

Skeptical pragmatic invariantism: good, but not good enough

Alexander Dinges

(forthcoming in *Synthese*)

1. Introduction

In this paper, I will discuss what I will call “skeptical pragmatic invariantism” (SPI) as a potential response to the intuitions we have about scenarios such as the so-called bank cases. SPI, very roughly, is a form of epistemic invariantism that says the following: The subject in the bank cases doesn’t know that the bank will be open. The knowledge ascription in the low standards case seems appropriate nevertheless because it has a true implicature.¹ The goal of this paper is to show that SPI is mistaken. In particular, I will show that SPI is incompatible with reasonable assumptions about how we are aware of the presence of implicatures. Such objections are not new, but extant formulations are wanting for reasons I will point out below. One may worry that refuting SPI is not a worthwhile project given that this view is an implausible minority position anyway. To respond, I will argue that, contrary to common opinion, other familiar objections to SPI fail and, thus, that SPI is a promising position to begin with. The structure of the paper is as follows: First, I will spell out SPI in more detail. This will already take care of some of the more obvious worries with the position. Second, I will discuss and reject a range of further, less obvious worries. Finally, I will present my take on the just mentioned awareness worry to show that SPI must be rejected in spite of its good standing so far.

2. What is SPI?

What is SPI? SPI is a form of epistemic invariantism. Thus, according to SPI, sentences of the form “*S* knows that *p*” express the proposition that *S* knows that *p* independently of the epistemic standard of the context of use.² SPI differs from other invariantist positions in that it entails a specific response to the puzzle from the bank case intuitions, so let us get clear on what

¹ SPI, so understood, is defended in (Conee 2005a: 52 f.), (Douven 2007) and (Davis 2004, 2007, 2010, 2015). Bach (2010) “suggests, without endorsing,” it. Fantl and McGrath (2009: 185–194) claim that the view is “not out of the question.” BonJour (2010: 78) defends what may be seen as a version of SPI, but presents his version of SPI only as a “tentative conclusion.” Schaffer (2004) defends a version of SPI, but has retracted the view by now. Note that Bach would couch his view in terms of implicatures rather than implicatures. The difference between implicatures and implicatures will mostly be irrelevant for my concerns, but see below for some further remarks.

² To be precise, this thesis distinguishes epistemic invariantism only from epistemic contextualism. Epistemic invariantism is also supposed to be incompatible with all forms of epistemic relativism. Thus, epistemic invariantism also entails, first, that the proposition expressed by knowledge sentences doesn’t vary with the epistemic standard of the context of assessment and, second, that the truth-value of that proposition is neither relative to the epistemic standard of the context of use nor to the epistemic standard of the context of assessment. See e.g. (MacFarlane 2005) for further discussion of the various forms of relativism.

the bank cases and the corresponding intuitions are. Here are the *bank cases* as presented by DeRose:

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 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 f.)

Assume that the bank will in fact be open in both of the above cases. Now, the intuitions we supposedly have about these cases are the following: On the one hand, it seems appropriate (or even true) for DeRose in Bank Case A (“the low standards case”) to say, “I know the bank will be open.” On the other hand, it seems appropriate (or even true) for DeRose in Bank Case B (“the high standards case”) to say, “I don’t know the bank will be open.” These data seem puzzling for the following simple reason: DeRose’s epistemic position seems to be the same in both bank cases. So, he should either know that the bank will be open in both cases or fail to know that the bank will be open in both cases. On either assumption, however, one of his knowledge claims is false, not true, as it intuitively seems.

Proponents of SPI respond to this puzzle in the following way: DeRose indeed fails to know that the bank will be open in the bank cases. This means that his knowledge denial in the high standards case is literally true and that his knowledge ascription in the low standards case is literally false (because invariantism is correct). The knowledge denial seems appropriate (or even true) simply because it is true. The knowledge ascription seems appropriate (or even true), not because it is true, but because it has a true implicature. These are the basic ideas behind SPI, but some clarificatory remarks might be in order.

One immediate worry with SPI could be that a true implicature can never make a literally false claim appear appropriate (or even true) because adding a truth to a falsity will result in just another falsity. This worry misconstrues SPI. To see this, it is useful to introduce a distinction between what we may call additive and substitutional implicatures (following (Meibauer 2009: 374)). *Additive implicatures* are implicatures where the speaker conveys what is said and

something else in addition.³ To take a familiar example, suppose you ask me where to get petrol, and I respond, “There is a garage around the corner.” Arguably, I will thereby implicate that the garage is likely to be open. This implicature is additive because I will convey that there is a garage around the corner *and* that it might be open. *Substitutional implicatures*, on the other hand, are implicatures where the speaker doesn’t convey what is said but something else instead. To take another familiar example, suppose I say, “The plane was a mile long!” In a suitable context, I will thereby implicate that the plane was huge. This implicature is substitutional because I will only convey that the plane was huge (not that it also was a mile long).

Obviously, additive implicatures can never make a literally false claim appropriate because, as just indicated, adding a truth to a falsity will lead to another falsity. Similar worries, however, do not apply to substitutional implicatures, for substitutional implicatures don’t add up to what is said but substitute it. So, our proponent of SPI will say that DeRose’s knowledge ascription in the low standards case seems appropriate (or even true) because it has a true *substitutional* implicature.⁴

What is the content of this implicature? In principle, the implicature could be spelled out in various ways, but proponents of SPI typically hold that the knowledge claim in the low standards scenario is a case of *loose use*.⁵ Thus, DeRose truly implicates that he is *close enough (for the purposes of the low standards case)* to knowing that the bank will be open. Similar implicatures arise in many other situations. For example, when I say, “It’s three o’clock,” I will often implicate that it is only *close enough (for present purposes)* to three o’clock. And when I say, “France is hexagonal,” I will often implicate that France is only *close enough (for present purposes)* to being hexagonal.⁶

³ Conveying that *p* is supposed to be roughly equivalent to intending to make one’s audience believe that *p*.

⁴ One may worry that even substitutional implicatures cannot make a literally false claim appear *true* (rather than appropriate). But this worry rests on shaky intuitions. First, it is unclear whether DeRose’s claim in the bank cases would indeed seem true rather than appropriate. See e.g. (Davis 2010: 1154). Thus, it is unclear whether the implicature needs to make his claim appear true. Second, it is unclear whether the relevant implicatures (namely, loose use implicatures, see below) cannot make a literally false claim appear true. See e.g. (DeRose 2012: 716) for the intuition that they can.

⁵ Schaffer (2004) is an exception. Note, though, that the worry I will level against SPI below does not depend on the above precisification of the position.

⁶ There are some controversies surrounding the correct analysis of loose use. Bach (1994), for example, treats loose use in terms of implicatures, Davis (2007) treats it in terms of implicatures. DeRose (2012: 714 ff.) suggests that loose use might be a semantic rather than pragmatic phenomenon. To sort out these issues, we would have to delve deeply into the discussion on the semantics-pragmatics distinction. This would go way beyond the scope of this paper. Since, however, the pragmatic account of loose use seems much more prominent than the semantic account, it should be fair to just assume that some such account is correct. Whether, in the end, the account involves implicatures or implicatures will mostly be irrelevant for the discussion to come. See below.

The idea of being close enough to knowledge can be spelled out in various ways. I will say below how I think it should be spelled out. Before that, however, let me issue a warning that two accounts of this notion that have come up in the literature won't do.

The first proposal has it that S is close enough (for the purposes of the low standards case) to knowing that p iff, *unless something unlikely happens*, S knows that p . (Here, the “unless”-clause is outside the scope of the knowledge operator.)^{7,8} This proposal is untenable because it doesn't yield a true implicature for the knowledge ascription in the low standards case: Assume that the above proposal holds and the corresponding implicature is true, that is, that unless something unlikely happens, DeRose knows that the bank will be open in the low standards case. Given that nothing unlikely happens in the low standards case, it follows that DeRose knows that the bank will be open in the low standards case. According to SPI, however, DeRose doesn't know that the bank will be open in the low standards case. Thus, the relevant implicature cannot be true in the low standards case.

The second proposal to interpret the loose use view has it that S is close enough (for the purposes of the low standards case) to knowing that p iff S knows that, *unless something unlikely happens*, p . (Here, the “unless”-clause is within the scope of the knowledge operator.)⁹ This proposal may deliver the right results for the bank cases as originally described. However, it fails to account for our intuitions about closely related scenarios. Suppose that, in a scenario otherwise like the low standards case, the unlikely happened and, unbeknownst to DeRose and his wife, the bank did change its hours and won't be open on the next Saturday. Clearly, it would now seem intuitively *false* for DeRose to say, “I know that the bank will be open,” for the bank just won't be open. On the present account, however, the claim would have to seem true. This is so because DeRose should still implicate that he knows that, unless something unlikely happens, the bank will be open. And this implicature should still be true. DeRose should still implicate the relevant claim because his conversational context is the same in the

⁷ Bach (2010: 122f) “suggests without endorsing” this proposal. To be precise, Bach uses the phrase “provided things are normal” instead of the phrase “unless something unlikely happens.” Nothing should depend on this modification however. I have made this modification only for the sake of unity.

⁸ Note that not just any “unlikely” event is going to be relevant in the present context. I will take it to be tacitly understood that we are only concerned with those unlikely events that entail that p is not the case. As regards DeRose's knowledge ascription in the bank cases, for example, a relevant unlikely event would be that the bank changes its hours. An irrelevant unlikely event would be, say, that an elephant sings the national anthem. The same goes for the second proposal to interpret the loose use implicature below.

⁹ Douven (2007: 333 n.) endorses this view. See also (BonJour 2010: 73 f.). Interestingly enough, Douven explicitly considers the first interpretation of the loose view discussed above but dismisses it without further argument. Bach, on the other hand, (who, remember, suggested the first interpretation) also considers the just mentioned second interpretation but dismisses this interpretation without further argument. (He just claims that “it might seem extreme to suppose that most of what ordinarily passes for categorical knowledge is really conditional in content.” (2010: 123))

original and the modified low standards case (and the presence of an implicature plausibly depends solely on the conversational context). The implicature should still be true because whether DeRose knows that, *unless something unlikely happens*, the bank will be open should be independent of whether the unlikely did happen. Note also that the proponent of SPI cannot appeal to the literal falsity of what DeRose says in the modified low standards case to account for the falsity intuition. For if she did, she could no longer coherently maintain that our intuitions about the original low standards case are unaffected by the literal falsity of what DeRose says in this case.

Here are the beginnings at least of what seems to be a more promising analysis of the loose use implicature: In order to be close enough to knowing that p , one must satisfy all the conditions for knowledge (whatever they are) except for the justification component. As regards this component, one must only be close enough to satisfying it.¹⁰ Assuming that knowledge is justified, true belief and assuming a relevant alternatives account of justification, we could, for example, provide the following definition: S is close enough (for the purposes of the low standards case) to knowing that p iff p , S believes that p and S can rule out all *likely* alternatives to p . Such a view would not be subject to the above criticisms: First, the knowledge ascription in the low standards case now has a true implicature, namely, that DeRose truly believes that the bank will be open and can rule out all likely alternatives to this claim (these alternatives presumably don't include the possibility that the bank changes its hours). Second, this implicature turns out false once we assume that the bank will not be open, for the present account has it that one can be close enough to knowing that p only if p . Of course, the account would have to be refined along many dimensions. The problems arising in this context, however, don't seem to go beyond the problems arising for anybody trying to define the notion of knowledge.

Given the above account of the loose use implicature, *literally* knowing something must amount to being able to rule out more than only the likely alternatives. After all, it should be easier to come close enough to knowledge than to literally know something. Just how many more alternatives do we have to rule out? SPI is compatible with a whole range of answers to this question. As far as the position itself goes, the only constraint is that, literally speaking, DeRose fails to know that the bank will be open in the bank cases. Thus, the additional alternatives will have to include some alternative that DeRose cannot rule out (for example, the alternative that the bank has changed its hours). The precise details are irrelevant for our present concerns, but I will suggest below (section 3.3) that, overall, it might be wise for the proponent

¹⁰ See (Davis 2007: 420) for considerations that go into a similar direction.

of SPI to construe knowledge as a demanding, but not excessively demanding epistemic relation.

Coming back to the alleged implicature in the low standards cases, how is this implicature supposed to be calculated? To fully respond to this worry, we would need a clear conception of what calculability amounts to. Besides, we would have to ask whether implicatures, in general, have to be calculable.¹¹ These issues cannot be addressed here. For now, it should suffice to show that there are various potential triggers for the calculation of the relevant implicature and that the close-enough implicature posited by the proponent of SPI is at least a plausible end point of the triggered calculation. The calculation could be triggered in at least two ways (I will not settle on a specific proposal). First, it could be triggered by the literal falsity of what is said in the low standards case. For DeRose and his wife both know DeRose's evidence, so they should know that it doesn't suffice for knowledge (assuming the SPI account of knowledge). Second, the calculation could be triggered by the overinformativeness of what DeRose says. In the low standards case, it doesn't matter whether DeRose knows or is only close enough to knowing that the bank will be open (because no error-possibilities have been mentioned and not much is at stake). Thus, specifying that DeRose knows the bank will be open would be giving more information than required. These considerations indicate already why DeRose's wife could plausibly settle with the implicature that DeRose is close enough to knowing the bank will be open in the low standards case: This seems precisely what she needs to know in the context described. Why don't we get a similar implicature in the high standards case? What is literally said in this case (that DeRose doesn't know the bank will be open) is true. Besides, being close enough to knowledge wouldn't suffice for the purposes of the high standards case because, in this case, much is at stake and an error-possibility has been mentioned. Thus, there seems to be no trigger for the search for an implicature in the high standards case.¹²

Summing up, SPI amounts to the following claims: (a) Epistemic invariantism is correct. (b) DeRose doesn't know that the bank will be open in the bank cases. As a consequence of (a) and (b), DeRose's knowledge ascription in the low standards case is literally false and his knowledge denial in the high standards cases is literally true. (c) DeRose's knowledge denial in the high standards case seems appropriate (or even true) simply because it is literally true. (d) DeRose's knowledge ascription in the low standards case seems appropriate (or even true) because it has a true substitutional implicature. (e) The content of this implicature is that

¹¹ See my (2015) for some discussion of the former issue.

¹² See (Davis 2007: 413) for related considerations. A common worry for pragmatic accounts of the bank case intuitions is that the relevant implicatures aren't properly calculable. The above considerations are too vague to fully address this worry. Even so, they should suffice to shift the burden of proof to those who think that SPI falls prey to some such worry.

DeRose is close enough (for the purposes of the low standards case) to knowing that the bank will be open. We now have a sufficiently clear conception of the version of SPI that interests us here to consider some objections to this view. Let's start with objections I don't consider decisive.

3. Objections and replies

3.1. Ruling out loose use

Here is the first objection. Assume that DeRose in the low standards says, "I know that the bank will be open *strictly speaking*." His claim will still seem true, but SPI no longer applies because DeRose now makes clear that he isn't speaking loosely. Thus, SPI fails to explain all relevant data.¹³

Compelling as it may seem, this objection isn't convincing because it rests on shaky intuitions. Conee (2005a: 52), for example, explicitly says that, when we ask whether the proposition that the bank will be open in the low standards case "is *really* known, or *truly* known, or *really and truly* known, fluent speakers have a strong inclination to doubt or deny this."¹⁴ In the same way, one can just deny that, intuitively, DeRose would be speaking truly if he were to say, "I know that the bank will be open *strictly speaking*."

3.2. Bank cases in thought

Baumann (2011) puts forward a very general objection to all possible views that seek to explain the bank case intuitions in terms of implicatures. In a nutshell, the objection is that, while implicatures may be used to explain our intuitions about the bank cases as outlined above, they cannot be used to explain corresponding intuitions about bank cases where, instead of asserting it, DeRose only thinks to himself that he knows (or doesn't know) the bank will be open. As applied to SPI, this objection could be spelled out as follows: Suppose that, in the low standards case, DeRose doesn't assert, "I know the bank will be open," but silently forms the belief that he knows that the bank will be open.¹⁵ As before, we would intuitively say that DeRose's belief is appropriate (or even true). This intuition, however, cannot be explained in terms of implicatures, for beliefs just aren't the right sort of thing to have implicatures. So, SPI fails to explain our intuitions about the belief version of the low standards case.

To respond, I am happy to grant that beliefs don't have implicatures. I am also happy to grant that we would intuitively evaluate DeRose's belief in just the same way in which we evaluate

¹³ See e.g. (Hawthorne 2004: 120) and (Blome-Tillmann 2013: 4302) for related worries.

¹⁴ See (Unger 1971: 214–216) and (Davis 2007: 430) for related remarks.

¹⁵ See (Baumann 2011: 160) for the resulting case.

his assertion. Even so, I don't think that SPI cannot explain these intuitions. The basic reason is that, even if beliefs don't have implicatures, belief *ascriptions* certainly do. Let me elaborate.

Suppose you read the following sentences in a suitable story: "The church bell rang five times. 'It's five,' Jill thought. 'I should be going home.'" You would presumably be happy to grant that Jill's belief is true. However, assume the implicature account of loose use according to which loose use does not figure at the level of what is literally said. On this view, what is literally said in the story is that Jill believed that it was *exactly* five o'clock. And this belief is most likely false. After all, Jill counted the rings of the church bell before she came to hold that belief; and after the last ring, it was presumably already after five. (Besides, church bells presumably aren't that precise anyway.) Does this show that the implicature account of loose use fails? I don't think so. A proponent of this account will argue as follows: The belief ascription "'It's five', Jill thought" implicates that Jill thought that it is *close enough* to five o'clock. For, it would be absurd in the story to ascribe to Jill the belief that it is *exactly* five o'clock. Now the former belief is most likely true. So we judge Jill's belief to be true when we read the above story.

Proponents of SPI can argue that the same happens when we consider a belief version of the low standards bank case: The description of such a case will contain a belief ascription like this: "DeRose believes that he knows the bank will be open." This belief ascription will implicate that DeRose believes that he is *close enough* to knowing the bank will be open. For, it would be absurd to impute to DeRose the false belief that he knows the bank will be open (given that he knows what evidence he has and given that this evidence doesn't suffice for knowledge). The former belief is correct. Thus, we judge DeRose's belief to be correct.

3.3. Sentence and speaker meaning

The following seems to be the most prominent objection to SPI:¹⁶ SPI has it that DeRose's knowledge ascription in the low standards case is literally false but carries a true substitutional implicature. Thus, the view presumably entails that DeRose in the low standards case doesn't use "knows" to convey what it literally expresses. If that is so, people will not use "knows" to convey what it literally expresses in most ordinary conversations, for most ordinary conversations in which we ascribe knowledge are very much like the conversation in the low standards case. In particular, the subject's epistemic position is roughly as good as DeRose's, there isn't much at stake and no error-possibilities are mentioned. However,

¹⁶ See (Hawthorne 2004: 121–123), (DeRose 2009: 125), (Fantl and McGrath 2009: 48) and (Blome-Tillmann 2013: 4304 f.).

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. (Blome-Tillmann 2013: 4304)

In other words, “eminently plausible” views on meaning determination entail that it is impossible for an expression to literally express something other than what it is used to convey in most ordinary conversations. Thus, DeRose in the low standards case must use “knows” to convey what it literally expresses, and SPI is mistaken.¹⁷

Response 1

This worry doesn’t seem decisive because the view on meaning determination described by Blome-Tillmann seems implausible, or questionable at least. For example, the view is incompatible with the common and plausible assumption that “three o’clock” is standardly used to refer to a range of points in time rather than the precise point in time it literally denotes. A less controversial view in the vicinity would be the following: What a term literally expresses is determined by what the term is standardly used to convey *when used literally*. This latter view, however, is unproblematic for SPI. It entails that there is “no room for a difference” between what a term literally expresses and what it is standardly used to convey *when it is used literally*. Accordingly, the proponent of SPI holds that, on the (potentially rare) occasions where “knows” is used literally (for example, when we speak of a priori truths or immediate perceptions), it does convey what it literally expresses; namely, a quite demanding epistemic relation.

One may worry that we must accept Blome-Tillmann’s meaning determination principle because the above alternative position is less controversial only at the cost of being pointless: To determine the literal meaning of a given expression, we can no longer just look at the use of this expression. We have to sort out non-literal uses first. But to do so, we have to know already what the literal meaning of the relevant expression is. To respond, note that, of course, the proposed view must be supplemented with independent tests for non-literality (which would presumably split up into tests for irony, metaphor, etc.). To take just one example, a test for irony could be that ironical uses are marked by a specific tone of voice. Tests along these lines should guide us in sorting out non-literal usages, not our previous knowledge of the literal meaning of the expression involved. (Devising the relevant tests admittedly is a massive project on its own.)¹⁸

¹⁷ Blome-Tillmann endorses the above meaning determination principle only for expressions that are not “composed out of simpler expressions.” I ignore this restriction because we will be dealing only with non-composed expressions anyway.

¹⁸ I will later argue that SPI must be rejected because the knowledge claims in the bank cases fail relevant tests for non-literality. Note, however, that this objection is entirely independent from the general considerations on

Response 2

Let me also discuss a second response strategy to the above objection. It is based on the idea that DeRose might not be as representative a knower as this objection suggests. In detail, it goes as follows: Consider the claim that a proponent of SPI is committed to the view that ordinary knowledge claims generally don't convey what they literally express. This commitment is supposed to follow from the thesis that DeRose doesn't know that the bank will be open in the bank cases. Note though that it only follows from this thesis if the epistemic position of ordinary "knowers" generally equals DeRose's in the bank cases (or is even worse). This is far from trivial, however. It may just as well be that, in general, we ascribe knowledge only to subjects whose epistemic position is better than DeRose's. This just seems to be a hitherto unsettled empirical issue.¹⁹ Thus, a proponent of SPI can respond to the present argument by just denying that ordinary knowledge claims generally don't convey what they literally express.^{20,21}

Here is a possible rejoinder to this response. Note that, according to the response, we often use "knows" literally. But if that is so, Cohen (2005: 58) objects, then the proponent of SPI can no longer explain why "competent speakers, under skeptical pressure, tend to deny that we know even the most conspicuous facts of perception, the clearest memories, etc." This objection can be explicated as follows: If, as the present response has it, we often use "knows" literally, then we most certainly do so in the cases Cohen describes (that involve "the most conspicuous facts of perception, the clearest memories, etc."). And, if we use "knows" literally in these cases, SPI has no explanation for why we still deny the relevant knowledge claims in the face of skeptical considerations. So, if we often use "knows" literally, Cohen's cases cannot be explained in terms of SPI.

This rejoinder is not convincing because it relies on what seems to be a dubious datum; namely, that, in the face of skeptical pressure, competent speakers deny knowledge even of "the

meaning determination above: If knowledge claims involve loose use, they should pass the tests for loose use whatever we think about the determinants of the meaning of an expression.

¹⁹ Recent empirical studies on the bank cases corroborate this assessment. It is striking that when people are asked about the truth-value of the knowledge claims in the bank cases, their responses generally range somewhere in the middle between "true" and "false." See e.g. (Hansen and Chemla 2013). These results strongly suggest that DeRose is at least not a paradigm case of an ordinary "knower."

²⁰ Conee (2005a: 52) and Davis (2007; 2010) suggest responses along these lines.

²¹ This move should also suffice to alleviate Hawthorne's worry that, if SPI is true, "no one—not even the philosophically unsophisticated—ever believes that he knows such ordinary proposition as that Manchester United beat Coventry City 2–1, or that the plane from Detroit is late, or that the towels are in the dryer." (2004: 119) Assuming that, according to SPI, knowledge claims like the ones mentioned by Hawthorne normally do convey what they literally express, SPI is perfectly compatible with the view that ordinary speakers have the pertaining beliefs. See (Douven 2007: 344 f.) for further discussion of this issue.

most conspicuous facts of perception, the clearest memories, etc.” Many people, in particular proponents and sympathizers of SPI, would reject this datum. Davis, for example, claims:

Careful reflection does not ordinarily prompt a retraction or correction of the claim to know one has a hand, the way it prompts Hannah [in our case DeRose] to retract her claim to know the bank will be open on Saturday. (Davis 2007: 436)

Relatedly, Fantl and McGrath write:

Try to get someone to concede he doesn’t know his spouse will be home by 6 p.m., and you will get your concession without too much work. Try to get someone to concede that he doesn’t know that George W. Bush wasn’t the first president—and good luck to you! (Fantl and McGrath 2009: 193f)²²

All three authors seem to agree that there is a whole range of claims that we wouldn’t retract under skeptical pressure. Thus, we can dismiss Cohen’s rejoinder because it relies on what seems to be a questionable datum.²³

Admittedly, at least some people (some philosophers, for example) become convinced by wide-ranging skeptical arguments. As a result, they think that they know hardly anything at all. This datum cannot be explained in terms of SPI once we accept the above response. However, I don’t think that this datum and the more mundane bank case reactions need to be explained in a uniform fashion.²⁴ Here is one candidate account of why skeptical arguments can lead some thinkers to hold that they know nothing: Good skeptical arguments are very difficult to counter. Even though unsound, these arguments involve argumentative fallacies that are very hard to spot.²⁵ It may be debatable whether such an account withstands scrutiny. I don’t see any reason, however, to think that the account would be incompatible with SPI. So it is an open question at best whether the above reactions to skeptical arguments provide troubling data for the proponent of SPI.

²² See also (Conee 2005a: 53) and (Douven 2007: Sec. 2).

²³ On this basis, we can also respond to the related worry that, unless we go for full-blown skepticism, there is no non-arbitrary way to set the boundary between knowledge and non-knowledge. See e.g. (DeRose 2012: 712 f.). We should set the boundary between knowledge and non-knowledge such that the knowledge ascriptions that we would retract under slight skeptical pressure do not count as true and those that we wouldn’t so retract do count as true. (BonJour (2010: 58) worries that, if we set the boundary between knowledge and non-knowledge in this not fully skeptical way, we cannot explain why knowledge is “a supremely valuable and desirable cognitive state, one whose possession marks the difference between full cognitive success and at least some degree of cognitive failure”. This may be so. But I, at least, fail to see why knowledge should be such a state to begin with. Knowing is to be in a very good epistemic position. But this does not mean that one’s epistemic position could not be better.)

²⁴ See (Conee 2005b: 66) and (Davis 2007: 436, 2015) for the idea of separating the account of our reactions to skeptical arguments and the account of the bank case intuitions.

²⁵ See (Sosa 1999: 148) for a fallacy that might be involved in skeptical arguments. He argues that knowledge requires “safe” not “sensitive” beliefs. On that basis, he goes on to suggest that skeptical arguments may sound compelling because “[s]afety and sensitivity, being mutual contrapositives, are easily confused.”

To be clear, if the proponent of SPI accepts the present response, her view no longer entails skepticism (on many common understandings of this notion). This is unproblematic however. SPI is *S(keptical)PI* because it entails that DeRose fails to know that the bank will be open in the bank cases. Stronger forms of skepticism were never supposed to be part of the view.²⁶

So far then, SPI looks like an attractive account of our intuitions about the bank cases.²⁷ Let us now turn to what I take to be a serious objection to SPI: the awareness objection. The idea behind this objection is not new.²⁸ I will improve on extant formulations, however, in various ways. First, I will restate the principles on which the objection relies in order to avoid a range of possible worries (see footnote 30 and pp. 15 f.). Second, I will underwrite the intuitions on which it relies by appealing to the idea that high standards contexts are “enlightened” contexts (see pp. 16 f.).

4. The awareness objection

The awareness objection is based on the observation that, in the case of substitutional implicatures (i.e. the implicatures relevant for SPI, that substitute what is literally said), people are aware of the discrepancy between what they say and what they implicate. Let’s refer to this idea as the *awareness principle*. The awareness principle can be cashed out as follows: Suppose someone makes an utterance that carries a substitutional implicature. Suppose further that an interlocutor objects to the literal content of this utterance either because she fails to realize that the utterance has a substitutional implicature or because she intentionally takes the utterance too literally—is being a “stickler for correct speech” to use a phrase from Grice (1989: 45). Now, the awareness principle says the following: A speaker confronted with such an objection

²⁶ Proponents of the present meaning determination objection may never have intended to argue against versions of SPI that don’t entail skepticism by definition. If that is so, the present response only shows that their objection doesn’t carry over to the form of SPI we are concerned with here (but see the first response above).

²⁷ I should briefly mention some potential worries for SPI that, I think, have been answered elsewhere already: One may want to argue against SPI on the grounds that the alleged implicature should be felicitously cancelable, but isn’t. Worries along these lines seem properly addressed e.g. in (Davis 2007: 411, 431): The implicature cannot be (explicitly) canceled because it is entailed by what is literally said (as is the implicature that it is close enough to three o’clock of an utterance of “It is three o’clock”). One may want to argue against SPI on the grounds that, by appealing to loose use, the proponent of SPI can no longer explain disagreement and retraction data. Worries along these lines seem properly addressed e.g. in (Davis 2007: 406 f.): It is perfectly appropriate to retract claims once it is pointed out that they are strictly speaking false. (See below for some further remarks on this issue.) Finally, one may want to argue against SPI on the grounds that loose use cannot explain why it would seem *inappropriate* for DeRose to *deny* knowledge in the low standards case. Worries along these lines seem properly addressed e.g. in (Davis 2007: 408 f.): Negative statements like “It is not three o’clock” can be used loosely to convey that it is not close enough to three o’clock. The same goes for “I don’t know the bank will be open.” This sentence may also convey that DeRose is not close enough to knowing that the bank will be open, which is false in the low standards case.

²⁸ Versions of the worry have been presented e.g. in (Hawthorne 2004: 104 f.), (MacFarlane 2005: 206 f.) and (Blome-Tillmann 2013: Sec. 3).

will straightforwardly be able to make sense of the objection (either in terms of the idea that the objector failed to see the implicature or else is being nitpicky) and respond accordingly.

Let's consider some examples to substantiate this principle. Suppose the following dialogue takes places:

A: The plane was a mile long.

B: That's absurd! No plane is a mile long!

In a suitable context, A's first utterance generates a substitutional implicature to the effect that the plane was huge. B's response, however, is an objection only to the falsity of the literal content of A's claim. Correspondingly, A, *qua* competent speaker, will straightforwardly realize that B is missing the point (intentionally or unintentionally). Thus, A can plausibly respond by saying,

A: Oh c'mon, you know what I meant!

This is what the awareness principle would predict. Analogous considerations hold for the following dialogues involving loose use:

A: It's three o'clock.

B: It's one minute past!

A: Oh, c'mon. You know what I meant.

A: France is hexagonal.

B: But what about this bump here and this bump there?

A: Oh, c'mon. You know what I meant!²⁹

The awareness principle thus seems to make the right predictions.³⁰

²⁹ The same goes for (substitutional) implicatures: According to Bach (1994), a typical utterance of "I have nothing to wear" expresses the falsity that the speaker has nothing to wear whatsoever but truly implicates that she has nothing *appropriate* to wear *for a given occasion*. Correspondingly, the following dialogue seems fine:

A: I have nothing to wear.

B: Sure you do. There are socks in the drawer, shirts on the shelf ...

A: Oh, c'mon. You know what I meant!

Similar dialogues could, I think, be devised for all relevant cases of implicatures. Thus, the present objection applies to SPI even if the proponent of SPI appeals to implicatures rather than implicatures.

³⁰ My presentation of the awareness principle closely follows the presentation in (Blome-Tillmann 2013: Sec. 3). Major points of divergence will be spelled out in the main text. Here I will briefly outline three minor points of divergences that should be mentioned nevertheless. First, Blome-Tillmann does not show that the awareness principle applies to *implicatures*. This makes the range of targeted positions unnecessarily narrow. Second, he does not identify the category of a *substitutional* implicature (or implicature) as the category of implicatures to which the awareness principle applies. Rather, he seems to think that the relevant class of implicatures involves violations of Quality₁ (the conversational maxim enjoining you to be sincere). See his principle (DP) and FN 39 on p. 4311.

Given this principle, the problem for the proponent of SPI is the following: According to SPI, DeRose's knowledge ascription in the low standards case carries a substitutional implicature. It should thus be natural for DeRose to respond to challenges to the literal content of what he says in the way predicted by the awareness principle. In particular, it should be natural for him to respond in this way to objections to the effect that he fails to rule out certain error-possibilities that are relevant for knowledge strictly speaking but not for being close enough to knowledge. Thus, the following dialogue should be fine in the low standards case:

(D1)

DeRose: I know that the bank will be open.

His wife: I doubt that. The bank could have changed its hours.

#DeRose: Oh, c'mon. You know what I meant!

However, many would say (and I would agree) that DeRose's second claim seems not particularly natural. As Hawthorne (2004: 105) points out, "our standard techniques for dealing with epistemic challenges that raise relatively farfetched possibilities are concession and, more rarely, sticking to one's guns." So, DeRose could plausibly *concede*, "You're right. I didn't

This is wrong, I think, because not all substitutional implicatures involve Quality₁ violations, and even so, they all allow for the construction of dialogues of the above sort. We only have to gear *B*'s objection to whatever maxim is violated at the level of what is said. For example, suppose I say, "Jane will come or she won't" in response to the question of what to do about Jane's absence. I do not violate Quality₁ because I do believe what I literally say. Still, I substitutionally implicate, say, that nothing can be done about Jane's absence because what I literally say is entirely uninformative (it's a mere logical truth that we all believe already anyway). Correspondingly, a stickler may object, "Well, these are all the options, aren't they? I wanted to know what to do about Jane's absence." And, again, it would be fine to respond, "Oh, c'mon. You know what I meant! (There's nothing we *can* do.)" Third, I have argued that *B*'s objections can be understood either in terms of the idea that *B* fails to see the implicature or in terms of the idea that she is being a stickler. Blome-Tillmann (2013: 4311 n.) envisages only the first kind of understanding. The stickler reading, however, often is the only plausible reading available (particularly when we consider implicatures).

think of that” or *stick to his guns* by saying “(Oh, c’mon.) How likely is that?” The above, *clarifying* response, however, seems unnatural. This is the awareness objection to SPI.^{31,32}

The present formulation of the awareness objection diverges from extant formulations in a subtle but important respect. In particular, the above dialogues have the first speaker respond to the second speaker’s objection by saying “You know what I meant.” In standard formulations, the speaker responds by saying “I was speaking loosely” or even by spelling out the supposedly implicated claim, for example, by saying something like “All I meant was that it is *approximately* three o’clock.” Both these alternatives are problematic. The second alternative assumes that speakers are always able to straightforwardly articulate what they implicate. This assumption, however, seems unwarranted. It seems that we do sometimes respond by saying “Oh, c’mon. You know what I meant” precisely when we did mean something (and take our audience to see what we meant) but aren’t in a position to straightforwardly come up with an articulation of what we meant. The first alternative is problematic because it is unclear whether a response along these lines would be unacceptable for DeRose in (D1). Blome-Tillmann, for example, suggests that “it might seem natural for [DeRose] to reply that [he] was speaking loosely.” (2013: 4303 n.) Similarly, consider the following dialogue:

DeRose: I know that the bank will be open.

³¹ One may worry that DeRose’s wife *is* addressing what DeRose means to convey in the dialogue above because she implicates that the possibility that the bank changes its hours not only exists but is likely enough to be relevant even for being close enough to knowledge. This would explain why the “You know what I meant” response is unacceptable. Even though I find it very plausible that DeRose’s wife does implicate something along these lines, I think the proponent of SPI cannot plausibly appeal to this implicature. First, this would threaten to make her view superfluous: Consider DeRose’s wife in the high standards case. As in (D1), she mentions the error-possibility that the bank has changed its hours. Given that in (D1), she implicates that this possibility is likely enough to be relevant for the epistemic position at issue in (D1) (that is, for being close enough to knowledge), it is very hard to deny that there should be a similar implicature in the high standards case. In particular, DeRose’s wife in this case should implicate that the mentioned error-possibility is likely enough to be relevant for the epistemic position at issue in the high standards case (that is, I take it, for knowledge strictly speaking). Once we grant this implicature, however, we can argue that DeRose’s wife changes DeRose’s epistemic position: In the high standards case, but not in the low standards case, he gains testimonial evidence that the bank is relevantly likely to change its hours. Correspondingly, we can say that our intuitions about the bank cases vary simply because DeRose’s epistemic position varies; SPI is no longer required. (See my (in press) for further elaboration on this kind of proposal and its application to so-called “third-person” bank cases.) Second, it doesn’t seem that the “You know what I meant” response would become any more appropriate if we let DeRose’s wife cancel the alleged implicature, for example, if we let her respond with “I doubt that. Not that this is likely to any relevant degree, but the bank could have changed its hours.” Hence, the awareness objection could easily be reinstated. Thanks to an anonymous referee for pushing the worry discussed in this footnote.

³² To be clear, I am happy to grant that, if the error-possibility DeRose’s wife mentions becomes ever more far-fetched, the “You know what I meant” response may become plausible at some point. For example, the response may seem fine when she inveighs, “I doubt that, we might be brains in a vat after all.” But this observation is of no avail to the proponent of SPI. So long as SPI is supposed to provide an answer to the bank case puzzle, it must be capable of dealing with less far-fetched alternatives such as the one in (D1). After all, these alternatives figure in the bank cases. (And they do so for a good reason: It is very much unclear whether more far-fetched alternatives would trigger the intuitions constituting the bank case puzzle. See e.g. (Gerken 2012: 141 f.))

His wife: I doubt that. The bank could have changed its hours.

DeRose: You're right. I don't know it.

It has frequently been pointed out that DeRose could now go on to say something like "I was speaking casually," particularly if his wife were to insist, "But you just said that you do know that!"³³

Neither of these worries affects the awareness objection once it is spelled out in terms of the "You know what I meant" response. First, on this construal, we no longer have to assume that speakers are always able to articulate what they meant. We only have to assume that they realize that they did not mean what they literally said (when the two come apart). Second, even if DeRose in (D1) could respond by saying, "I was speaking loosely/casually," it does not follow that the "You know what I meant" response is acceptable as well. To see this, consider the following dialogue:

DeRose: The bank will be open.

His wife: I doubt that. The bank could have changed its hours.

DeRose: I guess you're right. It might be closed.

Just as before, DeRose could now go on to say, "I was speaking casually/loosely," particularly if his wife were to insist, "But you just said that it will be open!" Obviously, though, DeRose's original assertion wasn't a case of loose use (in the relevant sense). More plausibly, it was just an ill-considered judgment.³⁴ And when DeRose says that he was speaking loosely/casually, he plausibly just wants to admit that he wasn't careful enough when he made his original statement (not that he didn't mean what he said). Correspondingly, it would *not* seem fine for him to respond, "Oh, c'mon. You know what I meant!"

One may worry that even if the acceptability of the "You know what I meant" response in (D1) does not follow from the acceptability of the "I was speaking loosely/casually" response, people may still find the "You know what I meant" response acceptable on independent intuitive grounds. This worry, however, is unfounded as the subsequent considerations show. Most authors should agree that high standards contexts feel "enlightened" in some sense. As Hawthorne puts it,

³³ See e.g. (DeRose 2009: 171), (Davis 2010: 1155) and (Dimmock and Huvenes 2014: 6 n.).

³⁴ Davis (2007: 410 f.) similarly distinguishes loose use from what he calls "sloppy" or "careless" or "unenlightened" use.

one thing seems clear and very important: we do have some tendency to suppose that, as more and more possibilities of error become salient to us, we are reaching an ever more enlightened perspective. (Hawthorne 2004: 164)

Cohen concurs:

There is no doubt that when one is in what the contextualist wants to view as a high standards context, one has a feeling of enlightenment regarding the correct application of the predicate in question. We feel as if we are seeing the truth of the matter that has, up until that point, eluded us. (Cohen 2005: 58)

Those who accept this idea should reject DeRose's response in (D1) for the following simple reason: In (D1), DeRose's wife objects to DeRose by mentioning an error-possibility. Presumably, she thereby raises the epistemic standard. Thus, given the assumption that high standards contexts are enlightened contexts, DeRose should feel enlightened in some sense after hearing his wife's objection. However, when DeRose responds to this objection by saying, "Oh c'mon. You know what I meant!" he is precisely denying that there is anything enlightening about this objection. To the contrary, he is suggesting that he had the mentioned error-possibility in mind all along, but just didn't bother to make that explicit. So, if we grant that raising epistemic standards is enlightening, then we should say that this response is inappropriate. Thus, most authors should be willing to accept my verdict about (D1).³⁵

Couldn't one defend SPI along the following lines? Merely mentioning an error-possibility—as DeRose's wife does in (D1)—does not suffice to raise the epistemic standard. At least, there are many potential reasons to doubt this assumption: Maybe one must also raise the stakes. Maybe the error-possibility must be made salient in some stronger sense. Maybe the error-possibility must be taken seriously, etc. Thus, we shouldn't expect DeRose to feel enlightened in (D1) even if we grant that high standards contexts are enlightened contexts. To respond, note that we can simply add all of the just mentioned dimensions to (D1). Once we have done so, it should be unproblematic to hold that DeRose should feel enlightened. As before, however, this verdict is incompatible with SPI, for even if, for example, DeRose takes seriously the possibility that the bank changes its hours, his failure to rule out this possibility concerns only what he literally said, not what he supposedly implicated. Thus, given SPI, the "You know what I meant" response should still be fine, and DeRose should not feel enlightened at all.

³⁵ To be clear, note that the phenomenon of enlightenment in high standards contexts, though closely related, must not be confused with the phenomenon of retraction: that we tend to retract knowledge claims when someone mentions an error-possibility. As indicated already, loose use may explain retraction. Still, when we retract claims because it is pointed out that they are strictly speaking false, we do so with what MacFarlane dubs an "exasperated grumble." (2011: 541) We precisely don't consider ourselves enlightened. So, loose use cannot explain enlightenment even though it can explain retraction.

5. Conclusion

To sum up, SPI is much more promising than it is standardly made out to be. In fact, the literature contains no compelling objections to this view. The awareness objection, however, can be developed into a serious challenge for SPI. So in the absence of a plausible response to this challenge, the view should be rejected after all.

Acknowledgements: I am grateful to Michael BlomeTillmann, Elke Brendel, Aurélien Darbelay, Mikkel Gerken, Beate Krickel, David Lanius, Dan López de Sa, David Löwenstein, Erik Stei, Emanuel Viebahn, Julia Zakkou and the participants of the Fifth Annual Graduate Epistemology Conference (Edinburgh, 2015) for very helpful comments on earlier drafts of the paper.

References

- Bach, K. (1994), 'Conversational implicature', *Mind & Language*, 9/2: 124–162.
- Bach, K. (2010), 'Knowledge In and Out of Context', in J. K. Campbell, M. O'Rourke, and H. Silverstein (eds.), *Knowledge and Skepticism* (Topics in Contemporary Philosophy, Cambridge, MA: MIT Press), 105–136.
- Baumann, P. (2011), 'WAMs: Why Worry?', *Philosophical Papers*, 40/2: 155–177.
- Blome-Tillmann, M. (2013), 'Knowledge and implicatures', *Synthese*, 190/18: 4293–4319.
- BonJour, L. (2010), 'The Myth of Knowledge', *Philosophical Perspectives*, 24/1: 57–83.
- Cohen, S. (2005), 'Contextualism Defended', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Contemporary Debates in Philosophy, 3, Malden, MA: Blackwell Pub.), 56–62.
- Conee, E. (2005a), 'Contextualism Contested', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Contemporary Debates in Philosophy, 3, Malden, MA: Blackwell Pub.), 47–56.
- Conee, E. (2005b), 'Contextualism Contested Some More', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Contemporary Debates in Philosophy, 3, Malden, MA: Blackwell Pub.), 62–66.
- Davis, W. A. (2004), 'Are Knowledge Claims Indexical?', *Erkenntnis*, 61/2/3: 257–281.
- Davis, W. A. (2007), 'Knowledge Claims and Context: Loose Use', *Philosophical Studies*, 132/3: 395–438.
- Davis, W. A. (2010), 'The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1, by Keith DeRose', *Mind*, 119/476: 1152–1157.
- Davis, W. A. (2015), 'Knowledge claims and context: belief', *Philosophical Studies*, 172/2.
- DeRose, K. (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1* (Oxford: Oxford University Press).
- DeRose, K. (2012), 'Replies to Nagel, Ludlow, and Fantl and McGrath', *Philosophy and Phenomenological Research*, 84/3: 703–721.
- Dimmock, P., and Huvenes, T. (2014), 'Knowledge, Conservatism, and Pragmatics', *Synthese*: 1–31.
- Dinges, A. (2015), 'Innocent implicatures', *Journal of Pragmatics*, 87: 54–63.
- Dinges, A. (forthcoming), 'Epistemic invariantism and contextualist intuitions', *Episteme: A Journal of Individual and Social Epistemology*.
- Douven, I. (2007), 'A Pragmatic Dissolution of Harman's Paradox', *Philosophy and Phenomenological Research*, 74/2: 326–345.

- Fantl, J., and McGrath, M. (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).
- Gerken, M. (2012), 'On the Cognitive Basis of Knowledge Ascriptions', in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions* (Oxford: Oxford University Press), 140–170.
- Grice, P. (1989), *Studies in the Way of Words* (Cambridge, Mass: Harvard University Press).
- Hansen, N., and Chemla, E. (2013), 'Experimenting on Contextualism', *Mind & Language*, 28/3: 286–321.
- Hawthorne, J. (2004), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- MacFarlane, J. (2005), 'The Assessment Sensitivity of Knowledge Attributions', in T. S. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology. Volume 1* (Oxford Studies in Epistemology, 1, Oxford: Oxford University Press), 197–234.
- MacFarlane, J. (2011), 'Relativism and Knowledge Attributions', in S. Bernecker and D. Pritchard (eds.), *The Routledge Companion to Epistemology* (New York: Routledge), 536–544.
- Meibauer, J. (2009), 'Implicature', in J. Mey (ed.), *Concise Encyclopedia of Pragmatics* (Amsterdam: Elsevier), 365–378.
- Schaffer, J. (2004), 'Skepticism, Contextualism, and Discrimination', *Philosophy and Phenomenological Research*, 69/1: 138–155.
- Sosa, E. (1999), 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 33/13: 141–153.
- Unger, P. (1971), 'A Defense of Skepticism', *The Philosophical Review*, 80/2: 198–219.