \(2016\)](#) for a general attack on Gricean pragmatic accounts of phenomena that can be restated in non-conversational, non-linguistic terms (and [Baumann \(2011\)](#) on the case of knowledge claims), and, e.g., [Douven \(2010\)](#) for an account of pragmatics in belief.

³⁷MPI is silent on what happens when knowledge sentences are embedded, e.g. under 'if...then'. Embedded implicatures notoriously cause trouble for Gricean implicature accounts. See [Kindermann \(2016\)](#) for the case against embedded implicatures from 'know that'. Of course, it's open to MPI to deny the closure of knowledge under entailment that is necessary for P2's truth (following [Dretske \(1970\)](#) and [Nozick \(1981\)](#)). On this way of dissolving the paradox, MPI would need an account of why P2 appears to be plausible despite its falsehood.

³⁸One might also take the situation to be more complex, with some speakers being firm dogmatists, others firm sceptics, and only some feeling any tension when prompted with (SA) (see e.g. [Blome-Tillmann \(2014\)](#)). I don't want to take a stand here, but I will assume for dialectical purposes that MPI is set to answer the very challenge from (SA) that contextualism is set to answer, given that Rysiew sees MPI to provide a better solution than contextualism.

³⁹Note that the problem for MPI isn't (any alleged peculiarity involved in) scepticism. The range of cases for which the error theory is needed extends beyond scepticism to more mundane cases. Thus, standard zebra and

of retraction and cross-contextual rejection is that they appeal to IPT: they require the attribution of an implausible form of speaker error, as I will argue below.

6 The Attribution of Pragmatic Error

Why is IPT implausible as an attribution of systematic error to speakers? Here's IPT again:

Ignorance and Projection Thesis (IPT)

- (a) In specific kinds of contexts, speakers are systematically ignorant of the fact that the meaning conveyed by knowledge claims is sensitive to the *context of utterance*.
- (b) In those contexts, speakers systematically project the relevant features of their own current context onto the context of utterance.

IPT is implausible for at least two reasons. First, IPT is implausible because it attributes an unprecedented form of speaker error. Speakers do not manifest similarly systematic lapses in analogous, paradigmatic cases of implicature. Recall that not any analogy to Gricean implicature will do. Knowledge denials in HIGH have substitutional exploitation implicatures. Consider two examples of substitutional exploitation implicatures – hyperbole (7) and metaphor (8) – occurring in arguments similar in structure to (SA):⁴⁰

- (7) 5th-graders have a ton of homework.
If 5th-graders have a ton of homework,
then 6th-graders have a ton of homework.

6th-graders have a ton of homework.

- (8) Abe can learn the ropes quickly.
If Abe can learn the ropes quickly, so can Ben.

Ben can learn the ropes quickly.

I submit that the arguments in (7) and (8) have no ring of paradox whatsoever, even when 'a ton of' is interpreted hyperbolically as meaning (roughly) 'a lot of' and 'learn the ropes' metaphorically as (roughly) 'learn the basics.' They seem innocuously valid, and where we have become convinced of the premises, further reflexion doesn't tempt us to reject the conclusion. But if paradigmatic examples of implicature do not give rise to IPT-like speaker error, then the needed

carpark cases, when put in argument form, provide cases of closure puzzles that resemble the modus-ponens version of (SA): 'I know this is a zebra. If I know this is a zebra, then I know it's not a painted mule. Therefore, I know this is not a painted mule.' While the air of paradox might be less pressing in such cases, the remaining puzzlement requires explanation, which MPI would need strong reasons not to find in the combination of their pragmatics plus an error theory.

⁴⁰These examples involve the exploitation of the maxim of Quality, which says: 'Try to make your contribution one that is true'. (Grice, 1975, 27) I'm not aware of paradigmatic examples of substitutional exploitation implicatures that are also involve Relation, as in the case of knowledge claims. Grice says surprisingly little about the maxim of Relation despite deeming it important in the generation of implicatures (though see Grice (1975, 35)). But perhaps generation by Relation is the feature we can drop most readily.

attribution of speaker error for ‘know-that’ looks like an *ad hoc* strategy to save the data that MPI’s combined semantic-cum-pragmatic account does not correctly predict.⁴¹ At the very least, an explanation is needed for why knowledge implicatures would give rise to speaker error while analogous paradigmatic cases of implicature do not.

Paradigmatic examples of implicature also do not give rise to retraction and cross-contextual evaluation in ways analogous to knowledge claims. Consider Esta’s intuitively felicitous rejection of Abe’s knowledge attribution again:

- (4) Abe_{LOW}: I know that Esta passed the exam.
(5) Esta_{HIGH}: No way. / That’s false. He doesn’t know that.

In order to capture the felicity of Esta’s falsity claim in (5), MPI needed to appeal to her wrongly projecting the high standards of her own context onto Abe’s context, mistaking his purely literal contribution for one in which he also implicates, as his assertion would in HIGH, that he can rule out the error possibilities salient at Esta’s context in (5). The error attribution involves Esta’s mistaken treatment of (4) as triggering an additive exploitation implicature (cf. section 3). But if we are looking for roughly analogous cases of cross-contextual evaluation, where the speaker is in a context in which the evaluated assertion would have a paradigmatic additive exploitation implicature, we do not find the kind of linguistic behaviour, and judgments, that call for the attribution of error. A paradigmatic example of an additive exploitation implicature is Grice’s well-known reference letter case, in which all a professor writes about a job candidate is that

- (9) Mr. X’s command of English is excellent, and his attendance at tutorials has been regular. (Grice, 1975, 33)

In the context of a reference letter, (9) is taken to convey not only what it literally says but also that the candidate is no good at philosophy. But suppose (9) is asserted by X’s philosophy professor in a conversation about whether X, a non-native speaker, has been able to follow class. Clearly, in this conversation, (9) would not generate the implicature about X’s lack of philosophical skill.⁴² Suppose further that Esta, who is wondering about X’s philosophical skill, overheard the assertion (9) made in this conversation, were to react by saying ‘No way/That’s false. X is the best philosophy student in his class.’ It seems clear that Esta’s reaction is not a felicitous evaluation of the professor’s assertion. But if Esta betrayed the same kind of error that led her to assert (5), then we might expect her to make this mistake with other, paradigmatic examples of additive exploitation implicatures, too. Moreover, if Esta’s evaluation in (5) seems felicitous to us, we are prone to IPT-error. So why then, are we not prone to IPT-like error in standard cases of such implicatures?

⁴¹ Note that (7) but not (8) shares with (SA) that the implicated content – roughly, that 5th/6th-graders have a lot of homework – involves context-sensitive material. But while what counts as a lot of homework may vary from one context to another, the factors influencing variation (such as comparison class: 5th-/6th-graders in Britain, Sweden, and Haiti, or in 1950, 1990, 2010) do not, despite often being implicit, lead to comparable error.

Note further that in presenting (7) and (8) as paradigmatic cases of substitutional exploitation implicatures, I have not argued that *no* standard cases of such implicatures give rise to puzzlement when occurring in short modus-ponens arguments. But with (7) and (8) being paradigmatic cases of such implicatures, the burden of proof is on MPI to provide examples that show that the kind of error needed for the explanation of (SA) is not limited to the knowledge implicatures.

⁴²It might generate an implicature that directly answers the question: X has been able to follow class.

Of course, the analogy to this cross-contextual reference letter case doesn't conclusively show that there are no standard cases of additive exploitation implicatures whatsoever that give rise to IPT-type speaker error. It does lend support to this conclusion, however, if standard additive exploitation implicatures are similar to the reference letter case. And, I submit, retraction-like cases with paradigmatic additive exploitation implicatures that are analogous to the zoo case in (6) also do not license retraction in the way the retraction of the preceding knowledge attribution is licensed in (6). I therefore conclude that the attribution of IPT's pragmatic error for (5) is unprecedented and *ad hoc*.

The second reason for IPT's implausibility comes with a principled worry. IPT's attribution of speaker error doesn't sit well with the orthodox understanding of speaker meaning. MPI requires the claim that speakers are systematically (i.e. persistently in retraction, cross-contextual evaluation and sceptical paradox contexts) wrong about what they themselves convey in communication: In contexts like (SA) (or when reflecting on (SA) and its conclusion), in retraction and cross-contextual evaluation, speakers are supposed to be wrong about what they pragmatically convey (or don't convey); in Gricean terms, they are wrong about their own speaker meaning. There might not be an aseptic distinction between pragmatics and semantics, and there might thus be difficulties with cleanly telling pragmatic error from semantic error. Yet, MPIers help themselves to Grice's rough-and-ready distinction between literal meaning (semantics, 'what is said'), what is implicated (pragmatics) and speaker meaning (which can be literal meaning, pragmatic meaning, or a combination of both). On standard accounts of speaker meaning, a speaker means *p* only if they *intend* to communicate *p* (put roughly).⁴³ According to MPI's error attribution, then, speakers are sometimes systematically wrong about their own intentions. Our own communicative intentions may sometimes be intransparent to us, but the thesis that in particular kinds of contexts, speakers are systematically in error about what they intended to communicate at the very least requires further argument.⁴⁴ Rysiew seems to share the worry but believe MPI safe from it:

I find it manifestly *implausible* to suppose that speakers can be wrong about what it is that *they* mean—about what their communicative intentions are—in uttering certain sentences. (Rysiew, 2001, 483; emphasis in original)

As I have argued, retraction, cross-contextual evaluation, and sceptical paradox cases show that MPI is committed to the stronger error attribution of IPT, committing MPI to exactly the claim that speakers are wrong about what their communicative intentions are in those cases.

7 Against Strict Pragmatic Invariantism

My target in this paper is the pragmatic strategy of calling on pragmatic contents to do service towards shifty patterns of language use in order to save the classical invariantist (semantic) analysis of 'know that'. I have focused on moderate pragmatic invariantism. But it should not come

⁴³See, e.g., Grice (1957, 1969) and Bach & Harnish (1979)

⁴⁴The worry is not that speakers are not aware of, or cannot articulate their communicative intentions. Our intentions very often do not offer themselves for exact linguistic articulation. But communicative intentions should be available for linguistic action including retraction, rejection, and truth-value judgments, and perhaps dispositions to cancel implicatures (cf. Dimmock & Huvenes (2014)). The worry is thus that competent speakers' linguistic actions require that they have different communicative intentions than those MPI is committed to attributing.

as a surprise that similar objections beset strict or sceptical pragmatic invariantisms (SPI), revealing a parallel commitment to IPT. Here, I air sample objections against Davis (2004, 2007, 2015), one of the main proponents of SPI.⁴⁵

SPI is the view that most of our ordinary knowledge claims are false. But while their literal, semantic content is false, ‘know that’ may be used *loosely* to pragmatically implicate that the subject is *close enough* to knowing. Davis (2007) explains the *Basic Variability* in our use of knowledge sentences in terms of differences between strict and loose usage of knowledge sentences. Consider a case of loose use:

(10) THE TIME MEASUREMENT CASE

A. Wondering how hard the final exam was, I ask Mike how long he took to finish. He answers ‘Two hours.’ B. When Nora says that she took two hours and four minutes to finish the exam, Mike responds ‘You took even longer than me. It took me two hours and two minutes.’ (Davis, 2007, 407)

On Davis’ account, if what Mike said in B is true, what he said in A is false. But in A, Mike used ‘two hours’ loosely. What he meant (conveyed, implicated) is that it took him *about* two hours to finish. Once the conversational purposes require more precision in B, Mike gives more precise information. In both conversational contexts, Mike’s utterance is perfectly appropriate. While the literal content of his utterance in A is false, the implicated content is true.

According to Davis, knowledge sentences can also be used loosely and strictly. Standards of strictness vary with contexts of utterance and allow for loose usage in Low contexts. While what is said by ‘I know that is a zebra’ in Low is false—subjects are hardly ever in good enough an epistemic position to count as knowing—what is implicated is that the speaker is close enough to knowing for contextually salient purposes, which may well be true. When the purposes of the conversation do not require us to commit to a subject’s knowing, it is sufficient to use knowledge attributions loosely and communicate that the subject is close enough to knowing. When the difference between knowing and being close enough to knowing matters to the purposes of the conversation, we use knowledge attributions strictly, i.e. we convey only their literal semantic content.

SPI runs into trouble with speakers who in a Low context intuitively reject knowledge denials made in HIGH as false. According to SPI, an assertion of ‘I don’t know that I have hands’ in HIGH makes strict, non-loose, use of the negated knowledge sentence. The literal content is true, and no implicature is generated. But if it’s natural for a speaker who is in Low and overhears the assertion in HIGH to reject the knowledge denial as false, that’s bad news for SPI.

Perhaps Davis can exploit his pragmatic story to explain rejection in this kind of case. Then he must claim that speakers erroneously take the negated knowledge sentence in HIGH to have been used loosely, *in the way it would have been used in the Low context*, implicating that the speaker in HIGH is not close enough to knowing for the purposes *salient in their Low context*. This implicature is false, hence the rejection. This explanation, however, requires the attribution of error as given in IPT.

Retraction—as in (6), repeated here—is also problematic.

⁴⁵This section is based on material in the third chapter of my dissertation (Kindermann, 2012), which includes a critical discussion of Schaffer’s (2004b) version of SPI.

- (6) a. Abe: (looking at a zebra in a normal zoo) I know that is a zebra.
 b. Esta: But can you rule out its being a cleverly painted mule?
 c. Abe: I guess I can't rule that out. I was wrong earlier. I didn't know that it's a zebra.

On the face of it, (6) should be data in favour of SPI. Abe in (6c) comes to realize that his earlier knowledge claim is strictly and literally false and hence retracts. However, SPI loses this straightforward explanation of retraction. Davis makes it clear that

[u]sing '*p*' loosely involves saying that *p* while meaning only that it is close enough to being the case that *p*. In loose usage, the speaker does not intend to commit himself to what he strictly speaking said. Hence he cannot be criticised for being sloppy or ignorant. (Davis, 2007, 411)

If speakers using knowledge sentences loosely in Low contexts never intend to commit themselves to what they literally say – that they *know* that they have hands – but only to what they mean, and if they cannot be criticised for being sloppy or ignorant, then they should not give in so easily and concede that they were *wrong*. After all, they cannot be criticised for their loose use. What they committed themselves to in the context of utterance was only the weaker claim that they are close enough to knowing for the purposes of their conversation at that stage. And that is still true. Thus, once under pressure, speakers should be at pains to point out, with words that they have available, that what they meant is true, even if strictly speaking, what they said is false. But if speakers do in fact use 'know' loosely, it must be that they do not have recourse to such defensive replies because they unwittingly project what the knowledge sentence would mean *were it used in their current context* onto the earlier context. Note, though, that they do not make this mistake in standard cases of loose use. If it was pointed out to Mike in (10) after his initial assertion of 'It took me two hours' that in fact it took him exactly two hours and two minutes, the natural reaction for Mike is not to concede that he was wrong earlier but rather to accept the precisification and to clarify that he only meant that it took him roughly two hours.⁴⁶

Finally, sceptical arguments also at first seem to play into SPI's hands. On SPI, arguments like (SA) are sound. Yet, if SPIers accept the challenge to explain the *appearance* of paradox, then they incur a commitment to IPT. Why do speakers balk at accepting the conclusion of (SA) as true—or why do they feel a tension between accepting the conclusion in the context of reasoning through (SA) and accepting its negation in ordinary, non-sceptical contexts (Low)? Presumably, this is because they interpret 'I know that I have hands' loosely as meaning that the speaker is close enough to knowing for the purposes salient in Low and they fail to realize that there is no tension between accepting that and accepting the negated sentence in (SA)'s HIGH, where only the strict literal content is expressed. They erroneously take the knowledge claims to express whatever they would express in the speakers' own current context.

8 Conclusion

I have argued that the pragmatic invariantist strategy of adducing pragmatically implicated contents to account for the contextual variability in speakers' use of knowledge sentences fails to

⁴⁶Cf. MacFarlane (2011, 541) and the 'awareness objection' to SPI in Dinges (2015b).

predict speakers' cross-contextual rejections and retractions of knowledge claims and their judgments leading to the appearance of paradox in sceptical arguments. Pragmatic invariantists are committed to the attribution of pragmatic error to speakers: speakers systematically project the relevant features of their own current context onto the context of utterance (if distinct from their own) and fail to be sensitive to the fact that the implicatures generated by (some) knowledge claims are sensitive to the claims' context of utterance; instead they mistakenly interpret knowledge claims as conveying what the utterer would have meant had they uttered it in the speaker's own context. The attribution of pragmatic error is implausible because it is an *ad hoc* attribution of error speakers do not exhibit in analogous cases of pragmatic implicature, and because, given standard accounts of pragmatic meaning as involving communicative intentions, speaker would be systematically wrong, in certain contexts, about what they themselves intend to communicate.

If my arguments are on the right track, then pragmatics offers no alternative to error theories, as needed by competing strategies. In particular, the pragmatic account of many proponents of MPI founders on exactly the data that troubles traditional contextualism. Here, I'm not making any attempt at adjudicating between the pragmatic strategy and the semantic, metaphysical, and error strategies. If *all* extant strategies require attributions of some form of speaker error,⁴⁷ then everyone is in need of a convincing (psychological) explanation of the error that fits a view's attributed pattern of error. So the error objection in this paper need not be seen as a knock-down objection, but can be understood as a challenge to pragmatic invariantists to provide a convincing (psychological) error theory that favours invariantism over its competitors.

Acknowledgement

For discussion and feedback on various versions of this paper and on precursor material, I'm indebted to Michael Blome-Tillmann, Cameron Bault, Jochen Briesen, Fernando Broncano-Berrocal, Alexander Dinges, Andy Egan, Harmen Ghijsen, Patrick Greenough, Frank Hoffmann, Torfinn Huvenes, Christoph Kelp, John MacFarlane, Mona Simion, Erik Stei, and audiences at the Universities of Konstanz and Leuven. I would especially like to thank Robin McKenna, Alexander Dinges, Chris Gauker and an anonymous referee for detailed comments on the paper. This research was supported by the Austrian Science Fund (FWF), project P 27587-G15.

References

- Bach, Kent. 1994. Conversational Implicature. *Mind and Language*, **9**, 124–62.
- Bach, Kent. 2010. Knowledge In and Out of Context. *Pages 105–36 of: Campbell, J. K., O'Rourke, M., & Silverstein, H. (eds), Knowledge and Skepticism. Topics in Contemporary Philosophy*. Cambridge, MA: MIT Press.
- Bach, Kent, & Harnish, Robert. 1979. *Linguistic Communication and Speech Acts*. Cambridge, MA: MIT Press.
- Baumann, Peter. 2011. WAMs: Why Worry? *Philosophical Papers*, **40**(2), 155–77.

⁴⁷See Kindermann (2013) and Montminy (2009b) on relativism and Kindermann (2012) for the universal claim.

- Beebe, James R. (ed). 2014. *Advances in Experimental Epistemology*. London and New York: Bloomsbury Publishing.
- Black, Tim. 2005. Classic Invariantism, Relevance and Warranted Assertability Manœuvres. *The Philosophical Quarterly*, **55**(219), 328–36.
- Black, Tim. 2008. A Warranted-assertability Defense of a Moorean Response to Skepticism. *Acta Analytica*, **23**(3), 187–205.
- Blome-Tillmann, Michael. 2013. Knowledge and Implicatures. *Synthese*, **190**(18), 4293–4319.
- Blome-Tillmann, Michael. 2014. *Knowledge and Presuppositions*. Oxford: Oxford University Press.
- BonJour, Laurence. 2010. The Myth of Knowledge. *Philosophical Perspectives*, **24**(1), 57–83.
- Brown, Jessica. 2005. Adapt or Die: the Death of Invariantism? *The Philosophical Quarterly*, **55**(219), 263–85.
- Brown, Jessica. 2006. Contextualism and Warranted Assertability Manœuvres. *Philosophical Studies*, **130**(3), 407–435.
- Cappelen, Herman. 2005. Pluralistic Skepticism. *Philosophical Perspectives*, **19**.
- Cappelen, Herman. 2008. The Creative Interpreter: Content Relativism and Assertion. *Philosophical Perspectives*, **22**, 23–46.
- Cappelen, Herman, & Lepore, Ernie. 2005. *Insensitive Semantics*. Malden, MA: Blackwell.
- Chrisman, Matthew. 2007. From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies*, **135**(2), 225–54.
- Cohen, Stewart. 1986. Knowledge and Context. *The Journal of Philosophy*, **83**(10), 574–583.
- Cohen, Stewart. 1987. Knowledge, context, and social standards. *Synthese*, **73**(1), 3–26.
- Cohen, Stewart. 1999. Contextualism, Skepticism, and the Structure of Reasons. *Noûs*, **33**, 57–89.
- Conee, Earl. 2005. Contextualism Contested. *Pages 47–56 of: Steup, Matthias, & Sosa, Ernest (eds), Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Davis, Wayne. 2007. Knowledge claims and context: loose use. *Philosophical Studies*, **132**(3), 395–438.
- Davis, Wayne A. 2004. Are Knowledge Claims Indexical? *Erkenntnis*, **61**(2), 257–81.
- Davis, Wayne A. 2015. Knowledge claims and context: belief. *Philosophical Studies*, **172**(2), 399–432.
- DeRose, Keith. 1992. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, **52**, 913–29.
- DeRose, Keith. 1995. Solving the Skeptical Problem. *The Philosophical Review*, **104**(1), 1–52.
- DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Oxford: Oxford University Press.
- Dimmock, Paul, & Huvenes, TorfinnThomesen. 2014. Knowledge, conservatism, and pragmatics. *Synthese*, **191**(14), 3239–69.
- Dinges, Alexander. 2015a. Innocent implicatures. *Journal of Pragmatics*, **87**, 54–63.
- Dinges, Alexander. 2015b. Skeptical pragmatic invariantism: good, but not good enough. *Synthese*, 1–17.
- Douven, Igor. 2007. A Pragmatic Dissolution of Harman's Paradox. *Philosophy and Phenomenological Research*, **74**(2), 326–45.
- Douven, Igor. 2010. The Pragmatics of Belief. *Journal of Pragmatics*, **42**, 35–47.
- Dretske, Fred I. 1970. Epistemic Operators. *The Journal of Philosophy*, **67**(24), 1007–1023.
- Fantl, Jeremy, & McGrath, Matthew. 2002. Evidence, Pragmatics, and Justification. *The Philo-*

- sophical Review*, **111**(1), 67–94.
- Fantl, Jeremy, & McGrath, Matthew. 2009. *Knowledge in an Uncertain World*. Oxford: Oxford University Press.
- Feldman, Richard. 2001. Skeptical Problems, Contextualist Solutions. *Philosophical Studies*, **103**(1), 61–85.
- Gerken, Mikkel. 2012. On the Cognitive Bases of Knowledge Ascriptions. *Pages 140–70 of:* Brown, Jessica, & Gerken, Mikkel (eds), *Knowledge Ascriptions*. Oxford: Oxford University Press.
- Gerken, Mikkel. 2013. Epistemic Focal Bias. *Australasian Journal of Philosophy*, **91**(1), 41–61.
- Greenough, Patrick. 2011. Truth-Relativism, Norm-Relativism, and Assertion. *In:* Brown, Jessica, & Cappelen, Herman (eds), *Assertion*. Oxford: Oxford University Press.
- Grice, H. P. 1957. Meaning. *The Philosophical Review*, **66**(3), pp. 377–388.
- Grice, Paul. 1969. Utterer's Meaning and Intentions. *The Philosophical Review*, **68**, 147–77. Reprinted in Grice (1989), 86–116).
- Grice, Paul. 1975. Logic and Conversation. *Pages 41–58 of:* Cole, Peter, & Morgan, Jerry (eds), *Syntax and Semantics 3: Speech Acts*. New York: Academic Press.
- Grice, Paul. 1989. *Studies in the Way of Words*. Cambridge, MA, and London: Harvard University Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hazlett, Allan. 2009. Knowledge and Conversation. *Philosophy and Phenomenological Research*, **78**(3), 591–620.
- Ichikawa, Jonathan. 2010. Quantifiers and Epistemic Contextualism. *Philosophical Studies*, **155**, 383–98.
- Kindermann, Dirk. 2012. *Perspective in Context. Relative Truth, Knowledge, and the First Person*. Ph.D. thesis, University of St Andrews, URL <<http://hdl.handle.net/10023/3164>>.
- Kindermann, Dirk. 2013. Relativism, sceptical paradox, and semantic blindness. *Philosophical Studies*, **162**(3), 585–603.
- Kindermann, Dirk. 2016. *Embedding Knowledge*. Manuscript.
- Knobe, Joshua, & Yalcin, Seth. 2014. Epistemic Modals and Context: Experimental Data. *Semantics and Pragmatics*, **7**(10), 1–21.
- Kompa, Nicola. 2002. The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien*, **64**, 79–96.
- Lewis, David. 1996. Elusive Knowledge. *Australasian Journal of Philosophy*, **74**(4), 549–67.
- Locke, Dustin. forthcoming. Implicature and Non-Local Pragmatic Encroachment. *Synthese*, 1–24.
- Lutz, Matt. 2014. The Pragmatics of Pragmatic Encroachment. *Synthese*, **191**(8), 1717–1740.
- MacFarlane, John. 2005. The Assessment-Sensitivity of Knowledge Attributions. *Pages 197–233 of:* Gendler, Tamar Szabó, & Hawthorne, John (eds), *Oxford Studies in Epistemology 1*. Oxford: Oxford University Press.
- MacFarlane, John. 2009. Nonindexical Contextualism. *Synthese*, **166**, 231–50.
- MacFarlane, John. 2011. Relativism and Knowledge Attributions. *In:* Bernecker, Sven, & Pritchard, Duncan (eds), *Routledge Companion to Epistemology*.
- MacFarlane, John. 2014. *Assessment Sensitivity. Relative Truth and Its Applications*. Oxford: Oxford University Press.
- Marmaridou, Sophia S.A. 2000. *Pragmatic Meaning and Cognition*. Amsterdam/Philadelphia: John Benjamins Publishing.

- Marques, Teresa. forthcoming. Retractions. *Synthese*.
- Meibauer, Jörg. 2009. Implicature. *Pages 365–78 of*: Mey, Jacob L. (ed), *Concise Encyclopedia of Pragmatics*. Amsterdam: Elsevier.
- Meibauer, Jörg. 2014. *Lying at the Semantics-Pragmatics Interface*. Berlin and Boston: De Gruyter Mouton.
- Miller, Zachary. 2016. The Reformulation Argument: Reining in Gricean Pragmatics. *Philosophical Studies*, **173**(2), 525–46.
- Montminy, Martin. 2009a. Contextualism, Invariantism, and Semantic Blindness. *Australasian Journal of Philosophy*, **87**(4), 639–57.
- Montminy, Martin. 2009b. Contextualism, Relativism and Ordinary Speakers' Judgments. *Philosophical Studies*, **143**(3), 341–56.
- Nagel, Jennifer. 2010a. Epistemic Anxiety and Adaptive Invariantism. *Philosophical Perspectives*, **24**(1), 407–35.
- Nagel, Jennifer. 2010b. Knowledge Ascriptions and the Psychological Consequences of thinking about Error. *The Philosophical Quarterly*, **60**(239), 286–306.
- Neta, Ram. 2003. Skepticism, Contextualism, and Semantic Self-Knowledge. *Philosophy and Phenomenological Research*, **67**, 396–411.
- Nozick, Robert. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pritchard, Duncan. 2010. Contextualism, Skepticism, and Warranted Assertability Maneuvres. *Pages 85–104 of*: Keim-Campbell, J., O'Rourke, M., & Silverstein, H. (eds), *Knowledge and Skepticism*. Cambridge, MA: MIT Press.
- Richard, Mark. 2004. Contextualism and Relativism. *Philosophical Studies*, **119**(1), 215–42.
- Richard, Mark. 2008. *When Truth Gives Out*. Oxford: Oxford University Press.
- Rysiew, Patrick. 2001. The Context-Sensitivity of Knowledge Attributions. *Noûs*, **35**(4), 477–514.
- Rysiew, Patrick. 2005. Contesting Contextualism. *Grazer Philosophische Studien*, **69**(1), 51–70.
- Rysiew, Patrick. 2007. Speaking of Knowing. *Noûs*, **41**(4), 627–62.
- Schaffer, Jonathan. 2004a. From Contextualism to Contrastivism. *Philosophical Studies*, **119**(1), 73–103.
- Schaffer, Jonathan. 2004b. Skepticism, Contextualism, and Discrimination. *Philosophy and Phenomenological Research*, **69**(1), 138–55.
- Schaffer, Jonathan, & Szabó, Zoltán Gendler. 2014. Epistemic Comparativism: A Contextualist Semantics for Knowledge Ascriptions. *Philosophical Studies*, **168**(2), 491–543.
- Schiffer, Stephen. 1996. Contextualist Solutions to Scepticism. *Proceedings of the Aristotelian Society*, **96**, 317–33.
- Stanley, Jason. 2005. *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Weatherson, Brian. 2005. Can We Do Without Pragmatic Encroachment? *Philosophical Perspectives*, **19**, 417–43.
- Weatherson, Brian. 2009. Conditionals and indexical relativism. *Synthese*, **166**(2), 333–357.
- Williamson, Timothy. 2005. Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *The Philosophical Quarterly*, **55**(219), 213–235.