

Epistemic Invariantism and Contextualist Intuitions

Dissertation
zur Erlangung des akademischen Grades

Doctor philosophiae
(Dr. phil.)

eingereicht
an der Philosophischen Fakultät I
der Humboldt-Universität zu Berlin

von Alexander Dinges, M.A.
geboren am 06.04.1984 in Bad Soden a. Ts.

Promotion erfolgte am 3. Juni 2015.

Der Präsident der Humboldt-Universität zu Berlin
Prof. Dr. Jan-Hendrik Olbertz

Der Dekan der Philosophischen Fakultät I
Prof. Dr. Michael Seadle

Gutachter

1. Prof. Dr. Geert Keil
2. Prof. Dan Lopéz de Sa
3. Prof. Jessica Brown

Epistemic Invariantism and Contextualist Intuitions

Alexander Dinges

Contents

1. Introduction	1
2. Relativism.....	21
3. Moderate Pragmatic Invariantism	47
4. Skeptical Pragmatic Invariantism.....	95
5. Projective Invariantism.....	112
Appendix: The Factivity Problem	129
References	147

Detailed Contents

1.	Introduction	1
1.1.	The Challenge to Invariantism	2
1.2.	Projective Invariantism Outlined.....	8
1.3.	Why not Contextualism?	9
1.4.	What was Invariantism Again?	11
1.5.	Chapter Summary	15
1.6.	Scope Limitations	17
2.	Relativism.....	21
2.1.	Relativism.....	22
2.2.	The Indirect Testimony Datum (ITD)	22
2.3.	An Argument Against Relativism Based on the (ITD)	26
2.4.	An Alternative to the (ITD).....	31
2.5.	Replies and Rejoinders	33
2.6.	Concluding Remarks	41
	Appendix 1: Scope Limitations	42
	Appendix 2: The Case Design	43
3.	Moderate Pragmatic Invariantism	47
3.1.	Conversational Implicatures: Additive and Substitutional.....	50
3.2.	Additive Implicatures and Seeming Truth-Values	52
3.3.	Substitutional Implicatures and Seeming Truth-Values.....	61
3.4.	Calculating Implicatures.....	69
3.4.1.	Observation and Exploitation Implicatures	69
3.4.2.	Calculating Substitutional Implicatures in the Bank Cases.....	87
3.5.	Conclusion.....	93

4. Skeptical Pragmatic Invariantism.....	95
4.1. What is Implicated?.....	96
4.2. Objections and Replies	99
4.3. Awareness.....	105
4.4. Conclusion.....	111
5. Projective Invariantism.....	112
5.1. The Contextualist Data	113
5.2. Objections and Replies	120
5.3. Conclusion.....	128
Appendix: The Factivity Problem	129
References	147

1. Introduction

It seems like an intuitively very plausible position that, if someone utters the sentence “Hannah knows that there is no life on Mars,” she thereby says that Hannah knows that there is no life on Mars. More generally, it seems that what is said, the proposition expressed, by knowledge ascribing sentences of the form “ S knows that p ” is the same in all contexts in which these sentences can be used; in all these contexts, “ S knows that p ” plausibly just says that S knows that p . I will call this position, or roughly this position, (*epistemic*) *invariantism*. Though intuitively very plausible, invariantism is far from universally accepted. Many philosophers reject invariantism and thus subscribe to (*epistemic*) *contextualism*, that is, roughly the view that knowledge ascribing sentences of the form “ S knows that p ” express different propositions in different contexts of use. What reasons are there to accept contextualism?

Historically speaking, the interest in contextualist positions has been sparked by the idea that they at least promise to provide a plausible response to longstanding and fundamental epistemological problems, most importantly, the problem of skepticism. Contextualism can be employed to ward off skeptical challenges to our everyday knowledge attributions by confining the success of skepticism to contexts where the skeptic manages to impose epistemic standards that are far more demanding than the standards governing our ordinary discourse. Though issues of this sort still play an important role in the contemporary discussion of contextualism, it has also become apparent that, by weighing invariantism and contextualism against each other, there is a lot to be learned about very general and equally fundamental issues in the philosophy of language. For example, we have to ask what linguistic data can be adduced to substantiate a semantic position such as contextualism, we have to ask when such data can be explained away in terms of more general pragmatic mechanisms, we have to ask how exactly to formulate a position such as contextualism and how to distinguish it from other positions in the vicinity (such as relativism), etc. It is these latter questions that will take center stage in the present dissertation.

The discussion will revolve around what seem to be the most important linguistic data put forward in support of contextualism, namely, our intuitions about scenarios such as

the so-called *bank cases*. When we consider the bank cases, it seems that the truth-value of sentences involving the term “knows” depends on the context in which these sentences are used, more specifically, on how much is at stake and which error-possibilities are salient for the knowledge ascribing subject. These data pose a challenge to invariantism because it is unclear how this seeming variability should be explained if not by assuming that the knowledge sentences in question express different propositions depending on the context in which they are used.

The goal of my dissertation is to show that invariantists can meet this challenge. In particular, I will argue that the invariantist should respond to this challenge by appealing to a novel version of invariantism that I will call *projective invariantism*. Projective invariantism says that we evaluate the knowledge claims in the bank cases differently not because they express different propositions in the different scenarios but simply because the epistemic position of the subject in question differs between these scenarios. More precisely, projective invariantism has it that readers of the bank cases will *project* certain assumptions into the scenarios such that, given these assumptions, the subject’s epistemic position differs in the relevant way.

As indicated already, the linguistic data from the bank cases are not the only possible objection to invariantism (skeptical arguments, for example, are another concern). So, even if my project succeeds, we cannot safely conclude that invariantism is correct. However, according to DeRose—one of the major proponents of contextualism—the data from the bank cases form the “best grounds” (2009: 47) for rejecting invariantism. So, if my project succeeds, we can at least be optimistic that a full defense of invariantism is forthcoming.

In this introductory chapter, I will elaborate on the challenge from the bank cases and say in slightly more detail what the projective invariantist solution to this challenge looks like. I will conclude by outlining the chapters to come and delineating the scope of this dissertation.

1.1. The Challenge to Invariantism

Consider the following scenarios, the so-called *bank cases*.

LOW: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they be deposited right away. Hannah suggests that they drive straight home and return to deposit their paychecks on Saturday morning. She remembers driving by the bank last Saturday and seeing that it was open until noon. Sarah says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays. Do you know the bank will be open tomorrow?” Hannah replies,

- ⊕ “I know the bank will be open tomorrow.”
- ⊖ “Well, no, I don’t know the bank will be open tomorrow.”

It turns out that the bank is open on Saturday.

HIGH: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Hannah and Sarah have just written a very large check, and if the money from their pay is not deposited by Monday, it will bounce, leaving them in a very bad situation with their creditors. And, of course, the bank is not open on Sunday. Hannah suggests that they drive straight home and return to deposit their paychecks on Saturday morning. She remembers driving by the bank last Saturday and seeing that it was open until noon. Sarah reminds Hannah of how important it is to deposit the check before Monday and says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays. Besides, banks can always change their hours. Do you know the bank will be open tomorrow?” Hannah replies,

- ⊕ “I know the bank will be open tomorrow.”
- ⊖ “Well, no, I don’t know the bank will be open tomorrow.”

It turns out that the bank is open on Saturday.¹

¹ This particular formulation of the bank cases diverges from their original formulation in (DeRose 1992) in that the bank cases above feature both a knowledge ascription and a knowledge denial in both *LOW* and

LOW and HIGH come in two versions, on the one hand, featuring a knowledge ascription, on the other, a knowledge denial. I will refer to the two versions of LOW as *the low standards cases* and to the two versions of HIGH as *the high standards cases*. What should we say about the truth-values of the knowledge ascriptions and denials in the low and the high standards cases? It may seem that the truth-value of these claims differs depending on whether they occur in the low or the high standards cases. In particular, the following data plausibly obtain.

Hannah's knowledge ascription seems true in the low standards cases but false in the high standards cases;² and, correspondingly,

Hannah's knowledge denial seems true in the high standards cases but false in the low standards cases.

Let's refer to these data as *the Contextualist Data*. Some (or even all) of these data may be doubted. Since the Contextualist Data have been endorsed by many philosophers, however, I will just take them for granted for the purposes of this dissertation.³

The Contextualist Data present a challenge for the invariantist for the following reason. *Prima facie*, at least, it may seem that there are only two ways to explain these data. First, we can say that Hannah's epistemic position differs between the low and the high standards cases and that, for this reason, we evaluate the knowledge claims in the bank cases differently depending on whether they occur in the low or the high standards cases. (Hannah's epistemic position, here, is supposed to be determined by all those facts that

HIGH. In the original bank cases, LOW only features a knowledge ascription and HIGH only features a knowledge denial. I have changed the scenarios in this way because intuitions about the knowledge ascriptions as well as intuitions about the knowledge denials will become relevant below.

² The assumption that Hannah's knowledge ascription in the low standards cases seems true may be considered problematic on the following grounds. Hannah claims to have knowledge about whether the bank *will* be open, thus, she claims to have knowledge about the future. However, it is at least a theoretical option to hold that there are no future facts. And if there are no future facts, there can be no knowledge about future facts. Thus, on such views, Hannah's knowledge ascriptions should turn out to be invariably false. For those who share these worries, I refer to (DeRose 2009: 4f) and (Hansen and Chemla 2013: 315f). These passages present bank case like scenarios that make no reference to future facts but equally trigger data like the Contextualist Data. (I don't use these cases because the bank cases are most familiar.)

³ Note that even if it turns out that the Contextualist Data cannot be sustained, it is interesting to discuss the theoretical options to explain these data. For, the challenge from the Contextualist Data for invariantism is just one instance of a so-called "context-shifting argument" that can be employed in all sorts of other domains. Much of what I say in what follows is general enough to have a bearing on all possible applications of context-shifting arguments. See e.g. (Cappelen and Lepore 2005: Ch. 2) for further applications of context-shifting arguments.

are relevant for the question of whether she knows that the bank will be open. What these facts are precisely will depend on one's theory of knowledge. They could include facts about Hannah's beliefs, her evidence, modal properties such as how easily she could have been mistaken, etc.) Second, we can say that the knowledge claims in the bank cases express different propositions in the low and the high standards cases and that, for this reason, we evaluate them differently depending on whether they occur in the low or the high standards cases. This would mean to adopt contextualism. More specifically, the idea here would be the following. First, the knowledge claims in the low standards cases express propositions to the effect that Hannah has or doesn't have what we may call *knowledge_L* (low-standards-knowledge) that the bank will be open. Similarly, the knowledge claims in the high standards cases express propositions to the effect that Hannah has or doesn't have what we may call *knowledge_H* (high-standards-knowledge) that the bank will be open. Second, Hannah knows_L but doesn't know_H that the bank will be open in all bank cases. Thus, the knowledge ascriptions are true only in the low standards cases, and the knowledge denials are true only in the high standards cases.

Clearly, the invariantist cannot accept the second account. For, as I said, this would mean to endorse contextualism and, thus, to give up invariantism. Can she accept the first account (according to which we evaluate the knowledge claims in the bank cases differently because Hannah's epistemic position differs between these cases)? This account may also seem problematic. For, there may seem to be very general reasons (that have nothing to do specifically with invariantism) to reject this account. Let me elaborate. The low and the high standards cases seem to differ only in what is at stake for Hannah and Sarah and which error-possibilities are mentioned in their conversational context. These features, though, may seem to have no bearing on Hannah's epistemic position, that is, on the question of whether she knows that the bank will be open. Thus, it may seem that Hannah's epistemic position cannot differ between the low and the high standards cases. Given that, our varying evaluations of the bank cases cannot be explained in terms of such a difference. In sum, then, there may seem to be no way for the invariantist to explain the Contextualist Data.

One possible response strategy to this challenge on the part of the invariantist is to hold that Hannah's discussion with Sarah in the high standards cases makes a difference to her

epistemic position after all. Such a view could be spelled out in various ways. For example, so-called subject-sensitive invariantists argue that one's epistemic position is partly determined by seemingly non-epistemic parameters such as how much is at stake in one's context or which error-possibilities are salient. Similarly, we may hold that Hannah will stop believing that the bank will be open (and therefore fail to know that this is so) because her discussion with Sarah and the heightened stakes decrease her confidence in this belief. We will discuss these views in slightly more detail in a minute. For now, let me only point out, following DeRose, that all such responses are bound to fail once we consider so-called "third-person" versions of the bank cases.

Third-person cases work just like the above "first-person" bank cases. The only difference is that Hannah does not ascribe or deny knowledge to herself but to a third party who is unaware of the conversation between Hannah and Sarah (in particular, the stakes and the mentioned error-possibilities in their conversation). More precisely, such cases look as follows.

*LOW*³: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Although they generally like to deposit their paychecks as soon as possible, it is not especially important in this case that they be deposited right away. Hannah calls Bill and asks whether the bank will be open on Saturday. Bill tells her that the bank will be open and that he remembers driving by last Saturday seeing that it was open until noon. Hannah says, "Bill told me that the bank will be open. Let's drive straight home and return to deposit the paychecks tomorrow morning." Sarah responds, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays. Does Bill know the bank will be open tomorrow?" Hannah replies,

- ⊕ "Bill knows the bank will be open tomorrow."
- ⊖ "Well, no, Bill doesn't know the bank will be open tomorrow."

It turns out that the bank is open on Saturday.

*HIGH*³: Hannah and Sarah are driving home from work on a Friday afternoon. They plan to stop at the bank to deposit their paychecks, but as they drive past the bank they notice that the lines inside are very long. Hannah and Sarah have just written a very large check, and if the money from their pay is not deposited by Monday, it will bounce, leaving them in a very bad situation with their creditors. And, of course, the bank is not open on Sunday. Hannah calls Bill and asks whether the bank will be open on Saturday. Bill tells her that the bank will be open and that he remembers driving by last Saturday seeing that it was open until noon. Hannah says, “Bill told me that the bank will be open. Let’s drive straight home and return to deposit the paychecks tomorrow morning.” Sarah reminds Hannah of how important it is to deposit the check before Monday and says, “Maybe the bank won’t be open tomorrow. Lots of banks are closed on Saturdays. Besides, banks can always change their hours. Does Bill know the bank will be open tomorrow?” Hannah replies,

- $$\left\{ \begin{array}{l} \oplus \quad \text{“Bill knows the bank will be open tomorrow”} \\ \ominus \quad \text{“Well, no, Bill doesn’t know the bank will be open tomorrow.”} \end{array} \right.$$

It turns out that the bank is open on Saturday.

As before, we can ask whether Hannah’s knowledge claims seem true or false in the cases just described. And, also as before, the following response may seem plausible. Hannah’s knowledge ascription seems true in the third-person low standards cases but false in the third-person high standards case. Correspondingly, Hannah’s knowledge denial seems true in the third-person high standards cases but false in the third-person low standards cases. Let’s refer to these data as the *third-person Contextualist Data*. Again, these data may be doubted. But, as before, I will take them for granted because they have been endorsed by many philosophers.

The third-person Contextualist Data are not amenable to the just proposed invariantist response strategies (according to which the discussion between Hannah and Sarah makes a difference to the relevant subject’s epistemic position). For, suppose we grant that Bill’s epistemic position depends, say, on how much is at stake for him. Even so, we cannot argue that there is a difference in Bill’s epistemic position between the third-person low and high standards cases. After all, even if Bill’s epistemic position depends on what is

at stake *for him*, it certainly does not depend on what is at stake *for Hannah and Sarah*. And the third-person bank cases differ only in what is at stake for Hannah and Sarah. DeRose puts the point in more general terms.

In fact, our subject's [Bill's] situation is identical in our two third-person cases except for the difference that, far away from her, different conversations, which she is and will remain oblivious to, are taking place about her. So, [if the invariantist wants to maintain that the subject's epistemic position differs between the third-person bank cases, she] must maintain that how strongly our subject is positioned with respect to the bank being open on Saturdays is substantially affected by such factors as what kind of far-away conversation is taking place about the subject. That's quite implausible on its face. (DeRose 2009: 64)

So, the challenge for the invariantist I want to address in this dissertation is the challenge to explain the third-person Contextualist Data. This challenge remains even if we grant that the discussion between Hannah and Sarah may make a difference to *Hannah's* epistemic position.

In order not to overload the terminology, I will from now on discuss only the third-person cases. Correspondingly, I will refer to the third-person low/high standards cases just as the "low/high standards cases." And I will refer to the third person Contextualist Data simply as "the Contextualist Data." What I say, however, should straightforwardly apply to the first-person cases as well.

1.2. Projective Invariantism Outlined

How can the invariantist respond to the challenge from the Contextualist Data? The goal of this dissertation is to defend projective invariantism as a response to this challenge. Projective invariantists hold that, even with respect to the third-person cases, the invariantist should claim that the epistemic position of the subject, in our case, Bill, differs between the low and the high standards cases. How can the projective invariantist substantiate this claim in the face of DeRose's worry above? In a nutshell, the strategy goes as follows. First of all, we have to recognize that the bank cases are underspecified even along dimensions that should uncontroversially make a difference to Bill's epistemic position. For example, it is left open how likely it is supposed to be that the bank has changed its hours. So, before we can evaluate whether Bill's epistemic position is the same in the low and the high standards cases, we have to ask how readers of the bank cases will fill in these aspects of the scenarios, what assumptions they will *project* into the empty, yet relevant slots. Once we address this question, however, it becomes

apparent that the most plausible view is that readers will fill in these features differently in the low and the high standards cases. For example, the fact that it occurs to Sarah that the bank may have changed its hours in the high standards cases but not in the low standards cases will plausibly be taken as evidence that this possibility is supposed to be more likely in the high than in the low standards cases. Thus, Bill's epistemic position will most likely be assumed to be weaker in the high than in the low standards cases. Projective invariantists claim, very roughly, that the Contextualist Data must be explained on the basis of mechanisms of this sort.

1.3. Why not Contextualism?

As I said, I will defend a specific form of invariantism in this dissertation. So, one may wonder why one should be interested in defending invariantism at all. What is the basic motivation for invariantism? What considerations speak against contextualism? Answering this question in detail would require another dissertation. Still, I want to indicate very briefly what I consider to be the main arguments against contextualism.⁴

Contextualism, to repeat, is the negation of invariantism. Thus, contextualists hold, roughly, that sentences like "*S* knows that *p*" express different propositions in different contexts of use. We have seen above that contextualism provides a straightforward account of the Contextualist Data (in terms of the idea that the knowledge claims in the low standards cases are about knowledge_L, while the knowledge claims in the high standards cases are about knowledge_H). So, why don't we just accept contextualism?

I take the basic problem to be this. Even though the Contextualist Data speak for contextualism, there is a range of other data that pull into the direction of invariantism. The most prominent of these data, the *Invariantist Data* as I will call them, concern speech reports and disagreement. Let's discuss them in turn.

As for speech reports, it seems that in whatever context someone, *S*, asserts "*S** knows that *p*," it will at least be true to report this claim by saying "*S* said that *S** knows that *p*." For example, whatever context I now occupy, it would seem true for me to say about the

⁴ Let me briefly note here that I use the term "contextualism" in a fairly narrow sense. Contextualism, on my understanding, does not include views according to which propositional truth is relative to an epistemic standard (such as e.g. non-indexical contextualism and assessment or genuine or radical relativism). In this dissertation, such views will be considered forms of invariantism. More on this below.

low standards cases “Hannah said that Bill knows/doesn’t know that the bank will be open.” This datum seems problematic for the contextualist. For, suppose that my context relevantly differs from the context described in the low standards cases. Assume, for example, that my context is more like the context in the high standards cases, such that, say, someone just pointed out to me that the bank may have changed its hours. If so, contextualism plausibly entails that my report “Hannah said that Bill knows/doesn’t know that the bank will be open” expresses the proposition that Hannah said that Bill knows_H/doesn’t know_H that the bank will be open in the low standards cases. For, by assumption, my context is like the context in the high standards cases and, according to contextualism, “knows” expresses the relation of knowledge_H in such contexts. However, what Hannah in the low standards cases actually says, according to contextualism, is that Bill knows_L/doesn’t know_L that the bank will be open. So, my report will turn out false, not true as it intuitively seems. This is the objection from speech reports.⁵

As for disagreement, it seems that two speakers disagree with, or contradict each other when the first speaker says “S knows that *p*” and the second speaker says “S doesn’t know that *p*.” And this seems to be so even if the two speakers occupy different contexts. For example, it seems that Hannah’s knowledge denial in the high standards cases contradicts her knowledge ascription in the low standards cases. Again, this is problematic for the contextualist. For, on this view, Hannah’s knowledge denial in the high standards cases expresses the proposition that Bill doesn’t know_H that the bank will be open, while Hannah’s knowledge ascription in the low standards cases expresses the proposition that Bill knows_L that the bank will be open. And these propositions are perfectly compatible. This is the objection from disagreement.⁶

These objections have been discussed extensively and contextualists have put forward a whole battery of possible response strategies.⁷ This is not the place to evaluate these strategies. All I seek to establish is that there is a certain sort of pull towards invariantism. This pull should be undeniable. For, given invariantism, speech report and disagreement data are just exactly as we would expect them to be. My report “Hannah said that Bill knows/doesn’t know that the bank will be open” will, according to invariantism, express

⁵ See e.g. (Hawthorne 2004a: Sec. 2.7).

⁶ See e.g. (Richard 2004: 215f).

⁷ See e.g. (Blome-Tillmann 2008) and (DeRose 2009).

the proposition that Hannah said that Bill knows/doesn't know that the bank will be open. And this is, indeed, what she says according to invariantism. Similarly, Hannah's knowledge denial in the high standards cases will contradict her knowledge ascription in the low standards cases because, in the one case, she says that Bill knows that the bank will be open, while, in the other, she says that Bill doesn't know that. These then are what I take to be the major reasons to look for invariantist alternatives to contextualism.⁸

1.4. What was Invariantism Again?

I have now answered the question of what I consider the basic motivation for invariantism. Another question that needs to be addressed is what exactly it is that makes a position a form of invariantism. So far, I have characterized invariantism as the view that sentences of the form “*S* knows that *p*” express the same proposition, or say the same thing, in all contexts of use, namely, that *S* knows that *p*. In the following paragraphs, I will correct and clarify this definition.⁹

First, a correction (or, maybe, a concession). As it stands, invariantism could be proven wrong by considerations that are entirely orthogonal to the epistemological issues addressed in this dissertation. For example, (semantic) eternalism might be true, that is, it may be that propositions have “eternal” truth-values. One straightforward consequence of this view would seem to be that sentences express different propositions at different times of use. For example, the sentence “Carla is hungry” as used at a given point in time will express the eternal truth (or falsity) that Carla is hungry *at the time of the utterance of the sentence*. If that is so, “*S* knows that *p*” will also express different propositions at different times of use. Hence, invariantism, as stated, turns out false. Similar problems arise if contextualism about vagueness is correct. According to this view, sentences containing vague terms express different propositions depending on how they are sharpened in the context of use. So, given that “knows” most plausibly is vague, invariantism stands refuted if contextualism about vagueness can be established.

⁸ In the appendix to this dissertation, I will discuss an entirely unrelated objection to contextualism, namely the factivity problem. I will argue that this objection should be dismissed. I should also mention Stanley's (2005b) gradability objection to contextualism which, I think, is satisfactorily answered in (Blome-Tillmann 2008).

⁹ Note that I am assuming throughout that “*S*” and “*p*” do not contain context-sensitive expressions.

To deal with these problems, we have to modify our definition of invariantism along the following lines. Invariantists do not hold that sentences of the form “*S* knows that *p*” expresses the same proposition in *all* contexts of use, that is, that there are no two contexts of use such that “*S* knows that *p*” expresses different propositions in these contexts. More precisely, they say that there are no two contexts of use such that “*S* knows that *p*” expresses different propositions in these contexts *due to a difference in epistemic standards*. In brief, according to invariantism, the proposition expressed by “*S* knows that *p*” does not vary with the epistemic standard of its context of use.¹⁰ Here, the *epistemic standard* of a given context is supposed to be a parameter determined by features such as what is at stake in that context and which error-possibilities are salient. Thus, there will, for example, be a difference in epistemic standards between the low and the high standards bank cases.¹¹ We will say that such an epistemic standard becomes *more demanding* or *higher* the more is at stake and the more error-possibilities are salient. Thus, invariantism has it that “*S* knows that *p*” expresses the same proposition independently of how demanding or high the epistemic standard of its context of use.

The above considerations allow us to meet the just stated challenge. For, even if the proposition expressed by “*S* knows that *p*” varies with, say, the time of its utterance, it does not follow that it also varies with the epistemic standard of the context of its utterance. Thus, eternalism, etc. no longer entail that invariantism is false. Note, however that, as before, invariantism does turn out false if the propositions expressed by Hannah’s knowledge claims in the bank cases vary between the low and the high standards bank cases. For, these cases differ precisely in the epistemic standard governing Hannah’s conversational context.

It may be held that what I call “invariantism” here is not a genuinely invariantistic view because it allows for all sorts of contextual variations. Correspondingly, the above

¹⁰ I am grateful to Julia Zakkou for helpful discussion of this issue.

¹¹ What exactly epistemic standards are will not be relevant for the project at hand. In particular, we need not bother here about what exactly it means to be salient. For our purposes, it will suffice to say that an error-possibility is salient in a given conversation just in case it has been mentioned in that conversation. A more precise account would have to add further conditions. For example, it might be that an error-possibility is salient only if the participants in the conversation attend to that possibility or take the possibility seriously or don’t challenge the mentioning of this possibility, etc. Similar considerations hold for the idea that more or less is at stake in a given conversation. For our purposes, it suffices to know that the low and the high standards cases differ with respect to what is at stake for Hannah and Sarah.

considerations may be seen not as corrections of the definition of invariantism but rather as concessions to the effect that I will not be in a position to fully defend invariantism in this dissertation. I am happy to grant that and to postpone a discussion of genuine invariantism for another occasion. The view I am interested in here is invariantism as defined above, whether or not it genuinely deserves the label “invariantism.”

We see that, with these *caveats* in place, it is no longer straightforward to specify the proposition that “*S* knows that *p*” expresses according to invariantism. If eternalism is true, for example, the proposition will not just be the proposition that *S* knows that *p* but something like the proposition that *S* knows that *p* at *t*. Similarly, if contextualism about vagueness is true, the proposition will be a proposition to the effect that *S* knows_s that *p*, where the subscript “*s*” indicates a given sharpening. However, for simplicity’s sake, I will sometimes still refer to this proposition just as the proposition that *S* knows that *p*, trusting that the reader will see how what I say could be made more precise.¹²

Second, some clarifications. I have said that, according to invariantism, *what is said* or *the proposition expressed* by knowledge claims is invariant. These notions, however, have lots of possible meanings. So, let me briefly and very roughly say how they are supposed to be understood in the present context. First, I use “the proposition expressed” and “what is said” interchangeably. Second, what is said by a given utterance must be distinguished from what is merely conveyed or suggested by that utterance. For example, when I say “I am sick” in response to your question “Will you come to my party?” I suggest that I won’t come. But that is not what I say. I only say that I am sick. Third, and relatedly, what is said must be understood as what is *literally* said by a given utterance. When I hyperbolically say “I am dying of shame!” I literally say that I am dying (of shame) even though, of course, this is not at all what I mean. Finally, I don’t want it to be part of the definition that what is said is something that can be absolutely true or false (or

¹² The above considerations were considerations to the effect that our original definition of invariantism was too strong. However, it may also be too weak. According to so-called *content relativism*, the proposition expressed by “*S* knows that *p*” does not vary with the epistemic standard of the context of use but with the epistemic standard of the context of assessment, that is, the context from which the sentence is evaluated. See (Cappelen 2008a, 2008b) for such a view. The resulting view should not count as a form of invariantism. Given the above definition, however, it does. So, to be even more correct, we would have to say that, according to invariantism, the proposition expressed by “*S* knows that *p*” neither varies with the epistemic standard of its context of use nor with the epistemic standard of its context of assessment.

even true or false only relative to a possible world). Hence, I do not want to rule out by definition the view that “*S* knows that *p*” expresses propositions with truth-values that are relative to an epistemic standard. Hence, relativism (which will be discussed in detail below) will be considered a form of invariantism. These stipulations may leave the notion of what is said unspecified along many further dimensions. However, I hope that these dimensions will not be relevant for the project at hand.¹³

To repeat, invariantism is the view that the proposition expressed, or what is said by sentences such as “*S* knows that *p*” does not vary with the epistemic standard of their context of use. This is what I will consider the official characterization of invariantism in what follows. There is a slightly different characterization of invariantism, however, that I will also make use of on some occasions. According to this definition, invariantism is the view that the *relation* expressed by “knows” does not vary with the epistemic standard of its context of use. I will use this definition because it will sometimes help to simplify the discussion. Still, what I say in terms of this definition should be taken with a grain of salt. For, strictly speaking, the alternative definition is not correct.

To see this, consider the following view on the semantics of “knows.” The relation expressed by “knows” is the same in all conversational contexts. It is a three-place relation, however, that has a context-dependent epistemic standard as one argument place.¹⁴ Such a view would entail that the proposition expressed by knowledge sentences varies with the epistemic standard of their context of use. Thus, it shouldn’t come out as a form of invariantism. The presently discussed alternative definition of invariantism, however, classifies it as form of invariantism. For, the view says that the *relation* expressed by “knows” does not vary with the epistemic standard of the context of use (it is always the same three-place relation). So, strictly speaking, the alternative definition

¹³ It may be somewhat unusual to consider relativism a form of invariantism. After all, relativism has it that the truth-value of knowledge propositions is by no means invariant. The rationale behind this classification, however, should be straightforward to appreciate. I have argued in the previous section that the basic motivation for invariantism are the data from disagreement and speech reports. At least on the face of it, these data can be explained once we hold that the proposition expressed by knowledge claims is invariant, whether or not the truth-value of this proposition is invariant as well. For example, if a sentence like “*S* knows that *p*” invariably expresses the proposition that *S* knows that *p*, then a use of this sentence on the part of *S** can truly be reported as “*S** said that *S* knows that *p*” in all contexts. This seems to be so whether or not the proposition that *S* knows that *p* has invariant truth-values. Analogous considerations seem to hold for the case of disagreement. So, if we take speech report and disagreement data to be the basic motivation for invariantism, it makes sense to classify relativism as a form of invariantism.

¹⁴ Schaffer (2004a) defends such a position.

of invariantism is not correct. It should be straightforward, however, to translate what I say in terms of this definition into a terminology that uses only the official definition.

1.5. Chapter Summary

By now, we should have a relatively firm grasp on the basic motivation and the nature of invariantism. I have also outlined the specific form of invariantism I want to defend in this dissertation: projective invariantism. It remains to be seen what reasons there are to favor projective invariantism over its invariantistic rivals. I will lay out these reasons in the rest of this dissertation. The strategy will be the following. The first three chapters will be devoted to discrediting rival invariantistic accounts. In the last chapter, I will show that projective invariantism provides at least *a* plausible account of the Contextualist Data. Taken together, these results should yield a relatively strong case for projective invariantism. Here is a slightly more detailed description of the content of the chapters to come. (In the course of the subsequent discussion, it may be useful to refer to Figure 1 on p. 20, which provides a taxonomy of possible accounts of the Contextualist Data.)

In the second chapter, I will discuss (*epistemic*) *relativism* as a first invariantist reaction to the Contextualist Data and the Invariantist Data. Relativists seek to accommodate these data by claiming that, even though sentences like “*S* knows that *p*” express the same proposition in all contexts of use, they have truth-values only relative to a given epistemic standard.¹⁵ I will argue that this view may succeed in explaining the data outlined so far. However, I will also argue that it cannot do justice to all manifestations of the invariantist pull. In particular, I will argue that it cannot explain cases of indirect testimony where one speaker comes to believe that *S* doesn’t know that *p* on the basis of the testimony “*S* said that she doesn’t know that *p*.”

In the third and fourth chapter, I will discuss *pragmatic invariantism* as a second invariantist reaction to the above data. Pragmatic invariantists don’t appeal to relative propositions. Rather, they hold that assertions involving sentences like “*S* knows that *p*” *implicate* different propositions when used in contexts with different epistemic standards (even though the proposition they express remains constant). These flexible pragmatic

¹⁵ Relativism, as understood here, is supposed to include both non-indexical contextualism and genuine or assessment or radical relativism. See below.

contents, then, are supposed to account for the Contextualist Data. Pragmatic invariantism comes in two varieties depending on which part of the Contextualist Data they seek to explain in terms of implicatures.

In the third chapter, I will discuss *moderate pragmatic invariantism*, according to which the data about the high standards cases must be explained in terms of implicatures. The idea here is that the knowledge ascription in the high standards cases seems false only because it has a false implicature, and that, correspondingly, the knowledge denial seems true only because it has a true implicature. I will argue that the demise of moderate pragmatic invariantism can be seen most clearly once we get clear on the nature of the implicature that is supposed to be doing the work here. In particular, I will distinguish additive from substitutional implicatures and show that, on either interpretation, moderate pragmatic invariantism must be rejected.

In the fourth chapter, I will discuss *skeptical pragmatic invariantism*, according to which the data about the low standards cases must be explained in terms of implicatures. The idea here is that the knowledge ascription in the low standards cases seems true only because it has a true implicature, and that, correspondingly, the knowledge denial seems false only because it has a false implicature. I will argue that, contrary to appearance, *skeptical pragmatic invariantism* is a more promising position than moderate pragmatic invariantism by defending it against a range of familiar objections. However, I will also indicate a shortcoming of this view. In particular, I will show that we would have to be aware of the implicatures posited by the proponent of skeptical pragmatic invariantism in a way in which we are not.

In the fifth and final chapter, I will show that projective invariantism can account for the Contextualist Data and defend it against potential objections.

In an appendix, I will briefly come back to the question of how to argue against contextualism. In particular, I will sustain the thesis that the Invariantist Data form the best grounds to reject contextualism by showing that another familiar objection, the so-called *factivity problem*, fails. Proponents of the factivity problem argue that contextualism must be rejected because it cannot be stated properly. The basic idea behind this charge is the following. On the one hand, contextualists seemingly *have to* say that

knowledge ascribing sentences like “*S* knows that *S* has hands” are true when used in ordinary contexts. On the other hand, they arguably aren’t *allowed* to say that because the sentences in question are not true by the standard of their own context. I will show that the argument to the factivity problem fails because it rests on the mistaken premise that contextualists are committed to the truth of particular ordinary knowledge attributions.

1.6. Scope Limitations

As the foregoing summary indicates, my dissertation is limited in two respects. First, I do not discuss all invariantist positions on offer. Second, I defend invariantism only against the challenge from the Contextualist Data but not against other challenges one might want to level against this view. Due to the vast literature on the topic, limitations of this sort seemed unavoidable to me. But, let me briefly describe what will be left out and why it will be left out.

First, *subject-sensitive invariantism* will be left out. Subject-sensitive invariantists neither appeal to relative propositions nor to merely conveyed contents. Instead, they hold that a subject’s epistemic position depends on features such as what is at stake for the subject or which error-possibilities she has in mind (that is, they deny what has been called *intellectualism*, roughly the view that knowledge depends solely on truth-conducive features).¹⁶ On that basis, they explain the Contextualist Data with respect to the first-person cases by simply pointing out that the subject’s epistemic position varies between these cases.

The basic reason why I don’t discuss this view has been indicated already. In this dissertation, I set myself to the task of explaining the Contextualist Data about the first-person as well as the third-person bank cases. Subject-sensitive invariantism, however, can only explain the data about the first-person bank cases: Even if Bill’s epistemic position depends on stakes and salient alternatives, it certainly does not depend on what is at stake *for Hannah and Sarah* or which error-possibilities *they* have in mind. Subject-

¹⁶ Versions of this view are put forward e.g. in (Hawthorne 2004a) and (Stanley 2005b).

sensitive invariantism, thus, is no genuine alternative to the view I myself will put forward below.¹⁷

The second invariantist position that will not be discussed in detail is the *belief removal theory*. This theory neither appeals to relative contents nor to merely conveyed contents nor to anti-intellectualism. Still, as regards the first-person bank cases, it argues that Hannah's epistemic position varies between the low and the high standards cases. The reason is supposed to be that, in light of the heightened stakes and the mentioned error-possibility in the high standards cases, she will stop believing that the bank will be open in these cases (even though her evidence would suffice for knowledge if she were to maintain that belief).¹⁸ As noted already, this strategy does not carry over to the third-person cases either. For, there is no reason why Bill should stop believing that the bank will be open in light of the heightened standards in Hannah and Sarah's conversation. As before, this is why I don't discuss this view.¹⁹

Finally, I don't discuss *psychological error-theories*. This family of theories is united by the idea that one or the other judgment about the bank cases is mistaken and must be explained away by appeal to general psychological biases. These theories, if correct, do have a good claim to explain all the relevant data. The reason why I still don't discuss them is that I have nothing to contribute to the recent debate. This debate increasingly focuses on results from psychology that I am not in a position to evaluate properly.²⁰

¹⁷ Note that these considerations don't show that subject-sensitive invariantism must be rejected. For, a subject-sensitive invariantist may apply her theory only to first-person cases and, at the same time, appeal to some other invariantist account to deal with third-person cases. Hawthorne (2004a: 162–166) and Stanley (2005b: 101–103), for example, supplement their theory with one or the other psychological error-theory to account for third-person cases.

¹⁸ Bach seems to be the only proponent of this view. See e.g. (Bach 2005: Sec. 5). The label is taken from (Hawthorne 2004a: 169).

¹⁹ Again, this does not show that the belief removal theory is wrong. For, as before, the proponent of the belief removal theory could apply her theory only to the first-person cases and explain the third-person cases in some other way. Bach, for example, freely admits that his theory is "a kind of error-theory" (2005: 78) when it comes to third-person cases.

²⁰ The most prominent form of psychological error-theory appeals to the so-called availability heuristic and is put forward in (Vogel 1990: 19f), (Feldman 2001: 76), (Hawthorne 2004a: 162–166, 2004b: 520ff) and (Williamson 2005a: 226). (The availability heuristic is a heuristic to evaluate the likelihood of given events on the basis of their cognitive "availability," that is, very roughly, the ease with which they can be imagined or remembered. It allegedly leads us to overestimate the likelihood of error-possibilities once they are made salient.) Nagel (2010) forcefully criticizes this approach on the basis of psychological results that indicate that the availability heuristic will not have the effects it is supposed to have. However, more recent approaches seem unaffected by this criticism. See e.g. (Nagel 2010, 2011) and (Gerken 2011, 2012).

The scope of this dissertation is not only limited in that I leave out the just discussed forms of invariantism. It is also limited in that I consider only one of potentially many arguments against invariantism.²¹ Most importantly, I don't discuss the idea that only contextualism provides a plausible response to certain kinds of skeptical arguments.²² A full defense of invariantism would have to address this issue. But, as will become apparent below, the discussion of skeptical arguments can plausibly be separated from my concerns here. So, it will be left for another occasion.

²¹ For an overview of further such arguments, see (Rysiew 2009).

²² See e.g. (DeRose 1995).

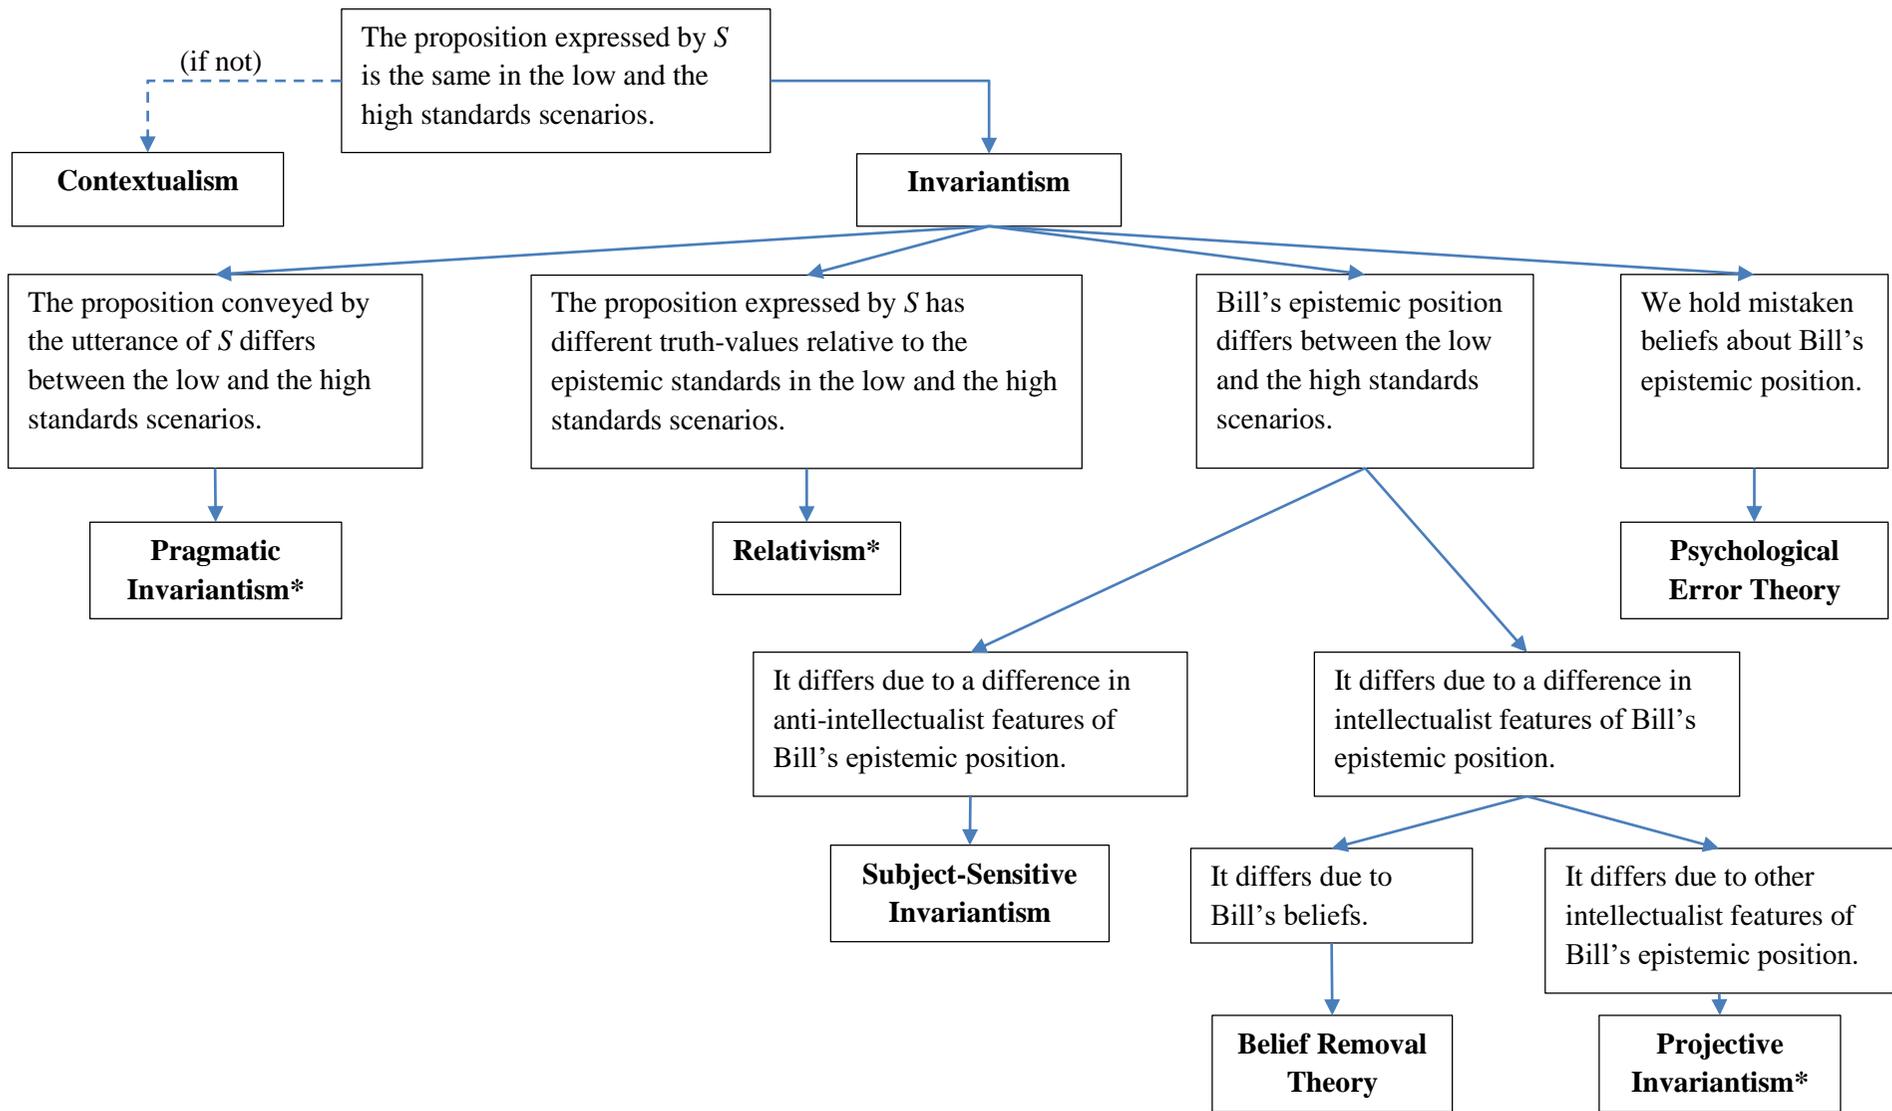


Figure 1: Candidate accounts of the Contextualist Data. “S” can stand for “Bill knows that the bank will be open” or “Bill doesn’t know that the bank will be open” as used by Hannah in the bank cases. “*” marks positions I will discuss in detail in this dissertation. Reading the chart: Start at one of the leafs. Follow the arrows leading to that leaf up to the top box. The position named in the leaf entails all the claims you pass (insofar as the position is applied to deal with the Contextualist Data; as indicated already, subject-sensitive invariantists, for example, may (and should) refrain from applying their theory to these data). Some remarks. First, the claims leading to a given position do not provide necessary and sufficient conditions for this position. They only describe what this position would say about the bank cases to explain the Contextualist Data (insofar as it is applied to these data). Second, the chart rightly suggests that all invariantist positions discerned are in principle compatible. Clearly, though, many combinations make no sense because they would yield redundant accounts of the data. Fourth, the chart is not exhaustive. It is supposed to cover only the views that have been prominently defended in the literature (and my own view). Finally, the classification of my view, projective invariantism, is not entirely correct. For, it is somewhat misleading to say that, according to this view, Bill’s epistemic position *differs* between the low and the high standards scenario. More precisely, the view has it that Bill’s epistemic position *will be assumed to differ* between these scenarios.

2. Relativism

In this section, I will discuss epistemic relativism, for short, relativism. Relativism is a view that seeks to explain the Contextualist Data and, at the same time, to do justice to the pull towards invariantism, in particular, to the Invariantist Data. In order to explain the Invariantist Data, relativists endorse invariantism, that is, that the proposition expressed by sentences like “*S* knows that *p*” does not vary with the epistemic standard of their context of use. In order to explain the Contextualist Data, they further hold that this proposition has truth-values only relative to a given epistemic standard. In particular, relativists hold that the proposition that Bill knows that the bank will be open is true relative to the epistemic standard in the low standards cases but false relative to the more demanding epistemic standard in high standards cases.¹ We need not bother about the details here because I will just grant that relativism, if true, can account for the Contextualist Data and the Invariantist Data.² Instead, I will argue that relativism fails because it does not explain further manifestations of the invariantist pull. More specifically, I will argue that relativism cannot explain our reactions to certain cases of *indirect testimony*, that is, to reports of the form “*S* said that *p*.”³ Even more specifically, I will argue that relativism cannot explain the following datum: When we find ourselves in a low standards context and someone reports, “*S* said that she (*S*) doesn’t know that *p*,” it is appropriate to form the belief that *S* doesn’t know that *p* as a response.

Here is an outline of the present chapter. First, I will briefly elaborate on the definition of epistemic relativism. Second, I will present the just mentioned datum from indirect testimony in more detail. Third, I will show that relativism is incompatible with this datum. Fourth, I will discuss and reject rejoinders to this objection. Finally, I will pick up some loose ends. In particular, I will discuss, first, why I neither apply the present objection to contextualism nor to other forms of relativism (such as relativism about

¹ See e.g. (Richard 2004; 2008), (Kompa 2005, 2011), (Brogaard 2008) and (MacFarlane 2005b; 2011) for proponents of relativism. Readers familiar with the debate will realize that relativism, as understood here, is neutral between so-called non-indexical contextualism and assessment (or genuine) relativism. This is so because the distinction between these two kinds of relativisms is irrelevant for the purposes at hand. My subsequent argument equally applies to both these views.

² For further discussion particularly of the disagreement issue, see e.g. (MacFarlane 2014: Ch. 6).

³ The terminology of “direct” and “indirect” testimony is taken from (Lackey 1999: 479f).

what's tasty). Second, I will explain why it is important for my case to consider indirect as opposed to direct testimony, why the reported knowledge claim must be a knowledge denial and why the report must be made in a low standards context.

2.1. Relativism

Relativism, to repeat, consists of the following two claims. First, the proposition expressed by “*S* knows that *p*” does not vary with an epistemic standard. It is always just the proposition that *S* knows that *p*. Second, this proposition has truth-values only relative to an epistemic standard. So, according to relativism, the proposition that *S* knows that *p* may be true relative to one standard but fail to be true relative to another. As a first pass, we can say that it will be true relative to a given standard just if *S* knows relative to that standard that *p*. What does it mean to know something relative to a given standard? The answer will depend on one's theory of knowledge. For example, from the perspective of a relevant alternatives framework,⁴ to know that *p* relative to a given standard means to be able to rule out the set of alternatives to *p* that is determined by the given standard. Here, more demanding standards will determine larger sets of alternatives than less demanding standards. The objection I will present in what follows figures on a level of abstractness that does not make it necessary to settle on a specific interpretation of what it means to know something relative to a given epistemic standard. So, I will continue to speak of knowledge relative to a given standard, leaving it to the reader to fill in her preferred theory of knowledge. Much more could be said about the nature of relativism. But the further details will be irrelevant for my present concerns. So, I leave them for another occasion.⁵

2.2. The Indirect Testimony Datum (ITD)

Let's now turn to the datum from indirect testimony, which I will later show to be incompatible with relativism. Consider the following scenario.

Knowledge Testimony Case: John and his journalist friend Sally sit in the pub talking about a criminal who has recently been caught by the police. The criminal was caught hiding in an office building. Sally and John are wondering how the building came into

⁴ See e.g. (Lewis 1996).

⁵ A thorough and wide ranging introduction into the topic can be found in (MacFarlane 2014: Ch. 1–6, 8).

the focus of the investigation. Sally was at a press conference where the police officer in charge was giving a statement and answered questions. She reports, “There was some evidence that the perpetrator would be in the building. But the police officer said that, when she gave the order to storm the building, she didn’t know that the perpetrator would be in there.” John reasonably considers Sally and the police officer sincere and well-informed (and everything else is normal).

I take the following to be an intuitive, pretheoretical datum that any theory of knowledge must accommodate. In the above scenario, it would be perfectly appropriate for John to respond to Sally’s report by forming not only the belief that the police officer *said* that she didn’t know that the perpetrator would be in the building but also the belief that the police officer did in fact not know that the perpetrator would be in the building. (At the very least, I take it, this is what ordinary speakers in John’s situation would often do. So, insofar as we do not want to say that ordinary speakers are mistaken in this behavior, we should grant that this behavior is appropriate.⁶) Let’s refer to this datum as the “Indirect Testimony Datum” (ITD). If we use “*P*” to refer to the proposition that the perpetrator was in the building, we can restate the (ITD) in the following way:

(ITD) In the *Knowledge Testimony Case*, it would be appropriate for John to form the belief that the police officer didn’t know that *P* in response to Sally’s report (for short, it would be appropriate for John to *accept* Sally’s indirect testimony).

To be clear, let me emphasize that the *Knowledge Testimony Case* is described such that John has good reasons to believe that the police officer and Sally are sincere and well-informed (and everything else is normal). The (ITD) is supposed to obtain only if these assumptions are in place. So, I am not denying that there are many circumstances in which we should not accept the indirect testimony of somebody else. For example, it would clearly be inappropriate for John to accept Sally’s testimony if, say, he had reasons to consider Sally or the police officer insincere or ill-informed. To repeat, this observation is perfectly compatible with the (ITD). For, this datum is supposed to obtain only in cases

⁶ One possible response to my objection to relativism would be to endorse an error-theory according to which people would form the belief above, but mistakenly so. Relativists, however, are fairly reluctant to endorse such error-theories. See e.g. (MacFarlane 2005b: 215f). So, I will ignore this response strategy throughout this chapter. Let me also note that even those who want to endorse an error-theory must provide an account of why we commit the relevant kind of error. It seems not at all trivial to provide such an account.

were the relevant speakers can be assumed to be sincere and well-informed. (The normality condition is supposed to rule out scenarios where John is not even listening to what Sally says or doesn't speak her language, etc. In such situations, it would also be inappropriate for John to form the relevant belief. But, again, this is perfectly compatible with the (ITD).)

A further clarificatory remark: It may seem that Sally somehow suggests that she thinks that the police officer is not to be trusted. For, this might seem to be the most straightforward explanation for why she only reports the police officer instead of directly endorsing what the police officer said. This, in turn, may seem to make it inappropriate for John to come to accept what the police officer supposedly said. Note, though, that we can happily grant that Sally suggests, for example, that the police officer wasn't sincere. As long as *John* considers the police officer sincere (and does not change his judgment in response to Sally's suggestion), he should still come to accept what the police officer said. We may imagine, for example, that John is already aware of the fact that Sally considers the police officer insincere but finds that the only evidence she has for this claim is preposterous. Under these assumptions, the (ITD) should hold.

To see the intuitive plausibility of the (ITD) more clearly, it is useful to consider the following comparison. Suppose the police officer hadn't said that she didn't know that the perpetrator would be in the building. Suppose, instead, that what she said had been a clearly non-relativistic proposition. For example, suppose that the police officer had said that there was some evidence that the perpetrator would not be in the building. And suppose further that Sally had reported this assertion in the following way, "There was some evidence that the perpetrator would be in the building. But the police officer said that there also was some evidence that the perpetrator would be somewhere else." In this case, it should clearly be appropriate for John to form not only the belief that the police officer *said* that there was some evidence that the perpetrator would be somewhere else but also the belief that there in fact was some evidence that the perpetrator would be somewhere else (assuming, as before, that John considers the police officer and Sally sincere and well-informed). The (ITD) is based on the idea that, intuitively, the situation is exactly the same when Sally makes the report she makes in the *Knowledge Testimony Case* above, that is, when she reports, "There was some evidence that the perpetrator

would be in the building. But the police officer said that she didn't know that the perpetrator would be in there." In this case as well, it would be perfectly appropriate for John to form not only the belief that the police officer *said* that she didn't know that the perpetrator would be in the building but also the belief that the police officer did in fact not know that the perpetrator would be in the building.

To see the distinctiveness of the (ITD), it is also useful to compare the *Knowledge Testimony Case* to a scenario where no such datum obtains. Suppose, for example, that John and Sally are wondering whether to go to an Indian restaurant that has opened recently. They haven't been there before. Sally says, "My sister said the food there is tasty." Now, it seems very clear that, even if John considers Sally and her sister sincere and well-informed, he will still not form the belief that the food at the Indian place is tasty, and it would also not be appropriate for him to do so. At best, John will form a belief to the effect that the food at the Indian place *is supposed to be* tasty. After all, it remains to be seen whether John and Sally's sister share a taste in food. So, in the taste case, we don't obtain a datum like the (ITD).

So much for the (ITD). By now, this datum should seem acceptable (if it doesn't, see section 2.4 below). But before we go on to address the question of why the (ITD) should be problematic for the relativist, let me emphasize that some aspects of the *Knowledge Testimony Case* that may seem accidental are, in fact, crucial for my case. I will explain later why these aspects are crucial. For now, let me only point out what these aspects are in order to forestall potentially distracting objection. First, the *Knowledge Testimony Case* involves *indirect* as opposed to direct testimony, that is, Sally does not herself endorse what the police officer said but only reports that the police officer said it. Second, what the police officer said is a *negated* knowledge claim. Third, Sally and John find themselves in a *low standards* situation where nothing is at stake and no error-possibilities are salient. To repeat, all these aspects are crucial for my case. Thus, I am happy to grant that the objection I will develop does not arise if we modify the *Knowledge Testimony Case* in any of these respects.

2.3. An Argument Against Relativism Based on the (ITD)

We will now turn to the question of why the (ITD) should be problematic for the relativist. In a nutshell, the argument for this conclusion goes as follows. First premise: Assuming relativism, John can appropriately form the belief that the police officer didn't know that *P* only if he is in a position to justify the belief that *it is true relative to his (John's) epistemic standard* that the police officer didn't know that *P*. Second premise: John is not in a position to justify the belief that it is true relative to his (John's) epistemic standard the police officer didn't know that *P*. (The only belief in the vicinity that he would be in a position to justify is the belief that it is true relative to the police officer's epistemic standard the police officer didn't know that *P*.) Conclusion: Assuming relativism, John cannot appropriately form the belief that the police officer didn't know that *P*. Thus, relativism is incompatible with the (ITD), according to which John can appropriately form that belief. In what follows, I will discuss the two premises of the above argument in more detail. Let's start with the first.

A Relativistic Justification Norm for Belief

To repeat, this premise has it that, assuming relativism, John can appropriately form the belief that the police officer didn't know that *P* only if he is in a position to justify the belief that *it is true relative to his (John's) epistemic standard* that the police officer didn't know that *P*. I take this premise to be relatively uncontroversial.⁷ So, in what follows, I will only very briefly motivate why one might be willing to accept it. In particular, I will first state a very general norm for the appropriateness of beliefs, namely, the justification norm. Then I will modify this norm so that it fits the relativist semantics. Finally, I will derive the above premise as an instance of the relativistic justification norm.

Consider the appropriateness of beliefs in general. Under what conditions does a belief count as appropriate? It should be relatively uncontroversial that speakers commit at least some sort of mistake if they come to believe what they cannot justify. So, given that mistaken beliefs aren't appropriate, it should also be relatively unproblematic to endorse a principle along the following lines.

⁷ See FN 9 for references.

One can appropriately form a given belief only if one is in a position to justify that belief.

Let's refer to this principle as the *justification norm (for belief)*. In what follows, I will take it for granted that a norm at least along the lines of the justification norm obtains. The question I want to address is how to adapt the justification norm to the relativist semantics.

To begin with, note that the relativist (and maybe others as well) will most likely hold that the level of justification required for appropriate belief depends on the context. So, to be more precise, we should substitute the justification norm by the following thesis.

In a given context, one can appropriately form a given belief only if one is in a position to provide a justification for that belief that satisfies the epistemic demands of the context.

This kind of contextualization, however, will not be relevant for my concerns. So, to simplify the discussion, I will just pretend that the standard of justification is the same in all contexts. So far, then, the relativist can stick to the justification norm as initially stated.

Assuming relativism, however, another issue has to be addressed: We have to understand what it means to justify a belief that has a relativistic content. Let's consider this issue in more detail. The justification norm as initially stated has it that one can appropriately form a belief only if one would be in a position to justify that belief. But suppose the belief in question has a relativistic content. Suppose, for example, that the belief has a content, call it "*r*," that is expressed by a knowledge sentence. Now, what does it mean to say that one would be in a position to justify the belief that *r*? Given that *r* is a relativistic content, *r* can be true relative to one epistemic standard but false relative to another. Correspondingly, one can be in a position to justify the belief that *r* is true relative to one epistemic standard even if one would not be in a position to justify the belief that *r* is true relative to another epistemic standard.⁸ Thus, it is not at all clear to which of these

⁸ Note that this is so even if the level of justification required for both beliefs is the same, that is, even if we accept the simplification made in the previous paragraph. The reason why one can be justified in holding one of the beliefs in question even if one would not be justified in holding the other is just that the beliefs in question have different contents. One of the beliefs is about whether *r* is true relative to one epistemic standard, the other belief is about whether *r* is true relative to another epistemic standard.

claims one is committed when one just plainly says that someone is in a position to justify the belief that r . So, the justification norm as initially stated allows for different interpretations when we consider beliefs about relativistic contents. Somewhat differently put, the relativist can accept the following neutral version of the justification norm:

One can appropriately form a belief that r only if one is in a position to justify the belief that r is true relative to the epistemic standard(s) X .

The question she needs to answer, however, is which epistemic standard(s) “ X ” is supposed to denote.

In principle, relativists can answer this question in various ways. In what follows, however, I will assume that the relativist says what seems most straightforward and most plausible, namely, that it is one’s own epistemic standard that is relevant for the appropriateness of one’s beliefs. Given that, we arrive at the following relativist version of the justification norm.

One can appropriately form a belief that r only if one is in a position to justify the belief that r is true relative to one’s own epistemic standard.

An instance of this norm is the following claim. John can appropriately form the belief that the police officer didn’t know that P only if he is in a position to justify the belief that *it is true relative to his (John’s) epistemic standard* that the police officer didn’t know that P . So, we have derived the first premise of the argument against relativism above.⁹

⁹ I have said above that this premise should be relatively uncontroversial among relativists. Admittedly, though, relativists haven’t yet considered the justification norm for belief and, correspondingly, haven’t directly endorsed the relativistic justification norm above. Still, relativists tend to accept a truth norm for belief that straightforwardly leads to the relativistic justification norm above. Let me elaborate. Relativists generally endorse (or at least express sympathies to) the following truth norm for belief. One can appropriately form a given belief only if that belief is true relative to one’s own epistemic standard. (See e.g. (Kölbel 2003: 70, 2004: 309) and (MacFarlane 2014: 116). Egan et al. (2005: 153) and Brogaard (2008: 448) endorse this norm at least for the case of assertion.) If the relativist accepts this norm, she will also have to accept that people commit at least some sort of second order mistake if they form a belief without *being in a position to justify* that this belief is true relative to their own epistemic standard. Thus, relativists who accept the just stated truth norm will also have to accept the relativist justification norm described in the main text. So, given that most relativists accept the truth norm, the relativist justification norm should be fairly uncontroversial as well.

Failing to Satisfy the Relativistic Justification Norm

Let's now turn to the second premise of my argument against relativism. To repeat, this premise has it that John is not in a position to justify the belief that it is true relative to his epistemic standard that the police officer didn't know that *P*. In this section, I will provide the first steps to defending that premise. To begin with, however, let me introduce some straightforward abbreviations that hopefully help to make the subsequent discussion more tractable.

I will say that the police officer didn't know-*E* that *P* just in case it is true relative to a given epistemic standard, *E*, that the police officer didn't know that *P*. I will refer to John's epistemic standard as "*J*" and to the police officer's epistemic standard as "*O*." Thus, I will say that the police officer didn't know-*J* that *P* just in case it is true relative to John's epistemic standard that the police officer didn't know that *P*. Correspondingly, I will say that the police officer didn't know-*O* that *P* just in case it is true relative to the police officer's epistemic standard that the police officer didn't know that *P*.

Given these abbreviations, the premise we are trying to establish amounts to the following claim. John is not in a position to justify the belief that the police officer didn't know-*J* that *P*. Why should this premise hold? The basic reason to accept this premise is this.

We can stipulate, I think, that before Sally makes her report, John has no information about the epistemic position of the police officer. Thus, if John is in a position to justify the belief that the police officer didn't know-*J* that *P*, this justification must somehow stem from Sally's report. But what can John learn from that report? Given Sally's report, John presumably is in a position to justify the belief that the police officer said that she didn't know that *P*. After all, Sally just told him that this is so. It may also be that John would therefore be in a position to justify the belief that the police officer didn't know-*O* that *P*. For, if the police officer said that she didn't know that *P*, it seems reasonable for John to assume that this claim was true relative to the police officer's epistemic standard, *O*.¹⁰ However, even if John would be in a position to justify the belief that the police officer didn't know-*O* that *P*, it just doesn't follow that John would also be in a position

¹⁰ I am assuming here that assertions are governed by a justification norm corresponding to the justification norm for beliefs derived above, that is, that one may assert that *p* only if one is in a position to justify the belief that *p* is true *relative to one's own* epistemic standard.

to justify the belief that the police officer didn't know- J that P . For, it may be that O and J differ. In particular, it may be that O was more demanding than J . (Remember that John's epistemic standard, J , is supposed to be very low.) And if that is so, it may be that the police officer didn't know- O that P even if she did know- J that P . Thus, it is very much unclear how John could come to be in a position to justify the belief that the police officer didn't know- J that P on the basis of Sally's report. Hence, it is very much unclear how John could come to be in a position to justify that belief at all.

Admittedly, John would be in a position to justify the belief that the police officer didn't know- J that P if he were in a position to justify the belief that his epistemic standard, J , and the police officer epistemic standard, O , are the same (or differ only marginally, I take this qualification to be tacitly understood throughout). For, if John would be in a position to justify that J and O are the same, then he could derive that the police officer didn't know- J that P from his arguably justified belief that the police officer didn't know- O that P .¹¹ However, it is very much unclear how John could come to be in a position to justify the belief that J and O are the same. After all, Sally does not tell him what was at stake or which error-possibilities were salient when the police officer made her assertion, and John wasn't at the press conference himself.

So, to sum up the discussion up to now, the objection I want to raise against relativism is this. Relativism is incompatible with the (ITD). According to this datum, it is appropriate for John to form the belief that the police officer didn't know that P in the *Knowledge Testimony Case*. However, relativism entails that it should be inappropriate for John to form this belief. For, relativism entails that this belief is appropriate only if John is in a position to justify the belief that the police officer didn't know- J that P . And John is not in a position to justify the belief that the police officer didn't know- J that P . The basic reason for this last claim, in turn, seems to be that John is not in a position to justify the belief that his and the police officer's epistemic standard are the same.

In what follows, I will address possible rejoinders to the present argument against relativism. In particular, I will address ways of rejecting the second premise of this

¹¹ In fact, it would even suffice if John was in a position to justify the belief that either O and J are the same or O is less demanding than J . For, if O is less demanding than J , then one can arguably derive knowledge- O from knowledge- J . But, to simplify the discussion, let's just assume that John's epistemic standard, J , is as low as it can get. Given that, we can ignore the possibility that O is less demanding than J .

argument (according to which John fails to be in a position to justify that the police officer didn't know-*J* that *P*). Before that, however, I will turn to those who still do not feel fully convinced by the (ITD), the datum on which the argument above is based. In particular, I will show that, in addition to the (ITD), there is a closely related and maybe more straightforward datum that could equally well be used to underwrite my argument.

2.4. An Alternative to the (ITD)

In what follows, I will briefly state the just mentioned alternative datum and indicate why it leads to the same problems as the (ITD) above.

Consider the *Knowledge Testimony Case* again. Remember Sally's following report.

Sally: There was some evidence that the perpetrator would be in the building. But the police officer said that, when she gave the order to storm the building, she didn't know that the perpetrator would be in there.

Suppose you reject the (ITD) and thus hold that John would refrain from forming the belief that the police officer didn't know that *P* in response to Sally's report. Even so, I take it, you may still be happy to grant that it would be very weird for John to indicate his unwillingness to form this belief by responding to Sally in any of the following ways.

#John: But was the police officer's epistemic standard the same as ours?

#John: But maybe the police officer considered more error-possibilities than we do.

#John: But I guess the issue was more important for the police officer than it is for us now.

So, I take it as another datum that any of these responses on the part of John would be inappropriate.

This datum leads to trouble for the relativist for more or less the same reasons as the (ITD). Here is why. As argued already, relativism entails that John can appropriately form the belief that the police officer didn't know that *P* only if John is in a position to justify the belief that the police officer didn't know-*J* that *P*. However, as also argued already, John only is in a position to justify the belief that the police officer didn't know-*O* that *P*. So, given relativism, at least one reason for John to refrain from accepting Sally's

testimony is that, for all he knows, his epistemic standard, *J*, and the police officer's epistemic standard, *O*, differ relevantly. The above responses are straightforward ways to articulate this reason. So, they should be perfectly appropriate if relativism holds. Thus, relativism seems incompatible with the idea that the responses above are inappropriate.

One might want to respond that, even if a difference in epistemic standards would indeed be a good reason for John not to accept Sally's testimony, it does not follow that we would find it appropriate for John to articulate this reason. Such a response, however, seems strangely *ad hoc*. Suppose, for example, that John thought the police officer was insincere (contrary to the stipulation in the scenario). In that case, it would clearly be natural for him to respond to Sally's report by saying something along the following lines, "But she's a liar. I don't believe a word of what she says." Or suppose that John considered the police officer ill-informed (again, contrary to the stipulation in the scenario). In that case, it would clearly be natural for him to say something along the following lines, "But she certainly just forgot the pictures taken by the drones. The perpetrator was clearly visible there." So, whatever reason John has for not accepting Sally's testimony, it seems perfectly appropriate for him to articulate it. So if, as the relativist has it, one reason for John not to accept Sally's testimony is that his and the police officer's epistemic standard may have differed, then it should seem perfectly appropriate for him to respond to Sally's testimony by pointing this out.

This last point can be further corroborated by considering the case of taste again. Remember the scenario where John and Sally talk about the Indian restaurant and Sally says, "My sister said the food there is tasty." I mentioned already that John should be unwilling to respond to this report by forming the belief that the food at the Indian restaurant is tasty. I have also indicated that the reason for this is that his and Sally's sister's taste may differ. Correspondingly, it is perfectly appropriate for John to articulate that reason in response to Sally's report. For example, it seems perfectly sensible for him to respond by saying things like this.

John: But does she like the same things as I do?

John: Then I guess she likes coriander? I hate it.

John: But I doubt that we have the same taste in food.

So, we have a further case in point to substantiate the general idea that it is perfectly appropriate to respond to testimonial reports by indicating why one is unwilling to accept them.

Given that we accept this general principle, we thus have an argument against relativism that is based on the same premises as the argument developed above, with the only difference that it does not appeal to the (ITD), but to the datum that it would be inappropriate for John to respond to Sally's report by pointing out that his and the police officer's epistemic standard may have differed. The discussion in the rest of this chapter will be couched in terms of the (ITD). But everything I say should apply *mutatis mutandis* to the datum presented in this section.

2.5. Replies and Rejoinders

Let us now turn to potential rejoinders to my above argument. As I said, the argument crucially depends on the premise that John is not in a position to justify the belief that the police officer didn't know- J that P . We have seen that this premise, in turn, is based on the assumption that John is not in a position to justify the belief that J equals O . So, one general strategy to respond to my argument is to reject the idea that John is not in a position to justify the belief that J equals O . In what follows, I will address and reject various ways to substantiate this response strategy.

John Can Take it for Granted that O equals J

The relativist could argue that, even though it is *possible* that O was relevantly more demanding than J , this is very unlikely. Thus, John can simply take it for granted that J equals O .¹² Such a response could be substantiated in various ways. For example, the relativist could hold that epistemic standards only change when certain institutional settings change, say, when we move to the court of law. Alternatively, she might want to say that epistemic standards change only when the philosopher enters the scene and presents alternatives that would otherwise be ignored.

¹² See e.g. (Bach 2005: 61) for a proposal along these lines on behalf of the contextualist as a response to certain aspects of the objection from speech reports, in particular, that we seem happy to disquotationally report knowledge claims even when we have no idea about the epistemic standard that governed the reported assertion.

These responses are possible. In the present dialectical context, however, they can be ignored. The basic reason is that, if we accept these responses, relativism will no longer have a bearing on the Contextualist Data. Let me elaborate. Clearly, the high standards bank cases neither involve a special institutional setting nor a philosopher who raises far-fetched error-possibilities. More generally, these cases do not seem to describe situations that are particularly unlikely to obtain. Certainly, John cannot just take it for granted that the police officer wasn't in a comparable situation. For example, it seems indeed quite likely that the whereabouts of the perpetrator were very important for the police officer. And it also seems quite likely that the reporters at the press conference questioned the evidence the police officer was presenting, thus making error-possibilities salient that John and Sally ignore. So, if the relativist wants to maintain that there is a relevant shift in epistemic standards between the bank cases, she cannot say that John can just take it for granted that there was no difference between his and the police officer's epistemic standard. Thus, if the relativist wants to explain the Contextualist Data, she must reject the responses above. So, since we are considering relativism *qua* account of the Contextualist Data, these responses can be ignored.

Sally Conveys that O equals J

Here is an alternative strategy to motivate the idea that John is in a position to justify the belief that *O* equals *J*. One might want to say that John is in a position to justify that belief because Sally somehow conveys that *O* equals *J*. Again, this proposal could be implemented in various ways. For example, we could say that Sally conversationally implicates that *O* equals *J*. Alternatively, we could say that she conventionally implicates or presupposes this proposition.

Even if we abstract away from the specific implementations of this proposal, however, a very general problem arises. Let's stipulate that, in the *Knowledge Testimony Case*, the stakes at the press conference were very high and that the reporters had raised error-possibilities that Sally and John do not consider. Given these stipulations, it should be the case that *O* does not equal *J*, in particular, *O* should now be significantly more demanding than *J*. Now suppose that, as stated above, Sally conveys that *O* equals *J* when she reports the police officer. It would follow that, in the situation just described, Sally's report should sound misleading. For, her report would now convey a falsehood. I would think,

however, that Sally's report would not sound misleading at all. To the contrary, it would sound just perfectly fine. So, the above account makes wrong predictions about the appropriateness of speech reports.

That said, I fully agree that Sally's report would indeed sound misleading if the police officer's standard had been *excessively* high. Suppose, for example, that the police officer had been discussing skepticism. In such a case, a qualified report along the following lines would sound much more appropriate than Sally's original report, "The police officer said that she didn't know that the perpetrator would be in the building. But she also said that she didn't know she has hands because she might be living in the Matrix. So, don't bother." But, even if Sally's original report sounds weird when the police officer's epistemic standard is assumed to be excessively high, it does not follow that it also sounds weird when the police officer's epistemic standards is assumed to be high but not excessively high.

Let me also note that I only want to maintain that a knowledge claim can be appropriately reported without qualification in cases where the reported knowledge claim is a *denial* and the reporting context is *less demanding* than the reported context (that is, cases such as the *Knowledge Testimony Case*). There are asymmetries depending on whether the reported claim is a knowledge ascription or a knowledge denial and on whether the reporting context is more demanding or less demanding than the reported context. DeRose (2009: 170f), for example, argues that "in cases of large, noticeable, and important shifts in standards" unqualified speech reports seem wrong. However, the only case he considers to substantiate this thesis is one where a knowledge *ascription* is reported without qualification in a case where the reporting context is *more demanding* than the reported context. I agree with DeRose that, in such a case, the unqualified report would seem somewhat wrong. But I deny that this observation generalizes.¹³

Sally Implicates that O equals J

We now have a very general objection to the proposal just mentioned, according to which Sally conveys that *O* equals *J*. The proposal leads to further problems once we specify its

¹³ The asymmetries here seem closely linked to the idea that high-standards contexts are generally considered to be more "enlightened" contexts. See e.g. (Hawthorne 2004a: 164). More asymmetries of this kind will be discussed below.

details. Suppose, for example, that we want to hold that Sally conveys that *O* equals *J* in terms of a conversational implicature. Given this assumption, two further problems arise.

The first problem is this. People are normally aware of what they implicate. Note the corresponding naturalness of the following exchanges where *B*'s first utterance has a certain implicature that *A* then makes explicit.

A: Will you come to the party?

B: I am sick.

A: So, you mean that you won't come?

B: Exactly.

A: Where does *C* live?

B: Somewhere in the south of France.

A: So, you don't know where exactly?

B: Right.

Natural conversations of this type seem to be available for all kinds of implicatures. But now suppose that John were to make explicit the alleged implicature that *J* equals *O* of Sally's report. The resulting dialogue would go as follows:

Sally: The police officer said that, when she gave the order to storm the building, she didn't know that the perpetrator would be in there.

#John: So, you mean that

- her epistemic standard was the same as ours?
- the issue is as important for us as it was for her?
- she didn't consider potential mistakes we are not considering right now?

#Sally: Exactly.

John's questions just as well as Sally's response seem very weird. This makes it very unlikely that Sally should implicate that *J* equals *O*.

The second problem is this. Implicatures are cancelable. However, even if Sally were to cancel the implicature that *J* equals *O*, John would still accept her testimony. In detail: Suppose that Sally would use one of the following claims to report the police officer.

Sally: Another failed operation would have meant the end of her career. Still, the police officer said that, when she gave the order to storm the building, she didn't know that the perpetrator would be in there.

Sally: In the course of a thorough evaluation of the evidence and its potential deficiencies, the police officer said that, when she gave the order to storm the building, she didn't know that the perpetrator would be in there.

I take it that, even if Sally were to use one of these sentences to report the police officer, it would still be appropriate for John to form the belief that the police officer didn't know that *P* as a response. (In fact, I think that, intuitively at least, the information that the issue was important for the police officer or that the police officer thoroughly evaluated the evidence should give John even more reason to accept Sally's testimony.) However, the sentences strongly suggest that the stakes or the error-possibilities relevant for the police officer were very different from John's corresponding context features. Given that, they plausibly also suggest that the police officer's epistemic standard was different. So, the implicature that the police officer's epistemic standard was the same as John's should be canceled. Thus, the implicature proposal cannot explain John's reactions to these sentences.

Sally Presupposes that O equals J

One may think that at least the last problem does not arise when we assume that Sally conveys that *O* equals *J* not in terms of a conversational implicature but in terms of a presupposition. After all, presuppositions aren't cancelable.¹⁴ The presupposition proposal, however, leads to other problems. Baker (2012) collects a range of tests for the presence of a presupposition. The alleged presupposition that *J* equals *O* fails all these tests. Let's briefly discuss them in turn.

First, there is the "Hey, wait a minute!" test. According to this test, if " π is presupposed by *S*, then it makes sense for an audience previously unaware of π to respond to an utterance of *S* by saying 'Hey, wait a minute, I didn't know that π .'" (Baker 2012: 117f)
Example,

¹⁴ See e.g. (López de Sa 2008) for such a proposal on behalf of the taste contextualist as a response to the disagreement objection.

A: The mathematician who proved Goldbach's Conjecture is a woman.

B: Hey, wait a minute. I had no idea that someone proved Goldbach's Conjecture.

It seems clear that our present putative presupposition fails that test. For, the following dialogue wouldn't seem to make sense,

Sally: The police officer said that, when she gave the order to storm the building, she didn't know that the perpetrator would be in there.

#John: Hey, wait a minute. I had no idea that

- the police officer's epistemic standard was the same as ours.
- the issue was as important for her as it is for us now.
- she didn't consider potential mistakes we are not considering right now.

Second, there is the "... and what's more ..." test. According to this test, if "S generates a presupposition that π , then it is infelicitous to follow S up with '... and what's more, π ,' or 'and what's more, S*' where S* contains no more information than the conjunction of S and π ." (Baker 2012: 118) Example,

#A: Sarah regrets smoking last night. And what's more, she smoked last night.

Again, the putative presupposition that J equals O fails that test. It seems perfectly felicitous for Sally to say,

Sally: The police officer said that when she gave the order to storm the building she didn't know that the perpetrator would be in there. And what's more,

- her epistemic standard was the same as ours.
- the issue was as important for her as it is for us now.
- she didn't consider potential mistakes we are not considering right now.¹⁵

¹⁵ Admittedly, there may well be *something* puzzling about Sally's above claims in that I wouldn't be surprised if ordinary speakers just couldn't make sense of why Sally would want to communicate that, say, the issue was as important for the police officer as it is for her and John. They might well not see the relevance of this claim. This, however, is clearly not the kind of infelicity we are looking for. The above utterance by A is certainly not infelicitous because one fails to see the relevance of the "what's more" claim. Rather, something (which is clearly relevant) is presented "as if it were new information" (Baker 2012: 119) even though it has been presupposed already in the foregoing sentence.

Third, there is the awkward cancellation test. According to this test, if “I assert S and presuppose π , then an immediate denial of π makes my collective utterance puzzling.” (Baker 2012: 119) Example,

#A: John has stopped leaving towels on the floor, although he never left towels on the floor.

But compare,

Sally: The police officer said that when she gave the order to storm the building she didn’t know that the perpetrator would be in there. But

- her epistemic standard wasn’t the same as ours.
- the issue was more important for her than it is for us now.
- she considered potential mistakes we are not considering right now.

These utterances don’t seem particularly puzzling. So, again, our presupposition fails the awkward cancellation test.¹⁶ (Given that conventional implicatures are not cancellable either, this last problem also arises if we hold that *Sally* conveys that O equals J in terms of a conventional implicature.)

Sally’s Report Shifts the Epistemic Standard

Here is the final potential response to the effect that John is in a position to justify the belief that J equals O . One might want to argue that, as soon as a knowledge claim is reported, the epistemic standard of the reporting context adapts (or at least may adapt) to the epistemic standard of the reported context. This would mean that, in virtue of *Sally*’s report, John’s epistemic standard, J , changes and becomes identical to the police officer’s standard, O . One could now say that John will be in a position to justify the belief that J equals O because he is aware of this context-shifting mechanism.¹⁷

¹⁶ Again, *Sally*’s claims may be somewhat puzzling because the relevance of the denial of the putative presupposition may remain unclear. As before, though, this is clearly not the kind of puzzlement involved in cases where a presupposition is denied. In these cases, we are certainly not puzzled because we don’t see the relevance of the denial of the presupposition but rather because there seems to be a tension between the denial of the presupposition and what has been said before.

¹⁷ Cohen (2005b: 203) suggests a similar proposal on behalf of the epistemic contextualist as a response to the speech report objection.

The difference between this proposal and the proposals above can be brought out as follows. The proposals above have it that John's epistemic standard is the same before and after Sally's report. Before and after the report, the epistemic standard is very undemanding. They go on to argue that, through Sally's report, John comes to be in a position to justify the belief that the police officer's epistemic standard also was very undemanding. Thus, they conclude, John is in a position to justify the belief that his and the police officer's epistemic standard are the same. The present proposal, however, says the following. John's epistemic standard is not the same before and after Sally's report. Before Sally's report, it is very undemanding. After Sally's report, it is identical to the epistemic standard of the police officer, whatever the police officer's epistemic standard was. Correspondingly, John does not come to be in a position to justify the belief that his and the police officer's epistemic standard are the same because he learns that the police officer's epistemic standard was as undemanding as his. Instead, he comes to be in a position to justify that belief because he is aware of the fact that Sally's report has changed his epistemic standard in such a way that it equals the police officer's epistemic standard, whatever this standard was.

There is an immediate problem with this proposal. The proposal allows us to claim that John is in a position to justify the belief that his epistemic standard (after Sally's report) is the same as the police officer's. Correspondingly, it allows the relativist to say that it would be appropriate for John to form the belief that the police officer didn't know that *P* after hearing Sally's report. However, it implausibly entails that John should abandon this belief as soon as he moves on to a conversation where no context-shifting occurred and the epistemic standard is undemanding. Let me elaborate.

Even given the above proposal, John has no clue whether the police officer's epistemic standard was demanding or undemanding (as indicated in the second to last paragraph). So, John would not be in a position to justify the belief that it is true relative to an undemanding epistemic standard that the police officer didn't know that *P*. (For all that John knows, the police officer may fail to know only relative to a very demanding standard.) Hence, given the relativistic justification norm, John must abandon the belief that the police officer didn't know that *P* once he finds himself in a context where the epistemic standard is undemanding (and no context-shifting occurs). This, however,

seems like a very implausible prediction. Suppose, for example, that, subsequent to the discussion with Sally, John joins the neighboring table where a discussion is raging about the police officer's conduct. In this situation, it would seem perfectly fine for John to think to himself, "The police officer didn't even know that *P*!" Relatedly, John would not start worrying that the police officer's epistemic standard may have differed from the epistemic standard that governs his new conversation.¹⁸

2.6. Concluding Remarks

To conclude, let me raise a very general challenge for the relativist who wants to adopt one of the above responses to my objection in spite of the problems just pointed out. We have seen that all of the above responses mimic strategies that contextualists have employed to respond to the challenges from speech reports and disagreement. So, if the relativist defends and employs these strategies herself, she will thereby at least partly undermine her position *vis-à-vis* the contextualist, who should now be allowed to employ these strategies as well. To what extent the relativist position is undermined will depend on the particular strategy she adopts and how much weight must be put on the data that she thereby grants to the contextualist. It may be an open question whether such a discussion would yield problematic results. However, I take it to be on the relativist to show that it would not. Less cautiously put, the (ITD) is not a random datum the relativist fails to predict. Rather, the (ITD) seems to be just another manifestation of the invariantist pull (already familiar in terms of disagreement and speech report data) that contextualism allegedly fails to accommodate and relativism is designed to explain. Thus, insofar as the relativist starts piling on the epicycles to deal with the (ITD), the suspicion arises that she hasn't gotten to the heart of the problem in the first place. Relativism begins to seem like a treatment of only some of the symptoms of a much deeper rooted disease. If we want to take seriously the pull towards invariantism, it seems, we should go for genuinely invariantistic positions, according to which not only the proposition expressed by knowledge claims remains constant but also their truth-value. We will turn to such views in the following chapters. Before that, however, let me pick up some loose ends.

¹⁸ Bach (2005: 61) suggests that there is a similar worry for the contextualist who wants to employ the context-shifting strategy to explain speech reports.

Appendix 1: Scope Limitations

In this appendix, I want to briefly indicate why I present the above objection as an objection to epistemic relativism but neither to other forms of relativism (such as relativism about taste) nor to epistemic contextualism.

I present the objection as an objection to *epistemic* relativism because I don't think it applies in other domains where relativism has been proposed. In particular, I don't think we obtain data like the (ITD) in these other domains. Consider, for example, the case of taste. Taste relativists claim that sentences like "This is tasty" express propositions that have truth-values only relative to the taste of a given speaker (or a given group of speakers). Thus, analogous to the knowledge case above, taste relativists must hold that it would be inappropriate for John to come to believe that the food at the Indian place is tasty when Sally says, "My sister said the food there is tasty." But, as indicated already, this prediction is borne out precisely by the data: It seems to be a plain fact that John cannot appropriately form that belief. Similar considerations may hold for other domains where relativism has been applied. Consider, for example, the case of ethics. I take it that it happens only very rarely that we come to believe that a certain action is good merely because someone (is said to have) said that the action is good. Thus, again, moral relativism would make the right predictions. It is unclear to me whether similar observations carry over to all other domains where relativism has been proposed. The given examples, however, should suffice to make my point.

Let's now turn to the question of why I present the objection only as an objection to epistemic *relativism* but not contextualism. There are two related reasons for this. First, even given that my objection applies to relativism, it seems quite difficult to evaluate whether it also applies to contextualism. For, we would first have to know how the contextualist treats knowledge sentences when embedded under "said that" and "believes that" operators. For example, the contextualist would have to tell us which proposition Sally is supposed to express when she reports, "The police officer said that she didn't know that *P*." One possibility would be that she expresses the proposition that the police officer said that she didn't know *relative to Sally's epistemic standard* that *P*. Another possibility would be that she expresses the proposition that the police officer said that she didn't know *relative to the police officer's epistemic standard* that *P*. Similar problems

arise when we try to understand the (ITD), according to which it would be appropriate for John to form the belief that the police officer didn't know that *P*. Again, the contextualist would have to tell us how she thinks the knowledge sentence behaves once embedded under the "believes that" operator. These questions are difficult to answer.¹⁹ Still, it should be apparent that they have to be answered before we can evaluate whether the above objection applies to contextualism. Second, I am not at all sure whether it would be an interesting result if the above objection did apply to epistemic contextualism. For, I am not at all sure that an objection along these lines would add any further complications to the already familiar objections from disagreement and speech reports.²⁰

Appendix 2: The Case Design

In this appendix, I will briefly explain the design of the *Knowledge Testimony Case*. So far, I have only stated that the following three aspects of the *Knowledge Testimony Case* are crucial for my case against relativism. First, the *Knowledge Testimony Case* involves *indirect* as opposed to direct testimony. Second, what the police officer said is a *negated* knowledge claim. Third, Sally and John find themselves in a *low standards* situation where nothing is at stake and no error-possibilities are salient. In what follows, I will explain why these aspects are crucial.

Consider, first, the fact that the *Knowledge Testimony Case* involves indirect as opposed to direct testimony. The reason for this is that relativists may be in a position to explain cases of direct testimony. To see this, let's modify the *Knowledge Testimony Case* so that it involves direct testimony. In particular, let's assume that Sally had herself endorsed the police officer's claim by saying something along the following lines, "There was some evidence that the perpetrator would be in the building. But, when she gave the order to

¹⁹ For further discussion of contextualism and the problem of embedded knowledge sentences, see e.g. (Cappelen and Hawthorne 2009: Ch. 2).

²⁰ I should note here that, as argued in (MacFarlane 2005a), subject-sensitive invariantism may also get into trouble once it is embedded into a theory of testimony. The problem described in the aforementioned paper, however, seems to be different from the problem discussed in the present chapter. In particular, I don't think that we can derive an objection to subject-sensitive invariantism on the basis of the (ITD). The basic reason for this is that the proponent of subject-sensitive invariantism need not relativize the justification norm discussed above in any way. She can just accept the standard formulation we started out with (or, maybe more plausibly, the first modified version, where we acknowledged a certain variability in the level of justification required for appropriate belief). Given this norm, we get a straightforward account of the (ITD) once we accept the plausible principle that, at least when our interlocutors are sincere and well-informed, testimony to the effect that *p* justifies us (at least to the degree relevant for John in the *Knowledge Testimony Case*) in believing that *p*.

storm the building, the police officer didn't know that the perpetrator would be in there." As before, I take it, we obtain the (ITD): It would be appropriate for John to form the belief that the police officer didn't know that *P* in response to Sally's claim. (This should be even clearer than in the case of indirect testimony.) This time, however, it is not at all clear whether the relativist cannot explain this datum. For example, it may be possible for the relativist to explain the datum as follows. Given that John considers Sally well-informed, it seems reasonable for him to assume that she didn't make a mistake when she said that the police officer didn't know that *P*. Thus, it seems reasonable for John to assume that it is true relative to Sally's epistemic standard that the police officer didn't know that *P*. It does not trivially follow that it is also true relative to John's epistemic standard that the police officer didn't know that *P*. But it is at least an option for the relativist to adopt what DeRose calls a "single scoreboard semantics." On this view, "there is a single scoreboard in a given conversation" and the epistemic standard of each participant of the conversation is determined by the epistemic standard registered on this conversational scoreboard.²¹ Given this view, Sally and John necessarily have the same epistemic standard because they participate in the same conversation. Thus, John can infer that it is true relative to his epistemic standard that the police officer didn't know that *P* from his arguably justified belief that it is true relative to Sally's epistemic standard that the police officer didn't know that *P*. Hence, there is at least a case to be made that, even according to relativism, John can happily form the belief that the police officer didn't know that *P*.

Admittedly, the single scoreboard view is not entirely uncontroversial.²² So, it is not entirely clear whether the relativist can adopt the above response strategy after all. However, this is not the place to adjudicate on the single scoreboard view. So, I will just grant this view to the relativist and, thus, grant as well that the relativist can explain cases of direct testimony. For those who reject the single scoreboard view, the problem for the relativist discussed in this chapter *may* arise already in the case of direct testimony. Whether the problem in fact arises will presumably depend on what exactly one proposes as the alternative to the single scoreboard view. But this is a topic for another occasion.

²¹ See (DeRose 2009: 135).

²² See e.g. (DeRose 2009: Ch. 4) for further discussion.

Let's now turn to the question of why Sally is reporting a negated knowledge claim and why John and Sally find themselves in a low standards context. The *Knowledge Testimony Case* is designed in this way because, again, I don't think the above objection goes through if we change the scenario in one of these respects. More specifically, some of these changes will make the (ITD) disappear, others leave the (ITD) intact but allow the relativist to explain this datum. Let me elaborate.

In the original *Knowledge Testimony Case*, John finds himself in a context with very low standards. No error-possibilities and no heightened stakes are mentioned. The situation of the police officer, on the other hand, is open. Her context may well have been more demanding. Furthermore, Sally reports a knowledge denial. In this constellation, it seems natural for John to accept Sally's indirect testimony and, thus, we obtain the (ITD). But let's turn the situation around. Suppose John finds himself in a high standards situation and Sally reports a knowledge ascription. Consider, for example, the following scenario. John is scrutinizing the evidence that led to the storming of the building. There is a photo where you can see the shape of a person in one of the windows. John considers the possibility that the shape is just a cardboard cutout and concludes that the police officer didn't know that the perpetrator would be in the building. Now reporter Sally comes along. She says, "You know, I was just talking to the police officer in the pub. She said that she knew that the perpetrator would be in the building." It seems quite clear that, in response to this contention, John will not form the belief that the police officer did know that the perpetrator would be in the building. Instead, he would most likely react by asking whether the police officer considered the possibility that the shape in the window was just a cardboard cutout. In this case, then, we don't obtain a datum like the (ITD). Hence, relativism seems to make the right predictions.

So far, we have considered a case where John finds himself in a low standards situation and Sally reports a knowledge denial and a case where John finds himself in a high standards context and Sally reports a knowledge ascription. What if we let Sally report a knowledge ascription in the low standards case and a knowledge denial in the high standards case? I take it that, in both cases, John could happily come to believe that the police officer knew or didn't know that *P* respectively (other than in the high standards case considered in the previous paragraph). Thus, in both cases, we would obtain a datum

like the (ITD). But I also think that these data are unproblematic for the relativist. Let's consider the low standards case in detail.

Suppose that John finds himself in a low standards context and Sally reports, "The police officer said that she knew that the perpetrator would be in the building." As I said, I take it that John could happily come to believe that the police officer knew that P in response to this claim. However, the relativist has a straightforward explanation for that. It goes as follows. Analogous to the original *Knowledge Testimony Case*, John is in a position to justify the belief that the police officer knew- O that P on the basis of Sally's report that the police officer said that she knew that P . Given that John's epistemic standard, J , is assumed to be very low, the police officer's standard, O , presumably was equally or more demanding than his. Correspondingly, knowledge- O presumably entails knowledge- J . Thus, John can derive that the police officer knew- J that P from his already justified belief that the police officer knew- O that P . Hence, even given relativism, it is perfectly fine for John to form the belief that police officer knew that P in the present scenario. Analogous considerations hold for the high standards case where a knowledge denial is reported. So, again, these cases cannot be used to get the above objection going. The situation is summarized in the tables below.

<i>data</i>	low	high
ascription	✓	✗
denial	✓	✓

Table 1: Data concerning the question of whether John can appropriately form the belief that the police officer knew/didn't know that P in response to Sally's respective indirect testimony depending on whether John finds himself in a low or a high standards situation.

<i>predictions</i>	low	high
ascription	✓	✗
denial	✗	✓

Table 2: Relativist predictions concerning the question of whether John can appropriately form the belief that the police officer knew/didn't know that P in response to Sally's respective indirect testimony depending on whether John finds himself in a low or a high standards situation.

3. Moderate Pragmatic Invariantism

Remember the bank cases and the corresponding Contextualist Data. In the bank cases, Hannah and Sarah wonder whether Bill knows that the bank will be open on the next Saturday. Bill has just informed Hannah on the phone that he was at the bank two weeks ago on a Saturday and that the bank was open then. Now the question is whether Hannah can truly ascribe or deny knowledge to Bill. The Contextualist Data say that the answer to this question depends on Hannah's context. In the high standards cases (where much is at stake and Sarah mentions the error-possibility that the bank has changed its hours) Hannah can truly say "Bill doesn't know that the bank will be open." In the low standards cases (where not much is at stake and no error-possibility is mentioned) Hannah can truly say "Bill knows that the bank will be open." The respective negations seem false.

In the previous chapter, I have rejected the relativist account of these data. In the next two chapters, I will discuss and reject the pragmatic invariantist account. Just like the relativist, a pragmatic invariantist endorses invariantism to explain the Invariantist Data (the data about disagreement and speech reports). However, to account for the Contextualist Data, she appeals to pragmatically conveyed contents, in particular, conversational implicatures, instead of relative truth-values.¹ The basic idea is that we evaluate the knowledge claims in the bank cases differently depending on whether they occur in the low or the high standards cases because they *implicate* different propositions in these respective contexts. Let me spell out this view in a little more detail.

Pragmatic invariantism comes in different guises depending on how demanding the knowledge relation invariantly expressed by "knows" is considered to be. In this chapter, I will be concerned with what I will call *moderate pragmatic invariantism* (MPI). According to this view, "knows" invariably expresses an undemanding relation such that

¹ In principle, pragmatic invariantists could appeal to other kinds of pragmatically conveyed contents (such as conventional implicatures or semantic presuppositions). So far as I can see, such views haven't been defended. Still, they are criticized and rejected in (Leite 2005: Sec. 2.2.1), (Fantl and McGrath 2007: 585ff) and (Blome-Tillmann 2013b: Sec. 6). Another potential form of pragmatic invariantism that will not be the direct target of my discussion is the view defended in (Pritchard 2010). This view seeks to explain the contextual variation of knowledge claims not on the basis of implicatures but on the basis of context-sensitive norms of assertion. Some of the subsequent objections may carry over to this version of pragmatic invariantism. But I will not address this issue here.

the knowledge ascriptions in the bank cases turn out to express a truth, while the knowledge denials turn out to express a falsity. For the proponent of MPI, it is relatively straightforward to explain the data about the low standards cases, that is, why the knowledge ascription in the low standards cases seems true, while the denial seems false. For, these data just reflect the truth-values of the propositions the respective knowledge claims express according to MPI. It is a challenge for the proponent of MPI, however, to explain the data about the high standards cases. For, when confronted with these scenarios, we tend to say that the knowledge ascription is false and the knowledge denial is true, contrary to the truth-value of the propositions allegedly expressed by these claims. A proponent of MPI endorses the following response to this challenge. The knowledge ascription in the high standards cases seems false because it *implicates* a falsity. Correspondingly, the knowledge denial in the high standards cases seems true because it *implicates* a truth. What is the content of these respective implicatures? We need not settle on a specific answer to this question for the purposes at hand. But one candidate idea would be that the knowledge ascription in the high standards cases implicates the falsity that Bill can rule out the salient alternative that the bank has changed its hours, while the knowledge denial in the high standards cases implicates the truth that Bill cannot rule out this alternative. The position is summarized in the table below.²

<i>MPI</i>	what is said	what is implicated
“Bill knows that the bank will be open.” (as uttered in the high standards cases)	Bill knows that the bank will be open.	Bill can rule out the salient alternative that the bank has changed its hours (e.g.).
	true	false
“Bill doesn’t know that the bank will be open.” (as uttered in the high standards cases)	Bill doesn’t know that the bank will be open.	Bill cannot rule out the salient alternative that the bank has changed its hours (e.g.).
	false	true

Table 3: A schematic depiction of the commitments of MPI. To be precise, the table presents only what I will call “the semantic/pragmatic thesis” of MPI. See below.

² The position is defended in e.g. (Rysiew 2001, 2005, 2007), (Brown 2006), (Black 2008a, 2008b) and (Hazlett 2009).

In the present chapter, I will argue that MPI, as just defined, should be rejected. To structure the discussion, it is useful to split up MPI into the following two subtheses. On the one hand, MPI entails that the knowledge ascription in the high standards cases expresses a truth but implicates a falsity and, correspondingly, that the knowledge denial in the high standards cases expresses a falsity but implicates a truth (see Table 3). Let's refer to this claim as *the semantic/pragmatic thesis*. This thesis by itself does not yield an account of the data about the high standards cases. It yields such an account only in combination with a very general claim to the effect that truth-value judgments often track the truth-value of the implicatures of the claim in question. More specifically, proponents of MPI must endorse a thesis along the following lines. There is at least a subclass of implicatures to which the alleged implicatures in the bank cases belong such that, if a claim has a false implicature of this sort, then the claim seems false (even if it expresses a truth) and, if a claim has true implicature of this sort, then the claim seems true (even if it expresses a falsity). Let's refer to this principle as *the bridging thesis* (because it bridges the gap between the semantic/pragmatic thesis and the data about the high standards cases). Now, the structure of the present chapter is as follows.

First, I will introduce a distinction between two kinds of implicatures, namely, the distinction between what I will call *additive* and *substitutional* implicatures. Second, I will argue that, if the proponent of MPI appeals to additive implicatures in her account of the data about the high standards cases, then the corresponding bridging thesis turns out to be mistaken. Third, I will argue that, if the proponent of MPI appeals to substitutional implicatures in her account of the data about the high standards cases, then the corresponding semantic/pragmatic thesis turns out to be mistaken (though the relevant bridging thesis may be correct). The basic idea in this latter context will be that the implicatures posited by the proponent of MPI are not properly calculable if they are construed as substitutional implicatures. The upshot of these considerations is that, whether the proponent of MPI appeals to additive or to substitutional implicatures, at least one of the claims that make up her view turns out to be mistaken. Given that the distinction between substitutional and additive implicatures is exhaustive, we can thus conclude that MPI is mistaken.

3.1. Conversational Implicatures: Additive and Substitutional

In this section, I will briefly elaborate on the notion of a conversational implicature to then introduce the distinction between additive and substitutional implicatures. The distinction should be relatively uncontroversial. It is a distinction, however, that is rarely made explicit. So, let me spend some paragraphs on laying it out.³

To begin with, let us get clear about what conversational implicatures, for short, just implicatures, are. I will not attempt to provide a general definition here. Instead, let me state three principles that should fairly uncontroversially apply, if not to implicatures in general, then at least to those implicatures that are supposed to play a role in the bank cases.⁴ First, the conversational implicatures of a given utterance are contents the speaker wants her audience to come to believe by making the utterance, for short, they are contents the speaker wants to convey (or, even shorter, conveys). Second, conversational implicatures are distinct from what is said (and other semantically encoded information such as conventional implicatures and semantic presuppositions).⁵ Third (and relatedly), conversational implicatures are calculable, that is, roughly, it must be possible for ones audience to “calculate” or “work out” their presence on the basis of the assumption that the speaker cooperatively engages in the conversation at hand (more on this below).⁶ For example, suppose I utter “I am sick” in response to your question of whether I will come to your party. I will thereby conversationally implicate that I will not come to your party. Correspondingly, this is something I want to convey without saying it. And my audience will normally be able to work out that this is something I want to convey because, very roughly, they will realize that, in the present context, my contribution makes most sense on this assumption. So much for the general notion of an implicature.

³ The labels “additive” and “substitutional” are taken from (Meibauer 2009: 374), where these terms are used in passing to mark a distinction that corresponds to the distinction I want to draw in this section. Bach’s (1994) distinction between what he calls “implicatures” and “non-literal contents” also seems to correspond to my distinction between additive and substitutional implicatures.

⁴ For a more complete discussion of the notion of a conversational implicature, see e.g. (Grice 1989), (Levinson 1983) and (Blome-Tillmann 2013a). All of the subsequent examples are taken from one or another of these publications.

⁵ It is a consequence of this, I take it, that, at least in general, implicatures can be canceled. Cancelability, however, will not play a role in the discussion to come. So, I don’t include it in the above list.

⁶ Note that Davis (2007: 413), a proponent of a skeptical version of pragmatic invariantism, rejects the last assumption. However, note also that, first, calculability won’t play a role in my case against skeptical versions of pragmatic invariantism. And, second, that Davis’ concerns are far from universally shared. See e.g. (Blome-Tillmann 2013a: 176f) for what seems to be a promising response to these concerns.

Here is the distinction between additive and substitutional implicatures. A given utterance has an *additive* implicature if it says one thing, implicates another and is used to convey both of these claims. For example, suppose I utter “I am sick” in response to your question whether I will come to your party, thereby saying that I am sick and implicating that I will not come to your party. This is a case of an additive implicature. For, I will most naturally intend to convey both that I am sick and that I will not come to your party by making my assertion; I will most naturally want you to come to believe both these propositions. Similarly, when I respond to your question “Where does Frankie live?” by saying “Somewhere in the south of France,” I will say that Frankie lives somewhere in the south of France, implicate that I have no more specific information and most naturally convey both propositions at once (assuming that more specific information would be relevant in the conversation at hand).

A given utterance has a *substitutional* implicature if it says one thing, implicates another but conveys solely the implicature. Suppose, for example, that I utter “The plane was a mile long.” What I say here is that the relevant plane was literally a mile long. What I most likely convey, however, is only that the plane was huge (I do not want you to come to believe that the plane was literally a mile long by making my assertion). So, this is a case of a substitutional implicature. Similarly, suppose I say “Either John will come or he won’t.” What I say here is that John will or will not come, what I most likely convey, however, is only that, say, there’s nothing we can do about John’s absence. As before, I do not want you to come to believe that John will or will not come by making my assertion. (To the contrary, I will most likely presume that you antecedently believe that this is so and, thus, are able to figure out that I want to convey something more informative.)⁷

This is the distinction between additive and substitutional implicatures. Given this distinction, we can ask: Are the implicatures posited by the proponent of MPI supposed to be additive or substitutional? In what follows, I will argue that, on either assumption, MPI must be rejected. In particular, I will argue, first, that, if we assume that the

⁷ I should note that, here, I don’t follow Grice (1989: 87) in his assumption that everything that is said must also be conveyed (or, in Gricean terms, “meant”). For Grice, cases of substitutional implicatures would be cases where speakers only “make as if to say” (1989: 34). It seems very common, though, not to follow Grice in this respect. See e.g. (Bach 1994: 143).

implicatures in question are additive, then the bridging thesis cannot be sustained (according to which, roughly, truth-value judgments are often guided by the truth-value of implicatures). Second, I will argue that, if we assume that the implicatures in question are substitutional, then the semantic/pragmatic thesis cannot be sustained (according to which, roughly, Hannah's knowledge claims in the bank cases have the relevant implicatures). The basic reason for this latter claim will be that the implicatures posited by MPI cannot be calculated in the way in which substitutional implicatures are calculated.

3.2. Additive Implicatures and Seeming Truth-Values

Let's assume that the implicatures appealed to by the proponent of MPI are supposed to be additive. In what follows, I will argue that, given this assumption, the bridging thesis entailed by MPI cannot be sustained. Recall what the bridging thesis amounts to: There is at least a subclass of implicatures to which the alleged implicatures in the bank cases belong such that, if a claim has a false implicature of this sort, then it seems false and, if a claim has true implicature of this sort, then it seems true. First, I will discuss the first part of the bridging thesis (according to which, roughly, a claim can seem false because it has a false implicature). I will argue that, when construed in terms of additive implicatures, this thesis may not be clearly wrong. Still, it is far from clear that it is true. Second, I will discuss the second part of the bridging thesis (according to which, roughly, a claim can seem true just because it has a true implicature). I will argue that this thesis is clearly wrong when construed in terms of additive implicatures.

Before we start, let me note that many authors just assume the first part of the bridging thesis. Even DeRose says that "a true assertion will be inappropriate, *and may seem false*, if it generates a false implicature" (2009: 114, my emphasis).⁸ But why should we believe that this is so? Consider the above example again where you ask me whether I will come to the party. Suppose I am indeed sick but plan to (and will) come to the party anyway. Will not my claim "I am sick" seem true even if it now has the false additive implicature that I won't come to your party?

⁸ For similar claims, see (Rysiew 2001: 496, 2007: 648), (Bach 2002, 2005: 62n) and (Schaffer 2004b: 146).

Alan Hazlett seeks to respond to this challenge by presenting the following two cases.

You know Katie didn't drive drunk last night. Slanderous Sammy is telling everyone: "Katie got drunk and drove home." "That's a lie!", you say, and it does seem to you that Sammy is speaking falsely. (Indeed, you think to yourself: Katie didn't get drunk and drive home! She drove home and got drunk.) But in fact it's true that Katie got drunk and drove home, just not in that order.

You pull into a small town, your tank almost empty, and ask a local if there is a gas station nearby. "There's a gas station around the corner" he says, and directs you to it. But when you get there it turns out the gas station has been closed for years. You think to yourself: that local lied to me!

He then goes on to say about these cases:

You intuit that Sammy and the local spoke falsely, but they didn't—their utterances were merely infelicitous. (Hazlett 2009: 610)

Let's focus on the case of Slanderous Sammy. I am inclined to agree that, in a given conversation, we would be prone to respond to the claim "Katie got drunk and drove home" by saying "That's false! (It was the other way round.)" But does this suffice to establish the bridging thesis under discussion?

I don't think so. The bridging thesis under discussion says, roughly, that claims can seem false when they have a false implicature. How should we understand the idea that a given claim "seems false" in this context? We cannot understand it as the thesis that, *in the context of a given conversation*, we would react to the claim by saying "That's false!" Instead, we must understand it as the thesis that we would judge the claim to be false *if the claim was presented in the context of a scenario description that settles all the relevant facts*. After all, the proponent of MPI is trying to explain the Contextualist Data. And these are data not about how we react to knowledge claims in a given conversation but when confronted with these claims in the context of a scenario description that (supposedly) settles all the relevant facts. So, Hazlett's example is a good one only if we would judge Slanderous Sammy's claim to be false, not only when confronted with it in the context of a given conversation, but also when it is embedded into a case description that settles all the relevant fact. This does not seem to be the case, however. Consider the following case description.

Katie drove home. Then she got drunk. Slanderous Sammy is telling everyone, "Katie got drunk and drove home."

Suppose now that we asked whether Slanderous Sammy's claim is true. Grice at least claims, after discussing a range of similar examples, that "most people would, I think, on reflection have a more or less strong inclination to say that to apply the word or phrase [in question] would be to say something *true* [...], however *misleading* it would be to apply the word or phrase thus." (Grice 1989: 9)⁹ I fully agree with Grice on this point. If I am correct, Hazlett's example cannot be used to underwrite the bridging principle required for the truth of MPI.¹⁰

Now, it may well be that Hazlett would just deny my (and Grice's) verdicts. And, if so, it is not entirely clear how to adjudicate who's correct. However, it seems at least plausible that our dispute ultimately comes down to the question of what competent speakers in general would say about cases like the one just described. To answer this question, it seems reasonable to consult empirical studies. The results of such studies, however, are squarely on my side. Larson et al. (2009a) conducted experiments just like the one just proposed for the conjunction case above. But only about 25% of their subjects judged the conjunction to be false. So, at least in general, people do not seem to share Hazlett's judgments. Consequently, it would seem dubious at least to use his case to underwrite the bridging thesis under discussion.¹¹

⁹ See (Blome-Tillmann 2013b: 4313n) for the same verdict.

¹⁰ One may wonder *why* our verdicts in a conversation differ from our verdicts about claims embedded in a case description. I do not have a specific theory to offer. To me, it just seems clear *that* there is such a difference. And this should suffice to make my point. Still, let me indicate two possible starting points for an account of why there is a difference. One idea might be that, by stating the facts that make the relevant claim literally true in the case description, we somehow prime readers on the literal content. Thus, they evaluate the literal content when asked for the truth-value of the claim in question. In a conversation, no such priming occurs. Thus, speakers may be more liberal in their evaluations. Another idea could be that a use of "That's false" in a conversation may be acceptable even though it is literally false because it can be used to convey that the claim in question only conveyed something false. For example, if it is obvious that the claim in question is true, the speaker can count on the hearers to figure out that a use of "That's false" is only meant to convey that the claim in question conveyed something false. When people are presented with a case description, however, they will, I take it, assume that precision is required. Thus, they will say "That's false" about a given claim only if, strictly speaking, the claim in question is false. Both of these account sketches would have to be spelled out in much more detail to be properly evaluable. This, however, is a topic for another occasion.

¹¹ Let me describe the just quoted study in some more detail. The precise question the participants of the study were asked was "Given the FACT, the underlined sentence is true/false?" In our conjunction case, the "FACT" would be the sentences stating that Katie first drove home and then got drunk. The "underlined sentence" in our case would be Slanderous Sammy's claim, "Katie got drunk and drove home." See (Larson et al. 2009a: Sec. 4.3.3). To be precise, this was only one of three related questions the subjects were asked. Another question they were asked was whether the sentence was "literally" true or false. Yet another question they were asked was whether the fictitious character "Literal Lucy" would say the sentence is true

Maybe there are other cases. For example, it has been suggested that, in many cases of so-called scalar implicatures, we intuitively judge claims to be false just because of their false implicatures.¹² Here are some examples that have been used (where “+>” precedes the relevant implicatures).

Gradable Adjectives: *A* likes *B*. +> *A* doesn't love *B*.¹³

Quantifiers: Some *Fs* are *Gs*. +> Not all *Fs* are *Gs*.¹⁴

Cardinals: *A* has two rackets. +> *A* doesn't have more than two rackets.¹⁵

In all these cases, it is claimed, we may intuitively judge the relevant claim to be false even if only its implicature is false. Again, I grant that we may judge these claims to be false in the context of a given conversation. The question at issue, however, is whether we would judge them to be false once they are embedded into a case description of the relevant sort. As before the results from (Larson et al. 2009a: Sec. 5.2) are extremely helpful when it comes to this latter question. Only 12% of their subjects judged the claims involving gradable adjectives to be false in the context of a corresponding case description. Similarly, only 33% judged the quantifier claims to be false. But,

or false. (Literal Lucy is previously described to them as, roughly, a person who tends to take everything too literally.) See (Larson et al. 2009a: Sec. 4.3.1 and 4.3.2). The 25% mentioned above are the median percentage of people who said that the sentence in the conjunction case was false across all question types just described. See (Larson et al. 2009a: Sec. 5.2). Unfortunately, Larson et al. do not provide the numbers only for the first question type (which would be of primary interest here). But their Figure 2 indicates that the differences between the question types were relatively small.

It may be worried that, by asking about the truth-value of a given *sentence* (as opposed to, say, a claim, an utterance, etc.) the subjects were specifically directed to evaluate only the literal content of the claims in question. It would be no wonder, then, that their evaluations are unaffected by implicatures. To obtain data relevant for the explanation of the Contextualist Data, it might be argued, we would have to ask the subjects whether they consider the *claim* (or the utterance, etc.) in question to be true. First of all, though, I see no reason to accept this thesis. I would be surprised if the empirical results appealed to above would be relevantly altered if we changed the question in the just proposed way. (Philosophers may (reasonably) draw distinctions between sentences, claims, utterances, assertions, etc. I don't think, however, that ordinary speakers do, at least as long as they are not specifically primed to heed these distinctions.) Second, even if the thesis is granted, it will be of no avail to the proponent of MPI. For, at least the Contextualist Data should remain unaltered if we asked whether the *sentences* Hannah uses in the bank cases are true or false. So, we could still conclude that the sentence based Contextualist Data cannot be explained in terms of implicatures. Thus, it would still follow that MPI is bound to fail.

¹² See e.g. (Levinson 1983: 132ff) for a partly introductory discussion of the relevant scalar implicatures. Levinson (1983: 127) also endorses the thesis that these scalar implicatures “will be hard to distinguish from the *semantic* content of linguistic expressions, because such implicatures will be routinely associated with the relevant expression in all ordinary contexts.”

¹³ See (Fantl and McGrath 2009: 42).

¹⁴ See (Davis 2007: 430).

¹⁵ Rysiew (2001: FN 32) refers to (Bach 1987: 78) where this example is used.

interestingly enough, 75% of their subjects said that the claims involving cardinals are false. So, implicatures of the cardinals type seem to be such that, if a claim has a false implicature of this sort, then that claim seems false in the relevant way.¹⁶

Does this result suffice to underwrite the bridging thesis under discussion? Not quite. Remember that, to be precise, this principle has it that there is a relevant class of implicatures *to which the alleged implicatures in the bank cases belong* such that, if a claim has a false implicature of this sort, then it seems false. To fully substantiate this thesis, the proponent of MPI still has to show that the alleged implicatures in the bank cases are of the cardinals type and not of any of the other types of implicatures where truth-value judgments are (more or less) unaffected by the truth-value of the implicature. As Fantl and McGrath (2009: 42) put it, “it is not adequate merely to remark that we do sometimes confuse or run together the semantic content and the propositions pragmatically implied. Why do we do it *here* rather than *there*?”¹⁷ It seems very much unclear, however, whether a plausible answer to this question is forthcoming. So, we should consider it an open question at best whether additive implicatures can be used to sustain the bridging principle under discussion.¹⁸

Let’s turn to the second part of the bridging principle that says, roughly, that claims can seem true just because they have a true implicature. We do not need to consider examples to see that additive implicature cannot be used to underwrite this principle. There is a systematic reason that they can’t. In a nutshell, a true additive implicatures can never make a literally false claim seem true because adding a truth to a falsity will lead to just another falsity. Let me elaborate. Consider the bridging thesis in its most general form, according to which the seeming truth-value of a given claim can be determined by the truth-value of its implicatures. The initial plausibility of this claim is certainly grounded

¹⁶ See (Larson et al. 2009b) for related results.

¹⁷ See also (Blome-Tillmann 2013b: 4313n).

¹⁸ It is often held that the implicatures of knowledge claims will be part of the intuitive truth-conditions of the claims in question because they are “generalized” or “standardized” or “default” implicatures, that is, roughly, because they are implicatures that are *normally* carried by a given sentence. See e.g. (Rysiew 2001), (Leite 2005: 110) and the Levinson quote in FN 12. But this is not sufficient. For, all of the above mentioned implicatures are highly generalized and, even so, most of them are not generally incorporated into the intuitive truth-conditions of the relevant claims. Besides, it is not at all clear why we should think that the implicatures of the knowledge claims in the bank cases are generalized implicatures. After all, they are supposed to vary whenever the epistemic standard changes. And it does seem that epistemic standards change quite a lot.

in the following idea. When we evaluate whether a given claim is true, we often evaluate not whether the proposition this claim *expresses* is true, but whether the proposition this claim *conveys* is true. The truth-value of the proposition that a given claim conveys, however, is at least partly determined by the implicatures of this claim (for they are part of the conveyed content). Thus, the truth-value of an implicature will often shape the truth-value we assign to a given claim. Given that this is the idea that grounds the bridging thesis, it follows directly that a claim that expresses a falsity will never seem true merely because it has a true additive implicature. To see this, consider some claim that expresses a falsity but additively implicates a truth. Suppose now that, as the above account has it, we evaluate whether the claim is true by evaluating whether the proposition it conveys is true. If so, the claim should definitely seem false (not true, as would be required). For, given that the claim has an additive implicature, the conveyed content of the claim will be the combination of the implicature and the proposition expressed. And, given that the proposition expressed is false, the combination of the implicature and the proposition expressed will also be false (whether or not the implicature is true). We thus have a very general reason to think that additive implicatures can never be used to underwrite the idea that a claim can seem true just because it has a true implicature.

In sum, at least the second part of the bridging thesis required for the truth of MPI cannot be sustained in terms of additive implicatures. Whether the first part can be sustained is an open question at best. What if we appeal to substitutional implicatures? Can these implicatures be used to underwrite the bridging thesis? Before we turn to this question, let me briefly discuss a potential rejoinder to the above considerations on the part of the proponent of MPI who wants to stick to additive implicatures.

Denying the Data

Instead of trying to substantiate the bridging thesis by providing examples of additive implicatures that shape our truth-value intuitions, the proponent of MPI could also just reject the troubling data about the high standards bank cases. In the present dialectical context, we are just assuming that these data obtain. So, we need not bother too much about this response strategy. Still, it should be addressed briefly, just to make clear that it does not provide an altogether easy way out for MPI.

At first sight, it may seem that it does. First, the proponent of MPI could say that the knowledge ascription in the high standards cases doesn't seem false but only inappropriate. This putative datum can easily be explained in terms of additive implicatures. For, even a true claim seems inappropriate if it has a false additive implicature. For example, it will be inappropriate to say "I am sick" in response to your question whether I will come to your party if I plan to come in spite of my sickness.

The same strategy will not be available in the case of the seeming truth of Hannah's knowledge denial. For, true additive implicatures do not only fail to explain why a claim seems true. They also cannot explain why a claim seems appropriate. As DeRose puts it, "we want to avoid falsehood both in what we implicate and (especially!) in what we actually say" (2009: 114). (This is true, of course, only for additive implicatures.) Suppose, for example, that you ask me whether I will come to your party. Suppose further that I plan not to come even though I am in perfectly good health. In this case, it will certainly be inappropriate to respond to your question by falsely saying "I am sick" even though this claim has the true additive implicature that I won't come.¹⁹

However, second, the proponent of MPI may even be happy to say that Hannah's knowledge denial in the high standards cases seems neither true nor appropriate. In fact, this is what many proponents of MPI do. Pritchard (2010: 88f), for example, writes that "in skeptical conversational contexts it seems inappropriate and false *either* to assert that an agent has knowledge *or* to assert that the agent lacks knowledge".²⁰ If that is so, additive MPI yields a perfectly plausible account of the data: The knowledge ascription in the high standards cases seems inappropriate because it has a false additive implicature. The knowledge denial in the high standards cases seems inappropriate because it is literally false.

¹⁹ For this objection see also (MacFarlane 2005b: 208) and (Fantl and McGrath 2009: 41f). The objection goes through only if we assume that the above utterance of "I am sick" is not an instance of what has been called a (potentially appropriate) "white lie." For a critical discussion of the idea of explaining the bank case data on the basis of white lies, see e.g. (Fantl and McGrath 2009: 41f).

²⁰ See also (Leite 2005: 123f), (Black 2005: 336) and (Hazlett 2009: 616).

As before, it seems difficult to criticize these responses because we just seem to have one intuitive verdict against the other. But, also as before, it may be useful to refer to the results of pertinent empirical studies.²¹

Hansen and Chemla (2013) asked subjects about the truth-value of the knowledge ascriptions and denials in the (first-person) bank cases (and some related scenarios). The subjects “were given the flexibility to provide their answers within a continuous range of options between false and true, by setting the right end of a red line between these two extreme anchors” (2013: 297f). The results are shown in the subsequent figure. (The “false” anchor was to the left, the “true” anchor to the right.)

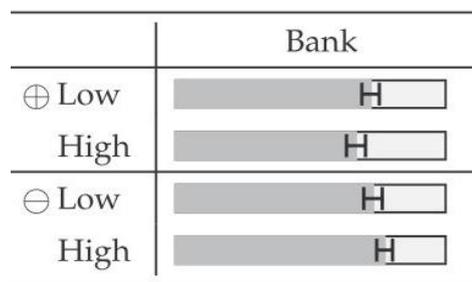


Figure 2: Mean responses to the bank cases.
Taken from (Hansen and Chemla 2013: 305).

As indicated in the figure, Hansen and Chemla found the following not particularly strong but still statistically significant effect:

The Contextualist Data_{empirical}:

The knowledge ascription in the low standards cases is, on average, judged as more true than the knowledge ascription in the high standards cases.

²¹ Blome-Tillmann (2013b: 4307) thinks that the above data denial response is bound to fail independently of empirical results. He holds that “Black and Pritchard [still need] an explanation for why the remaining theorists in the field” judge the knowledge denial in the high standards cases to be appropriate. However, this objection seems dubious to me. Insofar as we demand an account of the judgments of those philosophers who judge the knowledge denial in the high standards cases to be appropriate, we also need to acknowledge the necessity to provide an account of the judgments of those philosophers (namely, Leite, Black, Hazlett and Pritchard) who deny these data. It seems implausible to me, however, that any general semantic or pragmatic theory could explain both these judgments. After all, they contradict each other. To my mind, the only plausible response to such contradicting judgments is this. The judgments of the relevant philosophers on this particular occasion are skewed presumably because they have dealt with the relevant cases for so long that they can no longer say what’s appropriate and what isn’t. Thus, their judgments should be ignored. Instead, we should appeal to empirical results as to what competent and philosophically unspoiled speakers would say about the relevant cases. This is what I will be doing in what follows.

The knowledge denial in the high standards cases is, on average, judged as more true than the knowledge denial in the low standards cases.²²

Admittedly, it is relatively unclear what bearing these results should have on the plausibility of the Contextualist Data. For, when I speak of a claim as being judged as “more true” than another, all I want to say is that responses to the first claim were to the right of responses to the latter on the given truth scale. And it remains to be seen how best to interpret such outcomes. For example, ordinary speakers may indeed have a gradable notion of truth such that one of the claims is literally more true than the other. But it may also be that they just find one claim more appropriate than the other or that their confidence levels vary.²³ So, it is unclear whether the above results can be used to settle our dispute about whether the knowledge ascription in the high standards cases seems false or only inappropriate and whether the knowledge denial seems true or also inappropriate.

Independently of their relation to the Contextualist Data, though, the Contextualist Data_e present a problem for the proponent of MPI who appeals to additive implicatures. For, it is at least unclear whether the bridging thesis required to explain the Contextualist Data_e in terms of additive implicature can be upheld. Admittedly, it may well be possible to use additive implicatures to explain why the knowledge ascription in the low standards cases seems more true than the knowledge ascription in the high standards cases. For, as indicated already, the empirical results from (Larson et al. 2009a) suggest that it is easy to find examples where at least some people incorporate additive implicatures into the intuitive truth-conditions of the claims in question. So, assuming that the knowledge ascription in the high standards cases has a false additive implicature while the knowledge ascription in the low standards cases doesn't, the former should, on average, be judged as less true than the latter. However, the situation is far less clear when it comes to the datum that the knowledge denial in the high standards cases is judged as more true than the knowledge denial in the low standards cases. For, this time, the knowledge denial in the high standards cases has a *true* implicature while the knowledge denial in the low standards cases doesn't. And, even if this true implicature is incorporated into the intuitive

²² See (Hansen and Chemla 2013: 305). These data are confirmed in (Schaffer and Knobe 2012) and (Buckwalter 2014).

²³ See (Hansen and Chemla 2013: Sec. 5.2.1) for further discussion.

truth-conditions of the knowledge denial, it is unclear why the knowledge denial should thereby be evaluated as more true. For example, I cannot make a false claim more appropriate by merely adding a truth. Similarly, I won't feel more or less confident that a given claim is false once it is conjoined with an independent truth. The ultimate verdict on this issue will depend on how exactly we should interpret the notion of a claim being more true than another. The burden of proof, however, seems to be on those who want to put forward an interpretation according to which the Contextualist Data_e can be explained in terms of MPI.

So, even if we grant that the data about the high standards cases can be denied, the empirically confirmed Contextualist Data_e still seem troubling enough for the proponent of MPI, at least as long as she only appeals to additive implicatures.

3.3. Substitutional Implicatures and Seeming Truth-Values

Let's now turn to substitutional implicatures. Can these implicatures be used to underwrite the bridging thesis? Recall, again, what the bridging thesis amounts to. The bridging thesis says the following. There is at least a subclass of implicatures to which the alleged implicatures in the bank cases belong such that, if a claim has a false implicature of this sort, then it seems false, and if a claim has true implicature of this sort, then it seems true. We have seen that it is an open question at best whether the first part of the bridging thesis can be sustained in terms of additive implicatures. We have also seen that, for very general reasons, the second part of the bridging thesis cannot be sustained in terms of additive implicatures. So, in what follows, I will evaluate whether the bridging thesis can be sustained in terms of substitutional implicatures.

Let's start with the second part of the bridging principle, according to which, roughly, a claim can seem true because it has a true implicature. First of all, it should be noted that the principled problem we found for additive implicatures no longer applies when we turn to substitutional implicatures. The reason is this. In a nutshell, the problem with additive implicatures was that the overall conveyed content of a claim with an additive implicature contains the proposition expressed. As a result, the overall conveyed content is false if the proposition expressed is false, whatever the truth-value of the additive implicature. The same problem does not arise when we turn to substitutional implicatures. For, the

overall conveyed content of a claim with a substitutional implicature contains only the implicature. So, the conveyed content in these cases may be true even if the proposition expressed is false.

These considerations show that one particular argument against the idea that substitutional implicatures can make a literally false claim seem true does not apply. Clearly, though, this result does not suffice to establish the bridging principle under discussion. As before, we need to be given examples of cases where a claim seems true just because it has a true substitutional implicature. So, in what follows, I will, again, go through the examples that have been used in the literature to show that they are not convincing.

Tim Black offers an example of a substitutional implicature where truth-value judgments allegedly track the implicature. He says that

my assertion 'Joe Lieberman is a Republican' is false, but in the right sort of context it can seem true. We can explain the seeming truth of my assertion by appealing to the fact that it has a true implicature, something like this: Lieberman, *qua* legislator, often thinks or acts like a Republican. (Black 2008a: 193)

In a similar vein, Davis holds that “everyday metaphors” provide cases “in which what a speaker says appears true just because of a true implicature.” (2007: 413) An example Davis uses is this. “If my wife asks whether I made caffeinated or decaffeinated coffee, I might answer ‘It’s high-test.’” (2007: 411) As Davis notes, what his answer literally says is false “since coffee has no octane rating.” Still, he would hold, the claim may seem true because it has the true substitutional implicature that the coffee is caffeinated.

I am skeptical about Davis’ and Black’s contentions. Consider Black’s case. It may be that, in some conversational contexts, we would be prone to respond to the claim “Joe Lieberman is a Republican” by saying something along the lines of “That’s true! (He acts like a Republican.)” But, as in the case of additive implicatures, these are not exactly the data we need. We need claims that will be judged as true not only in a given conversation but also once embedded into a scenario description that settles all the relevant facts. So, suppose we presented subjects with a scenario description along the following lines.

Joe Lieberman is a member of the Democratic Party. However, he often thinks or acts like a Republican. [A description of “the right sort of context.”] Now A says, “Joe Lieberman is a Republican.”

Suppose further that we ask whether A’s claim is true. Independently of how we specify “the right sort of context,” I find it highly dubious that one would respond by saying that A’s claim is true. To the contrary, given that Liebermann is described as a member of the Democratic Party, A’s claim should seem false. So, it seems that Black’s case cannot be used to underwrite the bridging principle under discussion.

As before, Black might just deny my verdict. So, we would have one intuitive verdict against the other. But, also as before, empirical results seem to favor my view, as the subsequent discussion will show.

Though this is not the ultimate goal of the paper, (Keysar 1989: 378f) reports studies in which subjects were asked to evaluate the literal truth-value of metaphorical sentences after reading “context sections” containing information that either make the sentence metaphorically true or metaphorically false (whether the sentence is literally true is left open). For example, subjects were asked to evaluate the sentence “Bob Jones is a magician” in either of the following contexts,

Metaphorically True (M+) Sometimes it seems as if Bob’s money is made of rubber because he stretches it so far. How does he create such a healthy profit despite these expenses?

Metaphorically False (M-) Although Bob tries to budget carefully, it seems that to him that money just disappears into thin air. With such huge audiences, why doesn’t he ever break even?

Here, M+ is supposed to make “Bob Jones is a magician” metaphorically true because Bob is described as a person performing deeds that resemble magic tricks. Correspondingly, M- is supposed to make “Bob Jones is a magician” metaphorically false because there is nothing magical about Bob’s described behavior.²⁴ The subjects provided their responses on a seven-point scale for literal truth. It turns out that “the test sentences did not differ in ‘literal truth’ ratings with M+ and M- context sections”. In both cases,

²⁴ Admittedly, the just discussed context sections may not be the best way to indicate that the sentence in question is metaphorically true or false. For example, the reference of “these expenses” in M+ remains unclear. Correspondingly, it remains unclear whose “audiences” are being talked about in M-. However, Keysar tested not only the sentence just discussed but 11 other metaphorical sentences as well. He doesn’t provide the corresponding context sections. But it seems reasonable to assume that they do not generally suffer from similar defects.

they “were rated near the midpoint for literal truth”. Keysar concludes that “metaphorically related information did not affect the literal truth-value of test sentences.”

These results thus suggest that, contrary to Black, people would judge *A*’s claim in the Lieberman case above as literally false in spite of its true substitutional implicature (at least in the absence of arguments to the conclusion that Black’s case relevantly differs from cases of metaphor). So, Black’s example fails to provide a model for the bridging thesis. (The same goes for Davis’ example. The situation here is even clearer because there is no doubt that Davis’ case involves metaphor.)

One might want to object to my appeal to the above empirical study in the following way. The bridging thesis is a thesis about truth-value judgments. Keysar’s study, however, tested only how we evaluate the *literal* truth-value of a given claim. Thus, the study has no bearing on the question at hand.

This is an objection that would have to be investigated further. As things stand, however, it doesn’t seem particularly strong. First, the results in (Larson et al. 2009a: Sec. 5.1) suggest that, at least for the case of many additive implicatures, truth-value judgments and literal truth-value judgments coincide (even when there is a divergence between the truth-value of what is said and the truth-value of what is overall conveyed). Thus, we may reasonably expect that the results from Keysar’s study would be the same if we asked about the truth-value not the literal truth-value of the respective claims. Second, insofar as one is prone to accept that the knowledge denial in the high standards bank cases sounds true, one would most likely also be prone to accept that it sounds *literally* true. To explain this latter judgment, we need a bridging principle involving literal truth-values. And Keysar’s study shows at least that this latter principle cannot be sustained by means of Black’s and Davis’ examples. This seems problematic enough.

There are two possible response strategies to the above considerations on the part of the proponent of MPI who wants to appeal to substitutional implicatures. The first is, again, to deny the data about the high standards cases. The second is to try to provide further more convincing examples. I will discuss these strategies in turn. It turns out that the first

strategy founders when it comes to the empirically confirmed Contextualist Data_e. The second strategy, however, may succeed if we grant the existence of so-called implicatures.

Denying the Data

As before, let me briefly discuss the idea of just rejecting the data about the high standards cases. The proponent of MPI could argue that the knowledge denial in the high standards cases doesn't seem true but only appropriate and, similarly, that the knowledge ascription in the high standards cases doesn't seem false but only inappropriate.²⁵ These data could easily be explained on the basis of substitutional implicatures. For, on the one hand, it is inappropriate to say things with false substitutional implicatures (whether or not the literal content is true). For example, we wouldn't say "The plane was a mile long" if the plane wasn't even huge. And, on the other hand, it is perfectly appropriate to make a claim with a true substitutional implicature even if its literal content is false. For example, I could happily say "The plane was a mile long" even if the plane was not literally a mile long.

As before, it is not entirely clear how to settle a dispute where some hold that a claim seems intuitively true or false and others accept only that the claim seems appropriate or inappropriate. But, as before, it may be useful to refer to the results of pertinent empirical studies.

So, remember the Contextualist Data_e, according to which the knowledge ascription in the low standards cases is judged as more true than the knowledge ascription in the high standards cases and, correspondingly, the knowledge denial in the high standards cases is judged as more true than the knowledge denial in the low standards cases. As indicated above, it is relatively unclear what bearing these results should have on the plausibility of the Contextualist Data because it is unclear how to interpret the result that some claims are evaluated as "more true" than others. However, the Contextualist Data_e seem problematic for the proponent of substitutional MPI independently of their precise interpretation. Here is why. The results from Keysar's study seemed to show that substitutional implicatures have no effect on truth-value ratings when these ratings are

²⁵ See e.g. (Davis 2010: 1154). DeRose also suggests that data about appropriateness might rest on firmer grounds than truth-value data. As regards the knowledge ascription in the low standards cases and the knowledge denial in the high standards cases, he writes: "As we can sense, perhaps with even more certainty than with which we judge the truth-values of the claims, speakers do in fact use 'know(s)' in the way described, and appropriately so" (DeRose 2009: 50).

provided on a seven-point scale. Consequently, we may reasonably expect that substitutional implicatures should also show no effect when these responses are provided on a continuous scale. So, if, as the present version of MPI has it, the knowledge claims in the low and the high standards bank cases differ only with respect to their substitutional implicatures, then there should be no difference between their truth-value ratings on a continuous scale. According to the Contextualist Data_e, however, there are such (albeit small) differences. So, substitutional implicatures cannot be used to explain these data (whatever the correct interpretation of these data may be). Consequently, there seems to be a discrepancy between empirically confirmed data and the predictive capacities of MPI when construed in terms of substitutional implicatures.

Implicatures

Given that the strategy of denying the data on the part of the proponent of substitutional MPI doesn't seem very promising, the only hope she has can be to provide examples of implicatures that underwrite the bridging thesis after all. And, in fact, there is a slightly controversial but arguably not too controversial kind of substitutional implicatures that might do the trick, namely, so-called "implicatures."²⁶ In this section, I will spell out briefly what implicatures are and why implicatures, if they exist, can be used to underwrite the bridging thesis.

Implicatures may be considered a kind of implicature in that they are calculable conveyed contents that go beyond what is literally said.²⁷ What distinguishes implicatures from other implicatures is that implicatures are

a matter of leaving words out. [...] In implic-i-ture, one says something but does not mean that; rather, what one means includes an implicit qualification on what one says, something that one could have made explicit but did not. (Bach 2001: 252)

Here are some examples. The typical "implicit qualifications" or "left out words" appear in curly brackets.

You are not going to die {from that cut}.

²⁶ Bach coined the term "implicature" in his (1994). For appeals to implicatures on the part of proponents of MPI, see e.g. (Rysiew 2005: 62n, 2007: 629) and (Brown 2006: 416ff).

²⁷ Bach would not be entirely happy with the classification of implicatures as a kind of implicature. See e.g. (Bach 1994: 140). But at least for those implicatures that interest us here (namely, expansion implicatures, see below), he would, I take it, agree that they are calculable conveyed contents. This is the crucial assumption I am making here. The rest should be merely terminological issues.

I have eaten {recently}.

It is raining {here}.

Steel isn't strong enough {to carry this construction}.

Someone {on the course} solved the riddle.

Everyone {on the course} solved the riddle.

The final sentence, for example, supposedly says that everyone (in the whole world) solved the riddle. It implicates, however, that only everybody on the course solved the riddle.²⁸

There are two kinds of implicatures only one of which can be employed by the proponent of MPI. The kind of implicature that can be employed by the proponent of MPI involves what Bach calls “expansion.” An *expansion* implicature of a given utterance is an implicature of that utterance such that the literal content of the utterance already is a complete proposition. As Bach puts it, “there is already a complete proposition, something capable of being true or false” (1994: 133). The quantifier case is an example of this. The kind of implicature that cannot be employed by the proponent of MPI involves what Bach calls “completion.” A *completion* implicature of a given utterance is an implicature of that utterance such that the literal content of the utterance is “semantically underdeterminate.” An example is the steel case. According to Bach, “something must be added for the sentence to express a complete and determinate proposition (something capable of being true or false). [W]e need to know strong enough for what.” (1994: 127)²⁹ These latter implicatures cannot be employed by the proponent of MPI for the following reason. According to MPI, “Bill knows that the bank will be open” expresses a truth in the bank cases. Thus, it expresses something that is “capable of being true or false.” Hence, by definition, utterances of this sentence cannot have completion implicatures.

²⁸ These examples, particularly the last, straightforwardly indicate what's controversial about implicatures. In all the above cases, one might want to hold that the alleged implicated content is already part of what is said. Thus, this content would not count as an implicature according to the above general characterization of this notion (that has it that implicatures go beyond what is said). See e.g. (Stanley 2007). This is also the reason why e.g. Lewis (1996) and Ichikawa (2011) consider their views forms of contextualism even though they model the alleged context-sensitivity of “knows” on the basis of quantifier domain restriction. Resolving this issue, though, goes way beyond the scope of this project. So, I will just grant the idea of an implicature to the proponent of MPI and proceed from there.

²⁹ For more on this, see (Bach 2006b: 440f).

Since completion implicatures will thus play no role in what follows, I will refer to expansion implicatures simply as “implicatures.”³⁰

Now, if we allow for the existence of implicatures, it is very easy to find cases of substitutional implicatures that underwrite the bridging thesis. For example, when the teacher says to the class “Everyone solved the riddle,” this claim may well seem true even though the proposition expressed (that everyone in the whole world solved the riddle) will definitely be false. Similarly, when the teacher says to the class “No one solved the riddle,” the claim may easily seem true even if someone on an earlier course did solve the riddle. As Bach (2002: 30) claims, “intuitions tend [...] to be responsive to implicit qualifications”.

So, assuming implicatures, the only question that remains to be answered by the proponent of MPI is why the knowledge claims in the bank cases should be conceived of as cases of implicatures rather than cases of other kinds of implicatures where truth-value judgments do not track the implicated content in the relevant way. It may be possible to answer this question as follows. The proponent of MPI could say, for example, that the implicated content of Hannah’s knowledge claims in the high standards cases is that Bill knows or doesn’t know *for sure* that the bank will be open (where Bill knows for sure that the bank will be open only if he would be in a position to rule out the salient possibility that the bank has changed its hours).³¹ If so, the implicatures would be “a matter of leaving words out” (the left out words being “for sure”). Since, implicatures are characterized by being a matter of leaving words out, it thus follows that the bank case implicatures are implicatures.

³⁰ It is interesting to see that MPI is incompatible with completion implicatures only in virtue of its commitment to moderatism, the view that “Bill knows that the bank will be open” expresses a truth in the bank cases. Pragmatic invariantism by itself, however, is not incompatible with completion implicatures. Still, even Bach (2001: 252n) himself agrees that completion implicatures are “more controversial” presumably because it seems initially more plausible to treat the completed material as already part of what is said. In any case, it seems quite unclear that there would be any relevant difference between contextualism and a form of pragmatic invariantism that appeals to completion implicatures. See (Montminy 2007) for further elaboration of this last point.

³¹ For further discussion of the issue of so-called “clarification devices” and many more potential such devices, see (Hawthorne 2004a: 104), (Stanley 2005b: 68–70), (Ludlow 2005: 19–20), (Bach 2005: 61f), (Davis 2007: 418f), (DeRose 2009: 180ff) and (Montminy 2007: 114f, 2013: 2363).

In summary, if the proponent of MPI appeals to substitutional implicatures (in particular, implicatures), there is at least a case to be made that she can sustain the bridging thesis. In what follows, however, I will show that she will no longer be in a position to sustain the semantic/pragmatic thesis. In particular, I will show that Hannah's knowledge claims cannot have the substitutional implicatures our present proponent of MPI assigns to them because these implicatures are not calculable in the way in which substitutional implicatures are supposed to be calculable. Objections from calculability are not new. In fact, such objections have accompanied MPI throughout its entire history.³² Still, I think that the full force of these objection hasn't been appreciated. The basic reason for this is that they are normally presented against the backdrop of an underdeveloped account of how implicatures are calculated in general. In the following section, I will try to remedy this deficiency.

I will start out with a general discussion of how implicatures are calculated. In the course of this discussion, I will distinguish between two ways in which implicatures can be calculated. In particular, I will introduce and defend the distinction between *exploitation* and *observation* implicatures. I will then state a general principle in terms of this distinction, namely, the principle that all substitutional implicatures are exploitation implicatures. On that basis, I conclude that the implicatures posited by the proponent of MPI who appeals to substitutional implicatures must be calculable along exploitative lines. However, as a last step, I will argue that these implicatures aren't calculable along exploitative lines. It follows that substitutional implicatures cannot be used to sustain the semantic/pragmatic thesis required for the truth of MPI.

3.4. Calculating Implicatures

3.4.1. Observation and Exploitation Implicatures

Here is the distinction between exploitation and observation implicatures. To begin with, remember that the presence of an implicature is supposed to be calculable. What does calculability mean? Roughly, it means that the implicature can be derived on the basis of the assumption that the speaker is being cooperative. The notion of cooperativeness here

³² See e.g. (DeRose 2002).

can be spelled out in various ways.³³ In this dissertation, I will follow Grice in assuming that a speaker is *cooperative* roughly if she observes the following “conversational maxims.”

*Quality*₁: Be sincere.

*Quality*₂: Be justified.

*Quantity*₁: Be as informative as is required (for the current purpose of the exchange).

*Quantity*₂: Don’t be more informative than is required.

Relation: Be relevant.

Manner: Be perspicuous.³⁴

In what sense are implicatures supposed to be *derivable* from the assumption that speakers are cooperative (that is, they observe the above maxims)? According to Levinson, these “inferences come about in at least two distinct ways” (1983: 104). On the one hand, implicatures may “arise directly from the assumption that the speaker is observing the maxims.” (1983: 105)³⁵ On the other hand, implicatures may “come about by overtly and blatantly *not* following some maxim, in order to exploit it for communicative purposes.” (1983: 109)³⁶ Grice seems to draw the same distinction. On the one hand, he says, there are cases of implicatures “in which no maxim is violated, or at least in which it is not clear that any maxim is violated.”³⁷ On the other hand, he says, there are cases of implicatures in which, “though some maxim is violated at the level of what is said, the hearer is entitled to assume that that maxim, or at least the overall Cooperative Principle, is observed at the level of what is implicated.” These latter cases, Grice holds, “involve exploitation, that is, a procedure by which a maxim is flouted” (1989: 32f).³⁸ I will refer to these two kinds of implicatures as *observation* and *exploitation* implicatures respectively. The distinction between observation and exploitation implicatures seems

³³ See (Blome-Tillmann 2013a: 176) for a brief discussion of options.

³⁴ See (Grice 1989: 26f). The formulation of the maxims above diverges slightly from Grice’s original formulation. The reason for this should become apparent below.

³⁵ Or, as Levinson also puts it, if the “speaker is observing the maxims in a fairly direct way” (1983: 104). See also (Levinson 1983: 126, 157f).

³⁶ Or, as Levinson also puts it, when “the speaker deliberately and ostentatiously breaches or (as Grice puts it) *flouts* the maxims.” (1983: 104) See also (Levinson 1983: 126, 157f).

³⁷ Or, as he also put it, in which the violation of a maxim is only “apparent.” See e.g. (Grice 1989: 31).

³⁸ For similar distinctions, see (Hugly and Sayward 1979) and (Huang 2007: 27ff).

intuitively very plausible to me. Still, as it stands, it is not particularly clear. So, in what follows, I will spell it out and substantiate it further.

Paradigmatic examples

To get an initial grasp on the distinction, it is useful to consider paradigmatic examples of observation and exploitation implicatures. The following case provides a paradigmatic example of an observation implicature the generation of which is based on the maxim of Relation.³⁹ Someone responds to the question of where to get petrol by saying, “There is a garage around the corner.” She thereby implicates that the garage (most likely) sells petrol. This implicature is standardly taken to be observational. Very roughly, the reason for this is that the implicature can be calculated in the following way. The hearer assumes that the speaker is observing the maxims. She realizes, however, that the speaker would not be observing the maxims—in particular, Relation—if she did not think that the garage (most likely) sells petrol. Thus, the hearer concludes, the speaker implicates that the garage sells petrol. As one can see, the calculation of this implicature is based on the assumption that the speaker is observing the maxims. Thus the implicature is classified as an observation implicature.

The following case provides a paradigmatic example of an exploitation implicature the generation of which is also based on the maxim of Relation.⁴⁰ Someone responds to the question of what to do about John’s absence by saying, “John will come or he won’t.” She thereby implicates that there is nothing one can do about John’s absence. This implicature is standardly taken to be exploitative. Very roughly, the reason for this is that the implicature can be calculated in the following way. The hearer realizes that what the speaker says is a mere logical truth and, thus, completely irrelevant for the purposes at hand. Thus, what the speaker says “overtly and blatantly” fails to comply with the maxims, in particular, Relation. Consequently, the hearer concludes, the speaker must have meant to convey something other than what she said. And the most plausible candidate here is that there is nothing to be done about John’s absence. So this is what the speaker implicates. As one can see, the calculation of this implicature is based on the

³⁹ See (Grice 1989: 32).

⁴⁰ See (Levinson 1983: 111). In fact, Levinson considers the above implicature to be based on Quantity₁. But I think it can just as well be treated in terms of Relation. The subsequent discussion should bring this out.

understanding that the speaker “overtly and blatantly” fails to follow the maxims. Thus it is classified as an exploitation implicature.

Even given these examples, many details concerning the distinction between observation and exploitation implicatures remain obscure. In what follows, I will try to point out these obscurities and remove them to the extent required for the purposes at hand.

Observing the maxims at the level of what is said/conveyed

Here is the first obscurity. Consider the implicature in the case about John’s absence. Given that the case is supposed to involve an exploitation implicature, Levinson has to maintain that the speaker is “overtly and blatantly *not* following some maxim.” However, there is at least a sense in which the speaker does not fail to observe any maxim at all. For, in saying “John will come or he won’t,” the speaker conveys that there is nothing one can do about John’s absence. This latter claim is perfectly relevant, the speaker presumably believes it, etc. So, at least when it comes to what is conveyed, there is no maxim that the speaker fails to observe. Is Levinson just wrong then about his general characterization of exploitation implicatures? How do we have to understand his characterization of exploitation implicatures to get the cases right?

The proper response to this worry should be fairly obvious. Exploitation implicatures must be understood as implicatures where the speaker fails to observe a maxim *at the level of what is said* but not *at the level of what is conveyed*. Let us first get clear about what it means to observe a maxim at the level of what is said as opposed to the level of what is conveyed. As a second step, we will apply the results to the case of John’s absence to show that this case does indeed involve an exploitation implicature on the present understanding.

What does it mean to observe a maxim at the level of what is said? Given the above formulation of the maxims, this is straightforward to see. Consider, for example, Relation. This maxim says that you should be relevant. To observe this maxim at the level of what is said is to *say* (as opposed to merely convey) something which is relevant. Similarly, to observe the quality maxims at the level of what is said is to *say* (as opposed to merely convey) what one believes and is in a position to justify. And to observe the quantity maxims at the level of what is said is to *say* (as opposed to merely convey) something

informative but not too informative. The notion of observing a maxim at the level of what is conveyed can be defined in an analogous fashion.

Before we go on, let's briefly turn to the maxim of Manner. The maxim of Manner is special. For, it is not clear how to make sense of the idea that one should be perspicuous at the level of what is said as opposed to the level of what is conveyed. Rather, it seems that one can be perspicuous only at the level of the sentences used. Consider, for example, the submaxim of Manner that says "Avoid ambiguity."⁴¹ It is quite unclear how one should avoid ambiguity other than at the level of the sentence used. So, I will assume in what follows that no distinction can be drawn between observing Manner at the level of what is said and observing Manner at the level of what is conveyed. When I say that a speaker is observing (all) the conversational maxims at the level of what is said/conveyed, I thus intend to be understood as follows. The speaker observes all maxims *except for Manner* at the level of what is said/conveyed *and she observes Manner* (at whatever level this maxim can be observed). Since Manner implicatures will only play a minor role in the discussion to come, these qualifications can mostly be ignored though.

Given the distinction between observing a maxim at the level of what is said and observing it at the level of what is conveyed, we can see more clearly why the claim "John will come or he won't" generates an exploitation implicature. The reason is this. Exploitation implicatures are generated when the speaker fails to observe the maxims *at the level of what is said*. When one says "John will come or he won't" in a context where the issue is what to do about John's absence, *what one says* is clearly irrelevant. Thus, we have an exploitation implicature in these cases, just as we should.

Observing the maxim of Relation

So much for the first unclarity. Here is a second unclarity I want to remove. Consider observation implicatures. Grice says of observation implicatures that, when they occur, there is "no maxim that is violated." Levinson says that, in these cases, speakers "observe the maxims in a fairly direct way."⁴² Clearly, we must understand these claims as claims

⁴¹ See (Grice 1989: 27) for this submaxim.

⁴² Note here that, as I will understand these terms, "observing a maxim" and "not violating a maxim" can be used interchangeably (where "not observing a maxim" is presumably the term that more stably refers to what I have in mind when I use these terms). I assume, however, that there is a distinction to be drawn

about observing or violating the maxims *at the level of what is said*. For, we have seen that, even in the case of exploitation implicatures, no maxim is violated *at the level of what is conveyed*. Hence, the difference between observation and exploitation implicatures would disappear if we said that, also in the case of observation implicatures, speakers only observe the maxims at the level of what is conveyed. So, observation implicatures are implicatures where the speaker observes the maxims at the level of what is said. But, given that, why should the petrol case above turn out to involve an observation implicature? Let us consider this case in detail. What the speaker says in this case is that there is a garage around the corner. The question at issue is where to get petrol. Now, clearly, there are garages that don't sell petrol. So, clearly, what the speaker says at least does not entail an answer to the question at issue. In a sense, then, what the speaker says is not relevant. Hence, in a sense, the speaker does violate Relation at the level of what is said. Thus, it may seem, the relevant implicature would have to be classified as exploitative rather than observational, contrary to what Grice and Levinson say. How should we deal with this problem?

One potential response would be the following. We could grant that the speaker in the petrol case violates Relation at the level of what is said. In order to maintain the idea that, nevertheless, the case involves an observation implicature, we could further admit that observation implicatures must be defined as implicatures where the speaker violates a maxim at the level of what is said. In order to maintain the distinction between observation and exploitation implicatures, we could then define exploitation implicatures as implicatures where the speaker not only violates but *overtly and blatantly* violates a maxim at the level of what is said.⁴³

This response may seem initially plausible. But it is also problematic in various ways. First, it is in tension with what Grice and Levinson say. To repeat, Grice, for example,

between observing (or not violating) a maxim *in a primary sense* and observing (or not violating) a maxim *in a secondary sense*. In the primary sense, one observes a maxim if, as a matter of fact, one doesn't violate it. For example, one observes Relation at the level of what is said in this primary sense if, as a matter of fact, one says something relevant. In the secondary sense, one observes a maxim already if one *believes* that one doesn't violate it. For example, one observes Relation at the level of what is said in this secondary sense if one *considers* what one says relevant. Both senses will figure in the subsequent discussion. The context should make clear on each occasion which sense I have in mind.

⁴³ Bach (2006a) seems to favor such a proposal when he writes that “[i]mplicatures (and implicatures) arise only when one or another maxim is flouted or violated”.

explicitly says that, in the case of observation implicatures, “no maxim is violated.” Second, it is not quite clear why the speaker in the petrol case should not also be said to overtly and blatantly violate Relation. After all, it does seem fairly obvious that what she says does not entail an answer to the question at hand.⁴⁴

These may not be knock-down arguments. But instead of discussing the above proposal in more detail, let me simply put forward an alternative proposal that I find more appealing overall. The alternative proposal is that observation implicatures are implicatures where the speaker observes the maxims at the level of what is said, while exploitation implicatures are implicatures where the speaker fails to observe the maxims at the level of what is said. This again prompts the worry that, contrary to Grice and Levinson, the petrol case turns out to involve an exploitation implicature, because the speaker seemingly fails to observe Relation at the level of what is said. In order to respond to this worry, I will proceed as follows. I will first clarify the notion of relevance. Second, I will restate Relation in terms of this clarified notion. Finally, I will show that, given this clarified version of Relation, the speaker in the petrol case turns out to be observing Relation at the level of what is said.

To get clear on the notion of relevance, we first have to understand the logical form of this relation. Clearly, relevance is at least a two-place relation. There is always *some claim* that is relevant *for some question* (or some set of questions). Thus we preliminarily arrive at the following form.

Claim *P* is relevant for question *Q*.⁴⁵

This, however, doesn’t seem to be enough. For, whether a given claim is relevant very often depends on our background beliefs. If I, for example, believe that all garages sell petrol, the claim that there is a garage around the corner will be relevant for me regarding the question of where to get petrol. For somebody else who, for whatever reason, believes

⁴⁴ For a further problem, see FN 63 below.

⁴⁵ Note that I am concerned only with the relevance of claims here. Of course, questions can also be relevant or irrelevant. The notion of relevance required to explain this phenomenon will diverge from the one developed below. Since I will not deal with questions in what follows, though, this shouldn’t be much of a problem.

that garages never sell petrol, the same claim will be irrelevant regarding the same question. So, it seems that we need a further argument place along the following lines.

Claim P is relevant for question Q relative to a set of background beliefs B .

Further argument places may be required. Since they will be immaterial to the project at hand, however, I will stick to the logical form above. Thus, we can provide the following rough and ready definition of what it means for a claim to be relevant.

Claim P is relevant for question Q relative to a set of background beliefs B just if P together with B entails (or at least makes likely) a certain answer (or at least a partial answer) to Q .⁴⁶

Whether or not we accept precisely that definition, we can use the three-place relation of relevance to clarify what it means to observe Relation at the level of what is said. Here is my proposal.

A speaker S making an utterance u in a conversational context C *observes Relation at the level of what is said* just if what S says by uttering u is relevant for the question at issue in C relative to the background beliefs that S holds.⁴⁷

With this definition at hand, we can turn back to the question of why the speaker in the petrol case doesn't violate Relation at the level of what is said. The answer should now be clear. According to the present definition, the speaker observes Relation at the level of what is said just if what she says is relevant for the question at issue relative to her background beliefs. Now, what the speaker says in the petrol case is that there is a garage around the corner. The question at issue is where to get petrol. Moreover, the speaker presumably believes that garages generally sell petrol. So, given the above definition, the speaker observes Relation if the claim that there is a garage around the corner, together with the background belief that garages generally sell petrol, makes likely a particular

⁴⁶ Note that I am assuming here that P must not figure vacuously in the relevant entailments. Otherwise one could say something relevant by merely having relevant beliefs.

⁴⁷ To be precise, we would presumably have to impose some kind of restriction on the background beliefs that can figure in the present definition. For, a speaker will presumably not be said to be observing Relation if the background beliefs that make her claim relevant are obviously unacceptable for her audience or too complicated for the audience to grasp in the course of an ongoing conversation. The details of such a restriction, however, will not be relevant for my concerns. So, I leave them for another occasion.

answer to the question of where to get petrol. And clearly, this is the case. The answer that is made likely is that one can get petrol around the corner. So, according to the above definition, the speaker observes Relation at the level of what is said. Hence, we get an observation implicature, just as we should.⁴⁸

Observing one maxim at the cost of another

So far, then, we can accept the idea that observation implicatures are implicatures where the speaker observes the maxims at the level of what is said, while exploitation implicatures are implicatures where the speaker fails to observe the maxims at the level of what is said. But there is one last thing we have to get clear about. Consider the case where someone responds to the question of where Frankie lives by saying “Somewhere in the south of France,” thus implicating that she cannot be more specific. Is this an observation or an exploitation implicature? We may want to say that it is an exploitation implicature. For, the speaker clearly says less than would be required. So, arguably, she violates the quantity maxims at the level of what is said. Such a case, however, is standardly not considered an exploitation implicature.⁴⁹

How do we get this verdict? I think we have to slightly modify our above definitions along the following lines. Exploitation implicatures are implicatures where the speaker not only fails to observe the maxims, but fails to observe the maxims *and could have done*

⁴⁸ One may worry that, given the above definitions, the speaker in the case of John’s absence will also turn out to be observing Relation at the level of what is said. This is not the case, however. Let me elaborate. Assume the above definitions hold. We can then argue as follows. The speaker in the case of John’s absence observes Relation at the level of what is said only if what she says is relevant as regards the question at issue relative to her background beliefs. So, the speaker observes Relation at the level of what is said only if what the speaker says together with certain of her background beliefs entails or makes likely a particular answer to the question at issue. Now, what the speaker says is that John will or will not come (a logical truth). The question at issue is what should be done about John’s absence. So, the speaker observes Relation at the level of what is said only if the logical truth that John will or will not come, together with certain of the speaker’s background beliefs, entails or makes likely a particular answer to the question of what to do about John’s absence. The speaker, however, does not hold background beliefs that would satisfy this condition (assuming that the logical truth must not figure vacuously in the relevant derivations). To see this, consider an example of a background belief that would satisfy the condition. Consider, for example the belief that, whenever John will or will not come, one should call him. This belief together with the logical truth that John will or will not come would entail a particular answer to the question of what to do about John’s absence (namely, that one should call him). However, the speaker certainly does not hold this belief. For, given minimal logical competence on her part, this would mean that she also holds the absurd belief that one should *always* call John. So, according to the definitions above, the speaker in the case of John’s absence does not observe Relation at the level of what is said.

⁴⁹ See (Levinson 1983: 107) and (Grice 1989: 32).

better.⁵⁰ More precisely, exploitation implicatures are implicatures where the speaker fails to observe the maxims at the level of what is said *and* there is a proposition p such that, if the speaker had said that p , she would have observed the maxims at the level of what is said (or would have been closer to observing the maxims at the level of what is said than she actually is)⁵¹. Relatedly, we should redefine observation implicatures as follows. Observation implicatures are implicatures where the speaker *either* observes the maxims *or, if she doesn't, couldn't have done better*. More precisely, observation implicatures are implicatures where *either* the speaker observes the maxims at the level of what is said *or* there is no proposition p such that, if the speaker had said that p , she would have observed the maxims at the level of what is said (or would have been closer to observing the maxims at the level of what is said than she actually is).

With these definitions at hand, we can easily see why the France case above involves an observation rather than an exploitation implicature. The speaker in this case provides less information than is required and thus fails to observe the quantity maxims at the level of what is said. But this is presumably so only because if she were to provide more information by saying, say, that Frankie lives in Cannes, she would say something she doesn't justifiably believe and thus fail to observe the quality maxims at the level of what is said. So, even though the speaker fails to observe a maxim at the level of what is said, she presumably couldn't have done better. There is no proposition p such that, if the speaker had said that p , she would have observed all maxims at the level of what is said. Thus, given the above definitions, we have an observation rather than an exploitation implicature, just as we should.

Now, I think that the above proposal is ultimately the correct way to deal with cases such as the France case. However, it also makes the present framework much more cumbersome and difficult to handle. In particular, we can no longer test whether an

⁵⁰ This is at least one sense, I think, in which speakers in the case of exploitation implicatures not only fail to observe the maxims but “flout” or “exploit” them.

⁵¹ It may not be entirely straightforward to spell out what exactly it is supposed to mean that a speaker could be “closer” to observing the maxims. Here I will just assume an intuitive grasp of this concept. One thing should be noted however. Whether a speaker is close to observing the maxims depends not only on the sum of the distances to each individual maxim (as it were). It depends on a *weighted* sum of these distances. As Grice puts it, it “is obvious that the observance of some of these maxims is a matter of less urgency than is the observance of others” (1989: 27). In particular, the distance to the quality maxims will be weighted more strongly than the distance to the other maxims.

implicature is observational by simply looking at whether the speaker observes all individual maxims. Instead, we now have to check whether she observes all individual maxims and, if she doesn't, whether there is anything she could have said that doesn't violate any maxim. So, to simplify the discussion, I will adopt an alternative solution to the problem of the France case. This solution may not be entirely correct, and it is certainly less encompassing than the solution presented above. Still, it should yield the same results for the cases we are interested in here. The proposal goes as follows.

We stick to the simple idea that observation implicatures are implicatures where the speaker observes the maxims at the level of what is said, while exploitation implicatures are implicatures where the speaker fails to observe the maxims at the level of what is said. In order to deal with the France case, we slightly modify the content of Quantity₁. So far, this maxim says that one should be as informative as is required. Instead, we say that, according to Quantity₁, one should be as informative as is required, *but only to the extent that this is compatible with the quality maxims*. More precisely, we define the observation of Quantity₁ as follows.

A speaker *S* making an utterance *u* in a conversational context *C* *observes Quantity₁ at the level of what is said* just if, from the propositions that *S* justifiably believes, what *S* says in uttering *u* is the most informative proposition regarding the question at issue in *C*.

This proposal allows us to stick to the simple definition of observation and exploitation implicatures, while, at the same time, yielding the right result for the France case. Consider again the France case. It does seem that what the speaker says in this case is the most informative thing she can say without violating the quality maxims. From all the things the speaker justifiably believes, what she actually says is the most informative proposition. Correspondingly, the speaker observes Quantity₁ at the level of what is said in the sense just defined. So, according to the above proposal, we have an observation implicature, just as we should.

Calculating observation and exploitation implicatures

The proposal presented in the last paragraph is almost what I want to accept, but not quite. The reason is this. So far, the definitions of observation and exploitation implicatures

make no reference to any kind of calculation. However, observation and exploitation implicatures were supposed to differ precisely in the way in which they are calculated (according to Levinson at least). So calculation should play a role in their definition. In what follows, I want to remedy this deficiency by developing a definition of observation and exploitation implicatures that is based on the way in which these implicatures are calculated. (The relation to the definitions developed so far should soon become apparent.)

Let's start with observation implicatures. How do we calculate observation implicatures? One may worry that observation implicatures cannot be calculated given what I have said so far. For, on the present view, speakers observe all maxims at the level of what is said when an observation implicature is generated. And it may be held that, if all maxims are observed, there is nothing that could trigger the search for an implicature.⁵² This is not the case however.

Consider the petrol case. How can we calculate the implicature in this case? We start out, I take it, with the assumption that the speaker is observing the maxims at the level of what is said when she says "There is a gas station around the corner." Next, we realize that the speaker can be observing Relation at the level of what is said only if she holds the background belief that the mentioned garage is likely to sell petrol. Since nothing speaks against the idea that the speaker holds that belief, we conclude that the speaker does hold that belief. Consequently, we assume that the speaker wants us to hold that belief as well; that is, that the speaker means to convey that the garage is likely to sell petrol. Hence, we take it to be an implicature of the speaker's claim that the garage is likely to sell petrol. This is a paradigmatic instance of the kind of calculation that occurs in the case of observation implicatures. No maxim is violated. Still, an implicature is generated. Correspondingly, we can define observation implicatures as follows.

An *observation implicature* is an implicature that can be calculated along the lines of the following pattern:

⁵² Remember that Bach (2006a), for example, states that "[i]mplicatures (and implicatures) arise only when one or another maxim is flouted or violated".

The speaker is observing the conversational maxims at the level of what is said. The speaker can be observing the conversational maxims at the level of what is said only if she thinks that *p*. There are no good reasons to hold that the speaker doesn't think that *p*. Thus, the speaker thinks that *p*. Thus, the speaker wants me to think that *p* as well. Thus, the speaker conveys that *p*. Thus, the speaker implicates that *p* (given that *p* differs from what is said).⁵³

Let's now turn to exploitation implicatures. How are these implicatures calculated? Consider the case of John's absence. As before, I take it, we start out with the assumption that the speaker observes the maxims at the level of what is said when she says, "John will come or he won't." We realize, however, that the speaker can be observing the maxims, in particular Relation, at the level of what is said only if she thinks that the logical truth she expresses is relevant for the question at issue. And this is something that the speaker clearly doesn't believe. Thus, we discard the original assumption that the speaker observes the maxims at the level of what is said. Still, we uphold the assumption that the speaker observes the maxims at the level of what is conveyed. Correspondingly, we conclude that what the speaker means to convey differs from what is said.⁵⁴ As a last step, we try to figure what it could be that the speaker wanted to convey (on the basis of the assumptions that it differs from what is said, that it should still have some connection to what is said, that it is supposed to be in line with the maxims, and other contextual cues).

⁵³ This pattern is at least closely related to the pattern Grice (1989: 31) states as a general means to derive implicatures (not only observation implicatures). However, when he later considers examples of implicatures where a maxim is flouted, the pattern he uses seems much closer to the pattern below for exploitation implicatures (see esp. Grice's discussion of irony on p. 34). Hugly and Sayward (1979) also observe that Grice's original pattern cannot be applied to what I would refer to as exploitation implicatures.

⁵⁴ Some may worry that one can never uphold the assumption that the speaker is observing the maxims at the level of what is conveyed once one has given up the assumption that the speaker is observing the maxims at the level of what is said. For, it may be worried that if the speaker doesn't want to convey what she said, she automatically violates Manner. As Wilson (2006: 1727) puts it, "whatever implicature is derived, the resulting interpretation would irrevocably violate the Manner supermaxim ('Be Perspicuous'), since the most straightforward way of conveying this implicated information would have been to express it directly." This reasoning, however, is misguided. It may be that to express something directly is always more perspicuous than to merely convey it as an implicature. The maxim of Manner, however, does not say that you should be as perspicuous as possible. It only enjoins you to be perspicuous. (In the same way, Quality₂ enjoins you to be justified. Certainly, though, it doesn't say that you should be justified to the extent that your justification could not be better.) So, even if saying something directly is more perspicuous than to convey it as an implicature, it doesn't follow that conveying something as an implicature is never perspicuous. If the hearer has no problems figuring out what it was that was supposed to be implicated, the speaker's claim may well be counted as perspicuous even if it would have been even more perspicuous if the speaker had directly said what she in fact only conveys.

In the particular case at hand, this leads us to the conclusion that the speaker wants to convey and thus implicates that there is nothing to be done about John's absence. This is a paradigmatic instance of the type of calculation that occurs in the case of exploitation implicatures. Correspondingly, we can define exploitation implicatures as follows.

An *exploitation implicature* is an implicature that can be calculated along the lines of the following pattern:

The speaker is observing the conversational maxims at the level of what is said. But the speaker can be observing the conversational maxims at the level of what is said only if she thinks that *q*. And there are good reasons to hold that the speaker doesn't think that *q*.⁵⁵ Thus, contrary to the initial assumption, the speaker is not observing the conversational maxims at the level of what is said. Still, the speaker is observing the conversational maxims at the level of what is conveyed. Thus, the speaker must mean to convey something other than what is said. All things considered, the most plausible candidate conveyed content is *p*. Thus the speaker conveys that *p*. Thus the speaker implicates that *p* (or at least that part of *p* that differs from what is said).⁵⁶

Further examples

Before we conclude, let's consider some further examples to substantiate the present interpretation of the distinction between observation and exploitation implicatures. So far,

⁵⁵ What are the candidate reasons to think that a speaker doesn't believe that *q*? I think the answer to this question provides an account of why exploitation implicatures are associated with the idea of "overtly and blatantly" not observing a maxim. So let me say a few words about this issue. One reason to assume that a speaker doesn't believe that *q* is that the speaker said that she doesn't believe that *q* (or said that not-*q*). But the *q* involved in the pattern above most likely isn't a proposition the speaker and the hearer have explicitly discussed already. So, what other reasons are there to think that a speaker doesn't believe that *q*? One reason that will typically figure in the pattern above is that *q* is just an absurd claim that no one would be willing to hold. For example, consider the case of John's absence. Here the hearers realize that they can uphold the assumption that the speaker is observing Relation only if they assume that the speaker believes that a mere logical truth has a bearing on the question of what to do about John's absence. Why don't they conclude that the speaker holds such a belief? Because the content of that belief is so absurd that no one would be willing to accept it. Thus, the reasons to think that the speaker doesn't believe that *q* that figure in the pattern above normally license the claim that the speaker "overtly and blatantly" fails to hold the relevant belief. Correspondingly, they license the claim that the speaker "overtly and blatantly" fails to observe the maxims. That, I think, is why exploitation implicatures are associated with the idea that the speaker "overtly and blatantly" fails to observe a maxim.

⁵⁶ This pattern is at least closely related to the pattern described in (Levinson 1983: 157f) for the calculation of exploitation implicatures.

we have mostly considered implicatures based on the maxim of Relation. Let's now consider implicatures that are based on other maxims. First, consider quality implicatures.

Exploitation implicatures of this kind are easy to find. Take, for example, the case where someone says "Queen Victoria was made of iron", thereby implicating, say, that Queen Victoria was tough. This is standardly considered an exploitation implicature.⁵⁷ Correspondingly, it can be derived as follows. The speaker observes Quality₁ at the level of what is said only if she believes what she said: that Queen Victoria was literally made of iron. Obviously, though, the speaker doesn't hold that belief. Hence, she fails to observe the maxims at the level of what is said. Nevertheless, she observes them at the level of what is conveyed. Thus, the speaker must mean to convey something other than what is said. The most natural candidate claim is that Queen Victoria is tough. So, this claim is implicated.

Observation implicatures based on the quality maxims are more difficult to find. Grice doesn't provide examples of this sort. But Levinson says that, when I say "John has two PhDs," I observationally implicate that "I believe that he has, and have adequate evidence that he has." (1983: 105) Whether these claims are cases of observation implicatures presumably depends on whether we want to classify them as implicatures to begin with.⁵⁸ If they are implicatures, however, they clearly are of the observational type. For, they can be calculated as follows. In saying "John has two PhDs," the speaker is observing the maxims at the level of what is said. She would not be doing so unless she believed that John has two PhDs and that she is justified in holding this belief. There is no reason to think that the speaker doesn't hold these beliefs. Thus, the speaker holds these beliefs. Hence, she wants me to hold them as well. Thus, she implicates that she believes that John has two PhDs and that she is justified in holding this belief.

Let's now turn to quantity implicatures. As an example of an exploitation implicature that is based on the quantity maxims, consider Grice's letter of recommendation case.⁵⁹ In a letter of recommendation for a philosophy job, someone writes "Dear Madame, Ms. X's command of English is excellent, and her attendance at tutorials has been regular. Yours,

⁵⁷ See (Levinson 1983: 110).

⁵⁸ A reason that straightforwardly speaks against this classification is that the alleged implicatures aren't cancellable.

⁵⁹ See (Grice 1989: 33).

etc.” This letter will generate an exploitation implicature to the effect that Ms. X is no good at philosophy. Correspondingly, the implicature can be calculated in the following way. We can uphold the assumption that the writer observes Quantity₁ at the level of what is said only if we assume that she thinks she cannot say more than she says without violating the quality maxims. However, assuming that Ms. X is a pupil of the writer, the writer certainly can say more. Thus, the writer does not observe Quantity₁ at the level of what is said. Hence, we can conclude that the writer must have meant to convey something other than what is said. And a plausible candidate here is that Ms. X is no good at philosophy. This is what is implicated.

Here is an example of an observational quantity implicature. Suppose someone says, “It’s warm,” thus implicating that it is not hot. This implicature is standardly taken to be observational.⁶⁰ Correspondingly, the derivation could go (roughly) as follows. The speaker observes the maxims at the level of what is said. If the speaker observes Quantity₁ at the level of what is said, she must think that it is not hot. For, if she (justifiably) thought it was hot, what she said would have been less informative than it could have been (in particular, the speaker could have made the stronger claim that it is hot instead of the weaker claim that it is warm). Nothing speaks against the idea that the speaker thinks it is not hot. Thus, the speaker thinks it is not hot. Hence, she wants me to think that too. So, she implicates that it is not hot.

Manner implicatures

Let’s finally turn to Manner implicatures. We have seen already that it makes no sense to distinguish between observing Manner at the level of what is said and observing Manner at the level of what is conveyed. It is a consequence of this that Manner implicatures cannot be exploitation implicatures in the sense just defined. In detail, the reason is this. In the case of exploitation implicatures, hearers first realize that the speaker fails to observe the maxims at the level of what is said. Then they fall back to the assumption that the speaker still observes them at the level of what is conveyed (and then they try to figure out what it is that is conveyed). This cannot happen in the case of Manner implicatures, however, if we assume that it makes no sense to distinguish between observing Manner at the level of what is said and observing Manner at the level of what is conveyed. For, if

⁶⁰ See (Levinson 1983: 106).

it makes no sense to distinguish between observing Manner at the level of what is said and observing Manner at the level of what is conveyed, then it also makes no sense to say that a speaker fails to observe Manner at the level of what is said but doesn't fail to observe Manner at the level of what is conveyed. So, Manner implicatures cannot be exploitative.

This conclusion is admittedly at odds with what Grice and Levinson say. For, they both think that there are exploitative Manner implicatures. At this stage, however, I want to take the liberty to say that they were wrong. The distinction they plausibly had in mind when they distinguished exploitation from observation implicatures just fails to allow for exploitative Manner implicatures.

There are two possible conclusions to be drawn from this result. Either all Manner implicatures are observational or the distinction between observation and exploitation implicatures is not exhaustive. I do not want to adjudicate between these options here. Let me only indicate that I am inclined to accept the first. There is at least a case to be made that all Manner implicatures are observational.

To begin with, consider the following example of a Manner implicature that clearly is of the observational type.⁶¹ Someone says, "Katie got drunk and drove home." She thereby implicates that Katie *first* got drunk and *then* drove home. This is an observational Manner implicature that can be calculated in the following way. The speaker is observing the maxims at the level of what is said. In particular, she is observing Manner. If she is observing Manner, she must think that the events she refers to happened in the order in which she refers to them. Thus, she thinks that this is so. Thus, she wants me to think that this is so as well. Hence, she implicates that the events she refers to happened in the order in which she refers to them; that is, Katie *first* got drunk and *then* drove home.

Let us now consider a Manner implicature that Grice and Levinson consider to be exploitative.⁶² In a review of a musical performance, someone writes, "Ms. Singer produced a series of sounds corresponding closely to the score of an aria from *Rigoletto*." The writer implicates that Ms. Singer's performance was awful. As I said, Grice and

⁶¹ See (Levinson 1983: 108).

⁶² See (Levinson 1983: 112) and (Grice 1989: 37).

Levinson hold that this case involves an exploitation implicature. There is a case to be made, however, that the implicature can be calculated along the lines of the pattern for observation implicatures. The calculation could go as follows. The speaker is observing the maxims at the level of what is said. In particular, she is observing Manner. If she is observing Manner, she must think that there was no way to express what she said (or something sufficiently similar to what she said) in a more succinct way without violating some other maxim at some level. In particular, the speaker must think that she would have violated a maxim at some level if she had just said, “Ms. Singer sang an aria from *Rigoletto*.” Nothing speaks against this assumption, however. In particular, the following seems quite likely. The speaker thinks that, if she had just plainly said, “Ms. Singer sang an aria from *Rigoletto*,” then she would have conveyed that there was nothing unusual about the singing, in particular, that the singing wasn’t totally awful. The speaker also thinks, however, that the performance indeed was totally awful. As a consequence, she thinks that she would have violated the quality maxims at some level if she had just plainly said, “Ms. Singer sang an aria from *Rigoletto*.” Thus, among other things, the speaker most likely thinks that Ms. Singer’s performance was awful. Thus, she wants me to think that too. Thus, she implicates that Ms. Singer’s performance was awful.

It seems quite plausible to me that all Manner implicatures can be derived in a similar way (and that the above calculation is in line with the general calculation pattern for observation implicatures). Thus, I am inclined to think that all Manner implicatures are observation implicatures. If it turns out that they are not, however, I am also happy to grant that they are neither observation nor exploitation implicatures but implicatures of a third kind.⁶³

⁶³ One might think that, in light of the fact that my proposal forces me to reclassify some Manner implicatures, it would be better after all to adopt the proposal indicated above, according to which observation implicatures differ from exploitation implicatures in that the first involve the violation of a maxim while the second involve the exploitation or flouting of a maxim. Note, however, that a proposal along these lines also forces us to reclassify some implicatures. Consider, for example, the Manner implicature in the case of Katie’s drunken drive. Levinson plausibly classifies this implicature as observational. However, there is no maxim whatsoever in this example that the speaker could be said to violate (let alone flout or exploit). In this example, the speaker *would* violate Manner *if* she didn’t believe that the events she recounts happened in the order in which she recounts them. But, given that the speaker does hold that belief, Manner is not violated. Thus, according to the proposal just mentioned, we would have to say that Levinson is wrong to classify this implicature as observational. So, even according to this proposal, some implicatures would have to be reclassified.

Principle

We should now have a firm enough grasp on the distinction between observation and exploitation implicatures to come back to the general principle mentioned at the outset of this section.

Principle: All substitutional implicatures are exploitation implicatures.⁶⁴

I take this principle to be justified simply because I fail to see counterexamples to it. However, it may also be possible to support it by general considerations of roughly the following sort. Suppose a given implicature is not an exploitation implicature. In that case, the speaker will have observed all maxims at the level of what is said. Correspondingly, there will be no reason to think that what is said is not supposed to be conveyed. Hence, no substitutional implicature is generated. So, whenever something is not an exploitation implicature, it is also not a substitutional implicature. Thus, by contraposition, whenever something is a substitutional implicature, it is also an exploitation implicature. More would have to be said to fully substantiate this reasoning. But I will leave that for another occasion.⁶⁵

3.4.2. Calculating Substitutional Implicatures in the Bank Cases

We now have a relatively firm grasp on how implicatures are calculated in general. Given that, we can turn back to the implicatures posited by the proponent of substitutional MPI.

⁶⁴ Levinson (1983: 105) similarly claims that “implicatures that arise from the assumption that the speaker is observing the maxims [...] simply amplify the communicated content in restricted ways”. This is to say that all observation implicatures are additive implicatures. Assuming that the distinction between additive and substitutional implicatures and the distinction between observation and exploitation implicatures are exhaustive, Levinson’s thesis is equivalent to Principle.

⁶⁵ Note that the other direction of Principle presumably does not hold, that is, that not all exploitation implicatures are substitutional. See, for example, Grice’s letter of recommendation case (Quantity₁). In this case, we have an exploitation implicature. Still, it may well be that, in addition to what is merely conveyed, the speaker also wants to convey what is literally said. Here is a further example of this sort (taken from (Grice 1989: 35)). Anna says, “Carla is an old bag.” Berta responds, “Huh, lovely weather for March, isn’t it?” Berta thereby implicates that Anna shouldn’t have insulted Carla (or something along these lines). In this case, Berta clearly fails to observe Relation; there is no belief she plausibly holds that would connect what she says to the questions at issue. Thus, we have an exploitation implicature. Nevertheless, we might argue that this implicature is additive. For, it does seem that Berta may want to convey what she said, that the weather is good (in order to direct the conversation towards this new topic). So, here we have an exploitation implicature that is additive. Thus it is not the case that all exploitation implicatures are substitutional. More generally, this last thesis can be motivated as follows. Whenever there is an exploitation implicature, the speaker fails to observe the maxims at the level of what is said. However, the fact that a speaker fails to observe the maxims at the level of what is said only vindicates the conclusion that the speaker cannot want to convey *only* what is said. The speaker may still want to convey what is said and something else in addition. So, an exploitation implicature may well be additive.

Are these implicatures properly calculable? To answer this question, we first have to ask whether the relevant implicatures are supposed to be observational or exploitative. Given Principle, the answer is clear. The proponent of substitutional MPI appeals to substitutional implicatures. Thus, her implicatures must be calculable along the lines of the pattern discerned for exploitation implicatures. In what follows, however, I will show that the implicatures posited by the proponent of substitutional MPI cannot be calculated along the lines of this pattern. (To simplify the terminology, I will from now on use “calculable” in the sense of “calculable along the lines of the pattern discerned for exploitation implicatures.”)

Let’s discuss the alleged implicature of the knowledge denial in the high standards cases first. According to MPI, the knowledge denial in the high standards cases expresses the falsity that Bill doesn’t know that the bank will be open. It supposedly seems true, however, because it has a true implicature along the lines of the claim that Bill doesn’t know *for sure* that the bank will be open. At first sight, it may seem straightforward to calculate the relevant implicature along exploitative lines by appealing to the quality maxims. The idea here would be the following. In saying “Bill doesn’t know that the bank will be open,” Hannah observes the quality maxims at the level of what is said only if she justifiably believes that Bill doesn’t know that the bank will be open. However, according to MPI, Bill does know that the bank will be open and Hannah should be perfectly aware of that (because Bill just informed her about his epistemic position). Thus, Hannah doesn’t believe that Bill doesn’t know that the bank will be open. Hence, she fails to observe the quality maxims at the level of what is said. Thus, she doesn’t want to convey what she says. So, she most likely implicates that, even though Bill knows that the bank will be open, he doesn’t know *for sure* that this is so.

This strategy may seem initially appealing. But I don’t think it will do. The basic reason for this has been pointed out by DeRose (2009: 122f). In a nutshell, his worry is that while *we*, as omniscient observers of the discussion between Hannah and Sarah, may be in a position to calculate the relevant implicature on the basis of the quality maxims, Hannah’s audience, namely, Sarah will not be in a position to do so. Let me elaborate. It seems immensely plausible that, in order for an implicature to be properly calculable, it must be

calculable by the speaker's audience, that is, by Sarah in the case at hand.⁶⁶ Thus, if the above calculation is supposed to go through, there must be good reasons *for Sarah* to reject the assumption that Hannah believes that Bill doesn't know that the bank will be open when she says so. But Sarah has no such reasons. After all, she doesn't know what Bill told Hannah on the phone. In particular, she has no reason to believe that Bill didn't just say that he doesn't know that the bank will be open. Thus, Sarah cannot calculate the relevant implicature on the basis of the quality maxims. Thus, the implicature isn't properly calculable on the basis of these maxims after all.

It is for this reason, I think, that proponents of MPI generally appeal to the maxim of Relation to calculate the implicature in question.⁶⁷ In what follows, I will show that this strategy is also bound to fail. If the alleged implicature of Hannah's knowledge denial in the high standards cases is calculable on the basis of Relation, it should be possible for Hannah to go through a reasoning along the following lines. If Hannah observes Relation at the level of what said, then she has belief *B*. However, there are good reasons to hold that Hannah doesn't have belief *B*. So, Hannah fails to observe Relation at the level of what is said, etc. I will argue that this cannot be done because there is no belief *B* that satisfies this pattern. In other words, I will show that there are no problematic beliefs that we have to ascribe to Hannah in order to uphold the assumption that she observes Relation at the level of what is said.

What beliefs do we have to ascribe to Hannah to uphold the assumption that she observes Relation at the level of what is said? We have seen that, in order to observe Relation at the level of what is said, what Hannah says will have to be relevant for the question at issue relative to Hannah's background assumptions. What beliefs do we have to ascribe to Hannah in order to uphold this assumption? The question at issue in the high standards bank cases presumably is whether Bill not only knows but knows for sure that the bank will be open (because the stakes are high and a far-fetched error-possibility has been mentioned). What Hannah says, according to MPI, is that Bill doesn't know that the bank will be open. So, our question is: Which beliefs do we have to ascribe to Hannah in order

⁶⁶ See e.g. (Saul 2010: Sec. 1.2.2) for a defense of the claim that conversational implicatures must be calculable by one's audience.

⁶⁷ See e.g. (Rysiew 2001: 491, 2007: 638f), (Brown 2006: 424–426), (Hazlett 2009: 606) and (Black 2008a: 196, 199–200).

to uphold the assumption that the claim that Bill doesn't know that the bank will be open is relevant for the question of whether Bill knows for sure that the bank will be open relative to Hannah's background beliefs? It seems that the only background belief we have to ascribe to Hannah to uphold this assumption is the more or less trivial belief that if Bill doesn't know that the bank will be open, then he is also doesn't know for sure that the bank will be open. Clearly, it is entirely unproblematic to ascribe this belief to Hannah. So, there are no implausible beliefs we have to ascribe to Hannah in order to uphold the assumption that she observes Relation at the level of what is said. Thus, the relevant implicatures cannot be calculated on the basis of Relation.⁶⁸

This reasoning might seem too quick. For, one might think that I illegitimately assumed that the question at issue in the bank cases is whether Bill knows for sure that the bank will be open. Alternatively, one could hold that the question at issue is only whether Bill can rule out the salient alternative that the bank has changed its hours.⁶⁹ Even on this assumption, however, there is no problematic belief we have to ascribe to Hannah in order to uphold the assumption that she observes Relation at the level of what is said. Given that the question at issue is whether Bill can rule out the salient alternative, Hannah observes Relation just in case the claim that Bill doesn't know that the bank will be open is relevant for the question of whether Bill can rule out the salient alternative that the bank has changed its hours relative to Hannah's background assumptions. The only belief we have to ascribe to Hannah to uphold this assumption, however, is the belief that, if Bill doesn't even know that the bank will be open, then he will most likely also not be in a position to rule out the alternative that the bank has changed its hours (which is, after all, more far-fetched than any alternative one has to rule out to obtain knowledge, according to MPI). Again, it is not at all implausible that Hannah should hold this belief. So, as before, we cannot calculate the relevant implicature on the basis of Relation.⁷⁰

⁶⁸ See (DeRose 2009: 122) for essentially the same point.

⁶⁹ See (Blome-Tillmann 2013b: 4309f) and (Dimmock and Huvenes 2014: Sec. 6) for proposals along these lines.

⁷⁰ Note, besides, that I find it highly implausible that the question at issue in the high standards bank cases should only be whether Bill can rule out the salient alternative. Here is why. The proponent of MPI has to agree that in the low standards bank cases Hannah's knowledge ascription is relevant (because no implicature is supposed to be generated in these cases). Hence, at least one of the questions at issue in the low standards cases must be whether Bill knows that the bank will be open. However, if this question is at issue in the low standards cases, then there is no reason why it should not also be at issue in the high

Now, it might be held that, at least if we hold that the issue in the bank cases is whether Bill can rule out the salient alternative, then it is possible to calculate the relevant implicature on the basis of Quantity₁. The reasoning would go as follows. If Hannah observes Quantity₁ at the level of what is said, then what she says must be as informative as is required. What she says, however, isn't as informative as is required. For, even if we grant that what Hannah says makes likely a particular answer to the question of whether Bill can rule out the salient alternative, it does not entail any such answer. Thus, Hannah fails to observe Quantity₁ at the level of what is said.

At this stage, however, we must remember what Quantity₁ actually demands. We have seen that, according to Quantity₁, a speaker should be only as informative as she can be *without violating the quality maxims*. In fact, then, we can only reason as follows. If Hannah observes Quantity₁ at the level of what is said, then what she says must be as informative as she can be without violating the quality maxims. It is not at all clear, however, that there are any reasons for Sarah to reject the consequent of this conditional. For, again, Sarah has no clue what Bill told Hannah on the phone. Correspondingly, it is a perfectly open possibility for her that Bill didn't say whether he can rule out the salient alternative that the bank has changed its hours. So, again, it is impossible for Sarah to calculate the relevant implicature on the basis of Quantity₁.

So much for Hannah's knowledge denial in the high standards cases. Let's now turn to her knowledge ascription. The situation in this case looks even worse. This time, the appeal to quality is straightforwardly ruled out because, according to MPI, Hannah says something she justifiably believes. We can thus turn directly to the calculation of the relevant implicature in terms of Relation. As before, we have to ask whether there are any implausible beliefs that we have to ascribe to Hannah to uphold the assumption that she observes Relation at the level of what is said. Again, I will argue that there are no such beliefs.

standards cases. After all, it would seem bizarre if raising the stakes and mentioning further alternatives should make knowledge irrelevant and leave as relevant only the mentioned salient alternatives. In high standards situations, we plausibly strive for *stronger* epistemic positions, that is, to rule out *larger* sets of alternatives; we don't just strive to achieve *different* epistemic positions, that is, to rule out *different* sets of alternatives. So, knowledge should be relevant in the high standards cases. Hence, the question at issue in the high standards cases should be whether Bill knows that the bank will be open *and* can rule out the salient alternative, that is, roughly, whether Bill knows for sure that the bank will be open.

What beliefs do we have to ascribe to Hannah to uphold the assumption that she observes Relation at the level of what is said? According to MPI, Hannah expresses the proposition that Bill knows that the bank will be open when she says “Bill knows that the bank will be open” in the high standards cases. The question at issue is, again, whether Bill knows for sure that the bank will be open (or, alternatively, whether Bill can rule out the salient alternative that the bank has changed its hours). So, Hannah will be observing Relation just if the claim that Bill knows that the bank will be open is relevant for the question of whether Bill knows for sure that the bank will be open (or, alternatively, whether Bill can rule out the salient alternative that the bank has changed its hours) relative to Hannah’s background beliefs. Which background beliefs do we have to ascribe to Hannah to uphold that assumption? Presumably, we have to ascribe to Hannah the belief that, if Bill already knows that the bank will be open, then there is a good (or at least improved) chance that he even knows for sure that the bank will be open (or, alternatively, that he can even rule out the salient alternative that the bank has changed its hours). Are there any reasons for Sarah not to ascribe this belief to Hannah? I don’t see any. After all, the belief is quite plausible. So, again, the Relation strategy fails.

As before, one might want to invoke Quantity₁. For, again, one might want to say that Hannah is not as informative as is required. After all, she doesn’t directly answer the question of whether Bill knows for sure that the bank will be open (or can rule out the salient alternative). As before, however, we can respond that, in order to observe Quantity₁, Hannah needs to be only as informative as she can be *without violating the quality maxims*. And, also as before, Sarah has no reason to believe that Hannah would be in a position to say something more directly to the point without violating the quality maxims. For example, it should be perfectly unproblematic for Sarah to assume that Hannah just cannot tell whether Bill knows for sure that the bank will be open (or, alternatively, can rule out the salient alternative). Thus, again, the Quantity₁ strategy fails.

In sum, neither the alleged implicature of Hannah’s knowledge ascription in the high standards cases nor the alleged implicature of Hannah’s knowledge denial in the high standards cases can be calculated along exploitative lines. Hence, we should reject the semantic/pragmatic thesis of MPI once it is construed in terms of substitutional implicatures.

3.5. Conclusion

To conclude, the proponent of MPI can either appeal to additive implicatures or to substitutional implicatures. If she appeals to additive implicatures, she cannot sustain the bridging thesis, according to which the seeming truth-value of a given claim is often determined by the truth-value of the implicatures this claim carries. If the proponent of MPI appeals to substitutional implicatures, she doesn't face this problem (assuming that we grant the existence of implicatures). Still, she cannot sustain the semantic/pragmatic thesis, according to which the knowledge claims in the bank cases have the relevant implicatures. The reason for this is that substitutional implicatures must be calculable along the lines of the pattern discerned for exploitation implicatures; and the implicatures posited by the proponent of MPI cannot be calculated along these lines. Thus, either way, MPI must be rejected.

Before we go on, I should note that the objections against MPI I have discussed in this chapter are by no means to only objections there are. Still, they have a major merit that sets them apart from other objections in the literature, namely, they don't appeal to any potentially controversial intuitive datum (other than the Contextualist Data, which, as indicated, are just presupposed in this dissertation and, besides, can be backed up by empirical results). In particular, the result that additive implicatures cannot be used to explain why a given claim seems true is based on very general considerations as to how we evaluate the truth-value of a given claim. Similarly, the result that substitutional implicatures cannot be calculated in the bank cases is independent of any appeal to intuitive data. The same cannot be said about the other objections present in the

literature.⁷¹ I do not want to claim that these objections have no force for this reason. However, they do seem to have a weak spot that may be difficult to remove.

⁷¹ There is the objection that the alleged implicatures of the knowledge claims in the bank cases cannot be canceled properly even though, *qua* implicature, they should be. See (Dimmock and Huvenes 2014: Sec. 5) and (Petersen 2014: 12). This objection relies on the intuition that the alleged implicatures cannot be canceled properly. Montminy (2007: 105) points out that this intuition is not universally accepted. There is the objection that, at least once the proponent of MPI appeals to implicatures, she will no longer be in a position to explain disagreement or retraction data. See (Montminy 2007: 118) and (Dimmock and Huvenes 2014: Sec. 4). This objection obviously relies on disagreement and retraction data. These are a mine field anyway. And there is the objection that, if knowledge claims had implicatures, people should be aware of this fact in a way in which they aren't. See (Montminy 2007: 117, 2009: 646f) and (Blome-Tillmann 2013b: Sec. 5). This objection relies on the intuition that we aren't properly aware of the alleged implicatures in the bank cases, which is rejected e.g. in (Ludlow 2005: 30f) and (Petersen 2014: 10). (The awareness objection will be discussed in more detail in the context of skeptical versions of pragmatic invariantism below.)

4. Skeptical Pragmatic Invariantism

We now turn to a second version of pragmatic invariantism, namely, *skeptical pragmatic invariantism* (SPI). This view is very similar to MPI, the major difference being that, according to SPI, “knows” invariably expresses a very demanding relation. In more detail, SPI amounts to the following claims. “Knows” invariably expresses the same relation. This relation is such that Hannah’s knowledge denials in the bank cases turn out to express a truth, while her knowledge ascriptions turn out to express a falsity. Thus, SPI straightforwardly explains the data about the high standards cases: Our judgments about these cases just reflect the truth-values of the propositions the proponent of SPI takes to be expressed by the claims in question. But how does SPI explain the data about the low standards cases, where our judgments diverge from the truth-values of the allegedly expressed propositions? As should be expected, the proponent of SPI responds to this question in the following way. The knowledge ascription in the low standards cases seems true because it has a true implicature, and the knowledge denial in the low standards cases seems false because it has a false implicature.

Similar to the case of MPI, we can split up SPI into two subtheses. On the one hand, there is the semantic/pragmatic claim that Hannah’s knowledge ascription in the low standards cases expresses a falsity but implicates a truth, while Hannah’s knowledge denial in the low standards cases expresses a truth but implicates a falsity. On the other hand, there is the bridging thesis, according to which, roughly, truth-value intuitions can be guided by the implicatures of the claims in question.

This time, I will just grant the bridging thesis. For, we have seen in the discussion of MPI that the bridging thesis can be sustained if the proponent of SPI appeals to substitutional implicatures, in particular, implicatures. And proponents of SPI generally appeal to such implicatures anyway.¹ The only question that needs to be addressed then is whether the semantic/pragmatic thesis can be sustained. To simplify the discussion, I will just identify

¹ Schaffer (2004b) defends a version of SPI in terms of substitutional implicatures that are not implicatures. But he has given up this view by now. Implicature SPI is defended in (Conee 2005a: 52f), (Douven 2007) and (Davis 2004: 273f, 2007: 427, 2010: 1155–1156). Bach (2010: FN 31) “suggests, without endorsing,” it. Fantl and McGrath (2009: 183–194; 2012: 693f) claim that the view is “not out of the question.”

SPI with this thesis in what follows. And I will use the term “implicature” to refer only to substitutional implicatures, in particular, (substitutional) implicatures.

So, is SPI a plausible position? My verdict on this issue will be somewhat split. On the one hand, I think that SPI is much more promising than MPI and, thus, deserves to be defended. On the other hand, I think that there are data that rebel against SPI, data, however, that may not be universally accepted and that I cannot back up by empirical results. Correspondingly, the structure of this chapter is as follows. First, I will briefly elaborate on the definition of SPI. In particular, I will discuss how we should interpret the content of the implicatures posited by the proponent of SPI in order to avoid some simple-minded objections. Second, I will discuss and reject what I consider to be more serious objections to SPI. Finally, I will present the intuitive datum that I think speaks against SPI. In particular, I will argue that, if the knowledge claims in the bank cases had the implicatures they are supposed to have according to SPI, then people should be aware of this fact in a way in which they are not. As indicated already, this datum may not be universally accepted. So, I will only somewhat tentatively conclude that SPI must be rejected.

4.1. What is Implicated?

What is the content of the implicatures posited by the proponent of SPI? I take it to be relatively uncontroversial among proponents and sympathizers of SPI that the knowledge claims in the low standards scenarios are supposed to be cases of loose use. Thus, the implicature of the knowledge ascription in the low standards cases is that Bill is *close enough* to knowing that the bank will be open, and the implicature of the knowledge denial in the low standards cases is that Bill is *not close enough* to knowing that the bank will be open. Similar implicatures are generated in many other situations. For example, when I say “It’s three o’clock,” I will often implicate that it is only *close enough* to three o’clock. And, when I say “It’s not three o’clock,” I may implicate that it is *not even close enough* to three o’clock.²

² Note, that, as before in the case of implicature MPI, loose use, *qua* implicature, has a somewhat dubious status. DeRose (2012: 714ff), for example, argues that loose use can be treated semantically, suggesting that, therefore, loose use accounts might just be forms of contextualism. For similar doubts, see (Leite 2005: 121). Davis (2007: 415–417) argues, obviously, that loose use is a pragmatic phenomenon. Again, this is a topic for another occasion.

What does it mean to say that one is close enough to knowing something? I will not engage in the project of spelling out this notion. Still, let me issue a warning that two proposals that have come up in the literature won't do.

The first proposal has it that *S* is close enough 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.)³ This proposal is problematic because it doesn't yield a true implicature for the knowledge ascription in the low standards cases. In more detail, the problem is the following. The low standards bank cases are such that the bank will be open on the next Saturday. So, in these scenarios, nothing unlikely happens. So, if it is true that, unless something unlikely happens, Bill knows that the bank will be open, Bill will know that the bank will be open in these scenarios. However, according to SPI, Bill doesn't know that the bank will be open in the low standards cases. So, it is not true that, unless something unlikely happens, he knows that the bank will be open in these cases. Thus, assuming that Hannah's knowledge ascription implicates that, unless something unlikely happens, Bill knows that the bank will be open, this implicature will be just as false as the proposition allegedly expressed. In order to sustain SPI, however, we need a true implicature for the knowledge ascription in the low standards cases. So, the proponent of SPI cannot accept the above interpretation of the loose use view.

The second proposal to interpret the loose use view has it that *S* is close enough 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.)⁴ This proposal is problematic because it does not respect plausible judgments to the effect that knowledge is factive. Suppose, for example, that, in a scenario otherwise like the low standards cases, the unlikely happened and, unbeknownst to Hannah and Sarah, the bank did change its hours and won't be open on Saturday. Clearly, it would now seem intuitively false for Hannah to say “Bill knows that the bank will be open on Saturday.” For, the bank just

³ Bach (2010: 122f) “suggests without endorsing” this proposal.

⁴ Douven (2007: 333n) endorses this view and explicitly distinguishes it from the first. But he doesn't say why he dismisses the first. Bach “suggests without endorsing” this view as well. But, he favors the first view above. As regards the second, he claims that “it might seem extreme to suppose that most of what ordinarily passes for categorical knowledge is really conditional in content.” But he doesn't give reasons for this contention either. Harman (2007: 176) also seems to propose a view along these lines. His remarks are very sketchy though and he takes his view to be a form of contextualism.

won't be open on Saturday. On the present account, however, the claim would have to seem true. The reason is this.

First of all, it should be clear that, in the modified low standards case just described (where, unbeknownst to Hannah and Sarah, the unlikely happened and the bank changed its hours), Hannah's knowledge ascription should have just the same implicatures as it has in the original low standards cases. For, Hannah and Sarah's conversational context is identical in these scenarios. And it would seem that the presence of an implicature depends on nothing but the conversational context. Thus, if we accept the above interpretation of the loose use account, Hannah's knowledge ascription should implicate that Bill knows that, unless something unlikely happens, the bank will open. But, second, this implicature will be just as true in the modified low standards cases as it is in the original low standards cases. After all, why should the answer to the question of whether Bill knows that, *unless something unlikely happens*, the bank will be open depend on whether the unlikely did happen and the bank did change its hours? So, to the extent that the knowledge ascription in the original low standards bank cases seems true because of the implicature just discussed, the knowledge ascription in the modified low standards cases should seem true for just the same reason. Clearly, though, it doesn't.

The above proposals should thus be dismissed. As I said, I will not engage in the project of trying to provide an alternative interpretation of the notion of being close enough to knowledge. But let me briefly indicate a promising account in terms of the already mentioned relevant alternatives framework. In terms of that framework, we can say that *S* is close enough to knowing that *p* iff *p* and *S* can rule out all alternatives to *p* that are at issue in the given context.⁵ Such a view would not be subject to the above criticisms. For, first, we now get an arguably true implicature for the knowledge ascription in the low standards cases, namely, the claim that the bank will be open and Bill can rule out all alternatives to this claim at issue in the low standards cases (these alternatives presumably include only moderately distant alternatives because not much is at stake and no far-fetched error-possibilities have been mentioned in the low standards cases). 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 know that *p* only if *p*. Thus, factivity intuitions are respected.

⁵ See (Davis 2007: 420) for considerations that seem to go into a similar direction.

We now have a sufficiently clear conception of the version of SPI that interests us here to consider some genuine objections. Let's start with objections I don't consider decisive.

4.2. Objections and Replies

Sentence and Speaker Meaning

To begin with, consider the following, most commonly raised worry.⁶ In a nutshell, this worry has it that SPI is incompatible with plausible ideas as to how words acquire their meaning. In more detail, it goes as follows. If, as the proponent of SPI has it, “knows” does not convey what it literally expresses in the low standards bank cases, then it will not 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 cases. In particular, the subject's epistemic position is roughly as good as Bill's and there isn't much at stake and no error-possibilities have been mentioned. However,

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 2013b: 4304)

Adapted to my terminology, “eminently plausible” views on meaning determination entail that it is impossible for an expression to literally express something other than what it conveys in most ordinary conversations. Thus, “knows” must convey what it literally expresses in the low standards cases. So, SPI is wrong.

Response 1

This worry doesn't seem decisive because there is not much of a pressure for the proponent of SPI to accept the view on meaning determination described by Blome-Tillmann in the quote above. For example, the proponent of SPI could accept the following closely related and seemingly equally plausible view instead. What a term literally expresses is determined by what the term standardly conveys *when used literally*. This view may amount to a pointless triviality if we take the literal uses of a term to be just those uses where the term is used to convey what it literally expresses. But we need

⁶ See (Hawthorne 2004a: 121–123), (DeRose 2009: 125), (Fantl and McGrath 2009: 48) and (Blome-Tillmann 2013b: 4304f).

not do so. Plausibly enough, there are independent tests for non-literality. For example, it seems plausible that there are independent signs of irony, hyperbole, etc.⁷ So, all that we can establish is 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*. But this principle does not threaten SPI. For, the proponent of SPI will hold that, on the rare occasions where “knows” is used literally (say, when we speak of *a priori* truths or truths we gather from immediate perception), it does convey what it literally expresses according to SPI, namely, a quite demanding epistemic relation.

Response 2

Let me also discuss a second response strategy to the above objection. It is based on the idea that Bill 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 Bill 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 subjects to which we ordinarily ascribe knowledge generally equals or is worse than Bill’s in the bank cases. 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 Bill’s. 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.^{8,9}

⁷ One may want to object at this point that the knowledge claims in the bank cases just don’t show these signs of non-literality. In fact, my objection to SPI below (in terms of the awareness principle) will be an objection precisely along these lines. Note, however, that this objection is entirely independent of the general semantic considerations discussed above. If, as the proponent of SPI has it, knowledge claims are often cases of non-literal speech, then they should show the signs of non-literal speech whatever we think about how expressions acquire their meaning.

⁸ The SPI accounts in (Conee 2005a: 52) and (Davis 2007, 2010) suggest responses along these lines.

⁹ This move may 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.” (2004a: 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: 344f) for some further discussion of this issue.

Here is a possible rejoinder to this response. Note that, according to this response, we often use “knows” literally. But if that is so, Cohen (2005a: 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.”¹⁰ 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.

I reject this rejoinder because it relies on what seems to be a dubious datum, namely, the datum that, in the face of skeptical pressure, competent speakers deny knowledge even of “the 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 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)¹¹

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

It must be granted that at least some people (some philosophers, for example) become convinced by wide-ranging skeptical arguments and, as a result, come to think that they know hardly anything at all. Similarly, there certainly are skeptical arguments that are

¹⁰ See (DeRose 2012: 712f) for a related worry.

¹¹ See also (Conee 2005a: 53) and (Douven 2007: Sec. 2) for related considerations.

very difficult to answer. These data cannot be explained on the basis of SPI once we accept the above response. However, I don't think that these data and the more mundane bank case reactions need to be explained in a uniform fashion. One candidate account for why skeptical arguments sound compelling to some people that would be perfectly compatible with SPI is that skeptical arguments simply involve one or the other easily missable argumentative fallacy.¹² This issue would have to be addressed in more detail. But I think it should be clear that the burden of proof is on those who want to hold that skeptical arguments do yield troublesome data for the proponent of SPI.

Expansion

So much for the first worry for SPI. Here is a second worry. It is a worry that Blome-Tillmann raises against the proponent of MPI who appeals to (expansion) implicatures. But it could just as well be raised against the forms of SPI we are considering in this section (that also appeal to (expansion) implicatures). Blome-Tillmann states this worry in the following quote. As indicated, he addresses MPI, but it should be obvious how to transfer his considerations to the case of SPI.

[I]f we took Hannah's utterance of ['Bill doesn't know whether the bank will be open' in the high standards cases] [...] to be a case of expansion, then it should—just as in the examples in (23) [which correspond to my examples of implicatures above]—be fairly obvious what the additive but suppressed lexical material of that utterance would consist in. That this is far from obvious, however, casts significant doubt on the view that the utterances at issue are instances of Bach's expansion implicature. (Blome-Tillmann 2013b: 4317)¹³

The objection, here, cannot be that it is *impossible* to expand Hannah's knowledge claims such that they express their alleged implicatures. For, according to SPI, Hannah's knowledge claims are supposed to implicate that Bill is (not) close enough to knowing that the bank will be open or, equivalently, that Bill is (not) close enough to being such that Bill knows that the bank will be open. And "*Bill is (not) close enough to being such that Bill knows that the bank will be open*" most certainly is an expansion of the sentence

¹² See (Conee 2005b: 66) and (Davis 2007: 436) for the idea of separating the account of our reactions to skeptical arguments and the account of bank case reactions. See (Sosa 1999: 148) and (Feldman 2001: 75f) for candidate fallacies that might be involved in skeptical arguments. Sosa, for example, argues that knowledge requires "safe" not "sensitive" beliefs. On that basis, he goes on to suggest that skeptical arguments sound compelling because "[s]afety and sensitivity, being mutual contrapositives, are easily confused."

¹³ See (Leite 2005: 122) for a similar objection.

“Bill knows that the bank will be open.”¹⁴ The objection rather is this. First, it is, in some sense, “far from obvious” what the additional but suppressed lexical material of Hannah’s utterances could be. Second, this should be fairly obvious if they are cases of implicatures. So, they are not cases of implicatures.

I grant the first premise. However, I deny the second. This premise may initially seem plausible. For example, when the parents say to the child “You are not going to die,” and the child responds “I am immortal then?” the parents will straightforwardly come up with the expansion “You are not going to die *from the cut*.” However, this observation doesn’t generalize. Consider the following dialogue,

Teacher: Everyone is there. Let’s start.

Student: Everyone? What about the pope?

Teacher: Well, everyone who’s participating in the course.

Student: What about John? (*Assume that John participates in the course but is absent because of a long hospital stay.*)

Teacher: Well, everyone who’s supposed to be here is there!

Student: But wasn’t the janitor supposed to be here to repair the blackboard? (*Suppose that’s true.*)

Teacher: You know what I meant! Let’s start.

Presumably, the teacher meant that everyone whose absence would be a reason not to start is there. But she couldn’t articulate it. Such things do happen. So, it seems that it is not generally obvious how to articulate implicatures. But if that is so, then it doesn’t seem to be a problem that it is not obvious what the expansions of Hannah’s knowledge claims should look like (so long as these expansions are available at all). Hence, Blome-Tillmann’s objection has no force.

Calculability

Consider next the idea of objecting to SPI on the basis of the idea that the alleged implicatures aren’t properly calculable. Given that the implicatures posited by the

¹⁴ By the way, the same goes if Blome-Tillmann’s objection is leveled against MPI. Blome-Tillmann himself points out, in a different paper, that we can say things like “*x* knows with very/quite/extremely good evidence/justification that *p*.” (2008: 46) These sentences clearly are expansions of the sentence “*x* knows that *p*” and at least some of them (or their negations) may well be used to express what Hannah’s knowledge claims implicate according to MPI.

proponent of SPI are supposed to be substitutional, they must be calculable along the lines of the pattern discerned above for exploitation implicatures. It is relatively straightforward, however, to calculate them along these lines.

Consider the knowledge ascription in the low standards cases first. Here the standard proposal is roughly this.¹⁵ If Hannah observes Quality₁ (“Be sincere”) at the level of what is said when she says that Bill knows that the bank will be open, then she must believe that Bill knows that the bank will be open. Given SPI, however, Hannah doesn’t believe that (given what Bill told her). Thus, Hannah fails to observe Quality₁. So, we get an exploitation implicature, plausibly to the effect that Bill does not know that the bank will be open but still is close enough to knowing that the bank will be open.

One might want to object, as we did before in the case of MPI, that Sarah need not be aware of the Quality₁ violation and, thus, will fail to calculate the relevant implicature in the way just proposed. In the present context, however, this response has no force. For, depending on how the line between knowledge and non-knowledge is drawn, the proponent of SPI may either argue that we hardly know anything at all or, at least, that we generally do not have knowledge of the kind that is involved in the bank cases (say, knowledge about complex, specific future facts). Either way, Sarah will be in a position to see the Quality₁ violation because she will plausibly be aware that we hardly if ever have the relevant kind of knowledge. So, the above derivation seems fine.

Consider the knowledge denial in the low standards cases next. On the basis of the above proposal, it is easy to see how the relevant implicature in this case can be calculated. The proponent of SPI can argue that it is just trivial that Bill doesn’t know that the bank will be open (because we hardly if ever have knowledge of this kind). Thus, what Hannah says is completely uninformative (and, given that she has just been talking to Bill about his epistemic position with respect to the claim that bank will be open, it should be clear even to Sarah that Hannah would be in a position to provide at least some relevant information). So, Hannah cannot be assumed to be observing Quantity₁ (“Be as informative as is required”) at the level of what is said. This will trigger the search for an

¹⁵ See (Schaffer 2004b: 140), (Davis 2007: 413) and (Douven 2007: 339).

alternative content she could have meant. And a plausible candidate here is that Bill is not even close enough to knowing that the bank will be open.

4.3. Awareness

Let's now turn to what I take to be a more serious worry. It is based on the observation that, if a given utterance has a substitutional implicature, then people are, in a yet to be specified sense, aware of the discrepancy between what they say and what they mean to convey, for short, they are aware of their implicatures. Let's refer to this idea as the *awareness principle*.¹⁶ We can cash it out in the following way. Note, to begin with, that, given Principle, all substitutional implicatures are exploitation implicatures. This means that, whenever a given utterance has a substitutional implicature, the literal content of that utterance violates at least one conversational maxim. So, if a hearer either fails to realize that an utterance has such an implicature or intentionally takes the utterance too literally—is being a “stickler for correct speech” to use a phrase from Grice—, it will be natural for her to object to the literal content of the relevant utterance on the grounds that it violates the relevant maxims. The idea that, in the case of substitutional implicatures, people are aware of their implicatures can now be spelled out in the following way. A speaker confronted with such an objection will straightforwardly be able to make sense of this 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. Suppose the following dialogue takes places.

A: The plane was a mile long.

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

Here, A's first utterance generates a substitutional implicature to the effect that the plane was huge. B's response is an objection to the literal content of A's claim, in particular, to its violation of Quality₁. Now, it seems pretty clear that A, as a competent speaker, will

¹⁶ See (Bach 2005: 73), (MacFarlane 2005b: 206f, 2011: 541) and (Turri 2014) for a similar observation as part of a case against SPI. See also (Hawthorne 2004a: 104f). Hawthorne only presents this observation as an objection to contextualism. However, as the examples he discusses suggest, he seems willing to classify certain forms SPI (those that appeal to loose use) as forms of contextualism. Fantl and McGrath (2012: 693f) also acknowledge a principle along the above lines (but deny that it cuts against loose use SPI). See (Montminy 2009: 646f) and (Blome-Tillmann 2013b: Sec. 5) for an application of the principle to the case of MPI.

straightforwardly realize that *B* is (intentionally or unintentionally) missing the point and, thus, respond accordingly, for example, as follows,

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

The present example involved a violation of Quality₁. But it seems that people are not only aware of Quality₁ substitutional implicatures, but of substitutional implicatures in general.

Consider, for example, a case where the relevant maxim is Relation. *A* and *B* are wondering what to do about John's absence. The following dialogue transpires.

A: John will come or he won't.

B: Well, these are all the options, aren't they? But I was wondering what we should do about his absence.

Again, the following response on the part of *A* should be readily at hand for all competent speakers.

A: Oh, c'mon. You know what I meant. (There is nothing we *can* do about his absence.)

Here is an example where the violated maxim is Quality₂. Consider the following case as described by Grice.

I say of X's wife, *She is probably deceiving him this evening*. In a suitable context, or with a suitable gesture or tone of voice, it may be clear that I have no adequate reason for supposing this to be the case. My partner, to preserve the assumption that the conversational game is still being played, assumes that I am getting at some related proposition for the acceptance of which I do have a reasonable basis. The related proposition might well be that she is given to deceiving her husband, or possibly that she is the sort of person who would not stop short of such conduct. (Grice 1989: 34f)

In the just described context, the following dialogue could occur.

A: She's probably deceiving him this evening.

B: How can you know that? You do not even know whether she went out today!

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

So far, we have considered only examples that involve implicatures that are not implicatures. But what I have said also goes for implicatures. With respect to quantifiers, Stanley for example claims that

the unrestricted interpretation is also available to competent language users, as the coherence of the following sort of discourse illustrates:

A: Every bottle is in the fridge.

B: Well, your fridge couldn't possibly be that large! There are bottles somewhere in the world that aren't in your fridge. (Stanley 2007: 211)

Given this dialogue, it would be perfectly natural for A to respond by saying "Oh, c'mon. You know what I meant."¹⁷ Related observations have frequently been made with respect to implicatures of the loose use type. Here as well, dialogues of the following type seem fine.

A: It's {about} three o'clock.

B: It's one minute past!

A: Oh, c'mon. You know what I meant.¹⁸

In the above two implicature cases, the objections targeted violations of Quality₁. But we can also take claims with implicatures too literally by objecting to violations of Relation. Consider, for example, the following case, where the teacher tries to defend herself against the charge that a given riddle was too difficult a homework.

Teacher: Someone even solved the riddle!

Parents: So? That was presumably you!

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

This dialogue also seems fine. So, it seems that we are aware of implicatures in just the same way in which we are aware of other implicatures.¹⁹

¹⁷ For similar remarks, see also (Grice 1989: 45), (Conee 2005a: 49) and (Montminy 2009: 646f).

¹⁸ See e.g. (Laserson 1999: 522), (Bach 2005: 73), (MacFarlane 2005b: 207, 2011: 541), (Blome-Tillmann 2013b: 4302) and (Turri 2014).

¹⁹ The idea of cashing out the awareness principle in terms of the naturalness of dialogues like the ones above is taken from Blome-Tillmann. Note some points of divergence though. First, I have argued that B's respective objections can either be understood 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 (2013b: 4311n) envisages only the first kind of understanding. The stickler reading, however, often is the only reading available, particularly when it

Given the awareness principle, the problem for the proponent of SPI is this. Consider Hannah's knowledge ascription in the low standards cases. According to SPI, this claim has a substitutional implicature (in particular, a loose use implicature) because what Hannah says violates Quality₁. So, given the above awareness principle, the following dialogue should be fine in the low standards cases.

Hannah: Bill knows that the bank will be open.

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

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

However, many would say (and I would agree) that Hannah's response seems not particularly natural. As Hawthorne (2004a: 105) claims, "our standard techniques for dealing with epistemic challenges that raise relatively farfetched possibilities are concession and, more rarely, sticking to one's guns." So, Hannah could plausibly *concede* "You're right. I didn't think of that" or *stick to her guns* by saying "(Oh, c'mon.) How likely is that?" The above, *clarifying* response, however, seems unnatural.^{20,21}

Consider also Hannah's knowledge denial in the low standards cases. The proponent of SPI has it that this claim has a substitutional implicature that is based on the violation of

comes to implicatures. (The concept of a stickler is taken from (or at least inspired by) (Grice 1989: 44f). It is also employed in (Fantl and McGrath 2012: 693f).) Second, Blome-Tillmann thinks that, in her second utterance, A should be able to "clarif[y] her initial intentions by admitting that she hasn't spoken literally but exaggerated, spoken ironically, etc." (2013b: 4311n) Similarly, one might think that A should be able to state what exactly it was that she implicated (Hawthorne's discussion may suggest that). I do not want to go so far as that. For, the categories of exaggeration, irony, etc. may be unfamiliar. And I don't think that speakers are generally able to tell what exactly it is they implicate (see the discussion of expansion above). Still, it seems reasonable to require that speakers be able to tell that the hearer was missing the point. This seems to be captured by the response "You know what I meant!" Third, Blome-Tillmann (just like all the other authors mentioned in the previous footnotes) considers only objections on the part of B to the truth-value of A's literal claim, that is, only cases of Quality₁ violations. I think that B's objections may target whatever maxim violation there is.

²⁰ See also (Bach 2005: 73), (MacFarlane 2005b: 207, 2011: 541) and (Turri 2014).

²¹ Let me briefly mention a closely related potential objection to SPI that I consider indecisive. It might be argued that claims that seem true because they are close enough to being true stop to seem true once we add something like "strictly speaking." For example, it may seem true to say "It's three o'clock" when it is actually ten seconds past. Still, in the same situation, it will arguably not seem true to say "Strictly speaking, it is three o'clock." Now, if Hannah's knowledge ascription in the low standards cases seems true because it is a case of loose use, a corresponding utterance of "Strictly speaking, Bill knows the bank will be open" should not seem true. And, now it may be argued that this just isn't the case. See (Hawthorne 2004a: 120) and (Blome-Tillmann 2013b: 4302) for related points. However, Conee (2005a: 52) (and, similarly, Davis (2007: 430)) holds that it would indeed seem false for Hannah in the low standards cases to say "Bill *really and truly* knows that the bank will be open." Thus, the present objection rests on questionable intuitive verdicts.

Quantity₁. Correspondingly, dialogues of the following type should be fine if the awareness principle holds.

Hannah: Bill doesn't know that the bank will be open.

Sarah: Of course he doesn't. No one knows such things. Is that all he told you on the phone?

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

Again, though, such a dialogue seems hard to make sense of.

In order to avoid distracting objections, let me emphasize that the "You know what I meant!" responses on the part of Hannah in the dialogues just described are subtly different from other possible responses on her part that I would find perfectly acceptable. Blome-Tillmann, for example, suggests that, in a dialogue like the one above, "it might seem natural for [Hannah] to reply that she was speaking loosely." (2013b: 4303n)²² I fully agree. Similarly, consider the following dialogue.

Hannah: Bill knows that the bank will be open.

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

Hannah: You're right. Bill doesn't know it.

It seems fairly uncontroversial that Hannah could now go on to say something like "I was speaking casually," particularly, if Sarah were to insist "But you just said that he does!"²³

The fact that these responses are acceptable, however, does not show that the above "You know what I meant!" response is acceptable as well. To see this, consider the following dialogue.

Hannah: The bank will be open.

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

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

Just as before, Hannah could now go on to say "I was speaking casually/loosely," particularly, if Sarah were to insist "But you just said that it will be open!" Obviously,

²² For somewhat related judgments, see (Fantl and McGrath 2012: 694).

²³ See e.g. (DeRose 2009: 171), (Davis 2010: 1155) and (Dimmock and Huvenes 2014: 6n).

though, her original assertion wasn't a case of loose use (in the relevant sense). More plausibly, it was just an ill-considered judgment.²⁴ And when Hannah says that she was speaking loosely/casually, she plausibly just wants to admit that she wasn't careful enough when she made her original statement, not that she didn't mean what she said. Correspondingly, it would *not* seem fine for her to respond "Oh, c'mon. You know what I meant!" The same, I think, goes for the dialogues above.

For those who do not share these verdicts, let me note that there is another frequently adopted intuitive datum that may create some pressure towards following my lead. It is the (admittedly somewhat vague) datum that high standards contexts are contexts where we feel "enlightened" in some sense.²⁵ Let me elaborate. Consider the following dialogue again.

Hannah: Bill knows that the bank will be open.

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

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

I have objected to SPI on the grounds that this dialogue sounds weird. If you accept the idea that high standards contexts are enlightened contexts, you should accept that verdict, for the following reason. In the above dialogue, Sarah objects to Hannah by mentioning an error-possibility. Presumably, she thereby raises the epistemic standard. Thus, given the assumption that high standards contexts are enlightened contexts, Hannah should feel enlightened in some sense after hearing Sarah's objection. However, when Hannah responds to this objection by saying "Oh c'mon. You know what I meant!" she is precisely denying that there is anything enlightening about this objection. To the contrary, she is suggesting that she 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, we should accept my argument against SPI.²⁶

²⁴ Davis (2007: 410f) similarly distinguishes loose use from what he calls "sloppy" or "careless" or "unenlightened" use.

²⁵ See e.g. (Hawthorne 2004a: 164) and (Cohen 2005a: 58) for the thesis that high standards contexts are enlightened contexts.

²⁶ 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, the fact that we tend to retract knowledge

4.4. Conclusion

To conclude, I am almost certain that some of my readers will neither accept the intuitive verdicts about the dialogues above nor the enlightenment intuition. I have nothing to say to these readers. The upshot of the present chapter for them should be that SPI yields a very attractive account of the Contextualist Data. For, apart from the awareness objection just discussed, there seem to be no decisive objections to this view on offer. For those who do share my verdicts about the dialogues above or, at least, the enlightenment intuition, the upshot of the present chapter should be that SPI must be rejected. In any case, I will now turn to what I take to be the correct account of the Contextualist Data, projective invariantism.

claims after an error-possibility has been mentioned. Loose use may explain retraction. See e.g. (Davis 2007: 406f). 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, as I said, loose use cannot explain enlightenment (even though it can explain retraction).

5. Projective Invariantism

In this chapter, I will propose an invariantist account of the Contextualist Data that does not suffer from any of the problems discussed so far. I call it *projective invariantism* because it is based on the idea that readers *project* certain assumptions into the scenarios they read. Before we start, let me briefly outline this view and discuss how it relates to the views considered so far.

Projective invariantism is, of course, a form of invariantism. Thus, projective invariantists hold that the proposition expressed by knowledge claims does not vary with the epistemic standard of their context of use. As a result, they face the challenge from the Contextualist Data. To respond to this challenge, they argue roughly as follows. We evaluate the knowledge claims in the bank cases differently because the epistemic position of the subject of these knowledge claims differs between the bank cases. This, in turn, is so because there are certain straightforward psychological mechanisms that lead readers of the bank cases to interpret these cases in relevantly different ways. In particular, these mechanisms lead readers of the bank cases to assume that it is more likely in the high than in the low standards cases that the bank has changed its hours because this possibility is mentioned only in the high standards cases.

Projective invariantism, thus, differs from all the views discussed so far. Consider relativism (chapter 2) and pragmatic invariantism (chapter 3 and 4). With these views, it shares only the commitment to invariantism. For, projective invariantists neither appeal to relative nor to pragmatic contents in the way these theories do. Consider subject-sensitive invariantism and the belief removal theory (chapter 1). With these views, projective invariantism further shares the commitment to the idea that the epistemic position of the relevant subject differs between the bank cases. However, according to projective invariantism, the difference in the subject's epistemic position is neither due to anti-intellectualistic features of the subject nor to differences in the beliefs she holds (as these views would have it). Projective invariantism presumably is most closely related to psychological error-theories (chapter 1). With these theories, it shares the commitment to invariantism and the appeal to psychological principles. Furthermore, as most versions of the psychological error-theory, projective invariantism also trades on how likely we

consider the salient error-possibility in the high standards cases to be. Still, there are important differences. In particular, projective invariantism is not an error-theory and it does not appeal to the psychological mechanisms psychological error-theorists have so far employed.

This is projective invariantism in rough outline. In this chapter, I will further spell out and defend this view. Correspondingly, the structure of the chapter is as follows. In the first section, I will lay out the projective invariantist account of the Contextualist Data in more detail. In the second section, I will show how to deal with various objections to this position.

5.1. The Contextualist Data

In this section, I will present the projective invariantist account of the Contextualist Data. We have seen that these data are problematic for the invariantist for the following reason. On the one hand, the invariantist wants to say that the propositions expressed by the knowledge claims in the low and the high standards bank cases are the same. On the other, it seems that the relevant epistemic facts are the same in these cases as well. As a consequence, it seems that our truth-value judgments about the knowledge claims in the low and the high standards cases should be the same. But, according to the Contextualist Data, they are not. The strategy that I think the invariantist should pursue to respond to this challenge is to deny its second premise. The invariantist should claim that, contrary to appearance, the epistemic facts are, or more precisely, will be assumed to differ between the bank cases, in particular, that Bill's epistemic position will be assumed to be stronger in low standards cases than in the high standards cases.

In light of the above discussion of the third-person cases, such a strategy may seem hopeless. Remember that the corresponding low standards bank cases seem to differ from the high standards cases only in that, in the high standards cases, more is at stake for Hannah and Sarah and Sarah mentions the error-possibility that the bank has changed its hours. The epistemic subject, Bill, however, is neither part of nor aware of the conversation between Hannah and Sarah. Thus, it seems very much unclear how the above differences should affect his epistemic position. So, let me first present a general

framework that opens up the possibility of claiming that Bill's epistemic position differs between the bank cases to then introduce a particular implementation of this framework.

Here is the general framework. I grant that as far as the explicit content of the scenario descriptions is concerned, there is no epistemically relevant difference between the low and the high standards cases. However, when we ask whether Bill's epistemic position differs between the low and the high standards cases, we should not only look at what's explicitly said in the description of these scenarios. This is so because these descriptions leave open epistemically relevant features (such as, for example, the likelihood of the possibility that the bank changes its hours) that are going to be filled in in one way or another by the readers of the scenarios. And epistemic differences may reside in these non-explicit features just as well as in the explicitly described bits of the scenarios. Hence, one way to respond to the above argument against the idea that Bill's epistemic position differs between the bank cases is to show that, even though the explicitly described bits of the scenarios do not differ in any epistemically relevant respect, the descriptions invite readers to fill in the scenarios such that there are epistemically relevant differences after all.

In what follows, I will present a specific implementation of this general framework. To do so, I will, first, highlight one specific epistemically relevant yet unspecified feature of the bank cases. Second, I will provide considerations that make it at least *prima facie* plausible that this feature will be filled in differently in the different bank cases.

As to the underspecified feature, note that nothing is said explicitly in the scenario descriptions about how likely it is supposed to be that Hannah and Sarah's bank changes its hours. Nothing is said, for example, about how frequently it changed its hours before, how frequently shops in general change their hours in Hannah and Sarah's environment, how reliable the director of the bank is supposed to be, etc. Nothing is said about all these aspects of the situation. Still, quite plausibly, they affect Bill's epistemic position with respect to the claim that the bank will be open. For, the more likely it is that the bank

changes its hours, the easier it will be for Bill to be wrong about the opening hours of the bank. And, plausibly, the easier it is to be mistaken, the weaker one's epistemic position.¹

So, here we have an epistemically relevant yet unspecified feature of the bank cases, namely, the likelihood of the possibility that the bank changes its hours. I will now present three mutually supportive, yet self-standing arguments for the claim that this feature will be filled in differently in the different bank cases, in particular, for the claim that the possibility that the bank changes its hours will be assumed to be less likely in the low than in the high standards bank cases. This result should suffice to establish that Bill's epistemic position will be assumed to be weaker in the high than in the low standards cases.

In a nutshell, the arguments go as follows. According to the first argument, the likelihood of the possibility that the bank has changed its hours will be assumed to be higher in the high standards cases because we are reminded of this possibility when Sarah mentions it in these scenarios. According to the second argument, the likelihood of the possibility that the bank has changed its hours will be assumed to be higher in the high standards cases because Sarah implicates that this possibility is likely when she mentions it. According to the third argument, the likelihood of the possibility that the bank has changed its hours will be assumed to be higher in the high standards cases because this possibility would not occur to Sarah unless it was likely in these cases. I take the first argument to be the most original contribution to the debate. The second and the third argument draw heavily upon proposals already present in the literature. I put them forth nevertheless because they haven't been embedded into the general framework described above.²

Before we can spell out the just mentioned arguments in more detail, we first have to ask how readers of a scenario generally fill in the missing details of a scenario. This is

¹ Note that, when I speak of likelihoods here and in what follows, I have in mind some kind of more or less objective sort of likelihood that is relatively independent of the evidence available to particular subjects. As indicated above, for example, the likelihood of the possibility that Bill is mistaken is supposed to depend on features such as how frequently the bank has changed its hours before, whether or not Bill is aware of how often the bank has changed its hours before. DeRose (2009: 22f) and Gerken (2011: 56), for example, endorse the idea that a subject's epistemic position depends on the likelihood of error-possibilities in this objective sense.

² See below for references.

ultimately an empirical question. But the following seems like a principle that should guide us at least to a certain extent.

Normalcy Projection (NP): Subjects will tend to fill in the missing details of a scenario in accordance with what they think would be normal in the context of the scenario as explicitly described.

The idea behind this principle is that if the author of a scenario fails to specify a relevant bit of the scenario, then, presumably, this part of the scenario is not supposed to be special in any way. Hence, we will just project what we consider normal into the empty yet relevant slots. As I said, (NP) is ultimately an empirical principle that stands in need of empirical confirmation. In this dissertation, however, I will just assume that something along the lines of (NP) is true and proceed from there.

The argument from reminding

Here then is a detailed version of the first of the above three arguments to the conclusion that the bank will be assumed to be more likely to change its hours in the high standards cases. Plausibly, when we don't think of an error-possibility, we tend to act as if we thought that its likelihood is zero even if we already know (in some sort of unconscious or tacit way) that the likelihood is not zero. For example, suppose that you know full well that robins and bluebirds are equally likely to be seen in your area. Suppose further that you see an orange-breasted bird that, for all you see, is either a robin or a bluebird. If you don't think of the possibility that the bird is a bluebird (maybe because it's been a while since you last saw a bluebird), you will presumably still say, "That's a robin!" This is so because you treat the likelihood of unthought-of possibilities as zero.³ When subjects read the low standards cases, they similarly don't think of the possibility that (normal) banks sometimes change their hours. Hence, given the previous observation, they will treat the likelihood of this possibility as zero, in particular, they will act as if they thought that the likelihood of this possibility is zero. Given (NP), they will thus project this likelihood into the scenario. So, they will assume that the likelihood of the possibility that Hannah and

³ I take it that the above principle is fairly obvious and hence need not be backed up by empirical results. Still, of course, it is in the end an empirical psychological principle. For those who thus seek empirical corroboration, I refer to (Gerken 2011: 50; 2012: 155). The right-to-left direction of what he calls the "Principle of Contextual Salience" seems to correspond at least roughly to the above thesis. Gerken discusses this thesis in a psychologically much more informed fashion than I am able to do.

Sarah's bank changes its hours is zero in the low standards cases. In the high standards cases, the situation is different. It seems clear that once we are reminded of a known-to-be-likely possibility, we no longer treat its likelihood as zero even if we haven't been taking it into account before. For example, when I remind you, "It could also be a bluebird," you'll immediately retract your original assertion of "That's a robin!" because you remember that this possibility is likely.⁴ Similarly, when subjects read the high standards cases, they are reminded of the possibility that (normal) banks sometimes change their hours because this possibility is mentioned in the scenarios. Hence, they also realize that its likelihood is, though maybe low, certainly not zero. Again, given (NP), they will project this assumption into the scenario and thus assume that it is likely at least to a certain degree that Hannah and Sarah's bank changes its hours. So, the assumed likelihood of this possibility will be higher in the high than in the low standards cases. (The just presented argument leads to an account of the Contextualist Data only on the assumption that Bill fails to know that the bank will be open in the high standards cases because the normal likelihood of the possibility that banks change their hours prevents him from knowing that. It may be worried that this idea leads to a problematic form of skepticism. I don't think it does. The issue will be addressed below.)

The argument from implicatures

Here is a detailed version of the second argument for the conclusion that the bank will be assumed to be more likely to change its hours in the high standards cases, in particular, for the partial conclusion that the likelihood of this possibility will not be considered zero in the high standards cases. It is based on the idea that, first, Sarah implicates that it is likely to a certain degree that the bank changes its hours when she mentions that possibility and, second, that this implicature will be assumed to be true in the high standards cases given (NP). Let's discuss these claims in turn.

Why should Sarah implicate that it is likely to a certain degree that the bank changes its hours? Here, we can draw on the work of Dougherty and Rysiew (2009). They pursue a

⁴ The above claim is closely related to the left-to-right direction of Gerken's Principle of Contextual Salience which has it that "[n]ormally, if for an agent, *A*, *q* is a contextually salient alternative to *S*'s knowledge that *p*, then *A* processes *q* as an epistemically relevant alternative to *S*'s knowledge that *p*." However, my contention here is more moderate and therefore hopefully uncontroversial. It roughly says that normally, if for an agent, *A*, an alternative, *q*, is salient and *A* knows that the likelihood of *q* lies above a certain threshold, then *A* "processes" *q* as epistemically relevant.

project that is closely related to mine, namely, they try to explain, from an invariantist (intellectualist and fallibilist) point of view, why so-called “concessive knowledge attributions” sound odd (a concessive knowledge attribution being a claim of the form “I know that p , but it is possible that q ” where q entails not- p). We need not bother about the details of their account here. All that is relevant is that, in the course of their discussion, they plausibly argue for the following claim: “In saying that p is possible one implies that there is some real chance that p .” (2009: 130) (Dougherty and Rysiew do not say much about what a “real” chance is supposed to be. Still, they assume that, if there is a real chance of error, then this “may well be enough to prevent one from knowing.”) Following Dougherty and Rysiew, the above implicature can be derived in the following way.⁵ The mere existence of a possible alternative arguably will not be relevant for the question of whether someone knows that a certain proposition is true. (Otherwise we would presumably be wedded to a fairly strong form of skepticism.) So, when, say, Sarah says “Banks can always change their hours,” she will say something irrelevant to the extent that all she says is that this possibility exists. The irrelevance of her claim will thus trigger the search for an implicature. And this implicature arguably is the thesis that the possibility that banks change their hours not only exists but is also likely to a relevant degree.⁶

Why should readers of the bank cases assume that this implicature is supposed to be true? Plausibly, we do not normally try to mislead people. So, given (NP), Sarah will have to be assumed to believe that there is a real chance that the bank changes its hours.

⁵ See (Dougherty and Rysiew 2009: 128f). See also (Bach 2008: 79) for a similar reasoning. As will become apparent in a minute, the subsequent derivation relies on assumptions about the semantics of possibility claims that may not be entirely uncontroversial (in particular, that their semantic content by itself is irrelevant for the question of whether someone knows a given proposition). Addressing this issue, however, goes beyond the scope of the present project. For further discussion, see e.g. (Stanley 2005a), (Dodd 2010) and (Dougherty and Rysiew 2011).

⁶ I tacitly made the transition from the idea that a given possibility is a “real” possibility to the idea that the possibility is likely. It may be doubted, however, whether this transition is legitimate. For example, one may think that whether a possibility is real depends on how much is at stake. Thus, one may hold, a possibility can be real even if it is unlikely, namely, when much is at stake. See e.g. (Fantl and McGrath 2009: 22f) for such a conception of the realness of a possibility. Here is my response to this worry. Even those who want to say that the realness of a possibility depends on how much is at stake do not want to say that the realness of a possibility depends *only* on what is at stake. Instead, they will say that a possibility is real if it is likely enough for the stakes at hand (or something along these lines). Thus, if a possibility is real, it still follows that it is likely at least to a certain degree. And this is all that is required for the above argument to go through.

(Otherwise she would convey something she does not herself believe. Hence, one would most naturally think that she is trying to mislead Hannah.) Now, in the absence of misleading evidence, we normally hold true beliefs about our everyday concerns (such as the opening hours of our bank). So, given that no reason is given in the bank cases to think that Sarah should have been misled in any way and given (NP), Sarah's belief will most likely be assumed to be true. Thus, we have a second argument for the conclusion that the possibility that the bank changes its hours will be considered more likely in the high than in the low standards cases.⁷

The argument from occurrence

Here is the detailed version of the third argument for this conclusion. As Bach plausibly points out, if “our cognitive processes are operating well, generally the thought of a possibility contrary to something we're inclined to believe occurs to us only if it's a realistic possibility, not a far-fetched one” (2010: 121). Thus, the possibility that the bank changed its hours should normally occur to Sarah only if this possibility is likely to a certain degree. So, given that this possibility does occur to Sarah in the high standards cases, we have to assume that it is likely there in order to preserve normalcy. Hence, given (NP), we have a further argument for the claim that the possibility that the bank changes its hours will be assumed to be less likely in the low than in the high standards bank cases.⁸

Note that the presently described mechanism is importantly different from the implicature mechanism above. For example, the implicature mechanism arguably won't apply when we consider bank cases where the error-possibility is not mentioned in the high standards cases but only considered in thought by the knowledge ascribing subject. For, thoughts (as opposed to mentionings) plausibly don't have implicatures.⁹ The present mechanism, however, straightforwardly does apply to these cases. For, to repeat, an error-possibility normally occurs to somebody (in thought or in speech) only if this possibility is likely.

⁷ For a suggestion that goes in in the direction of this second argument, see May's comments on <http://certaindoubts.com/?p=1373>.

⁸ Again, I am tacitly making the transition from something being a “realistic” or “not far-fetched” alternative to it being a likely alternative. See FN 6 for why this transition is compatible with the idea that whether a possibility is realistic or far-fetched depends on how much is at stake.

⁹ See e.g. (Baumann 2011) for such bank cases (and an argument against pragmatic invariantism that is based on these cases and the just mentioned idea that thoughts don't have implicatures).

To sum up, we have seen that there are three straightforward reasons to hold that the possibility that the bank changes its hours will be assumed to be less likely in the low than in the high standards bank cases. Hence, we have good reasons to assume that Bill's epistemic position will be assumed to be better in the low than in the high standards cases. Thus, we have a straightforward explanation of the Contextualist Data.

5.2. Objections and Replies

In what follows, I will discuss potential objections to projective invariantism and show how they can be answered. Here is the first objection.

Skepticism

Schaffer and Knobe (2012: 27) object to a view similar to projective invariantism that it would commit the invariantist to “the quasi-skeptical view that in fact Hannah [in our case, Bill] does not know that the bank will be open on Saturday” and suggest “that most invariantists will already reject the strategy at this point.”¹⁰ They would presumably level the same objection to projective invariantism in the following way. To begin with, note that projective invariantism is not supposed to be an error-theory. Thus, according to projective invariantism, both our judgments about the low as well as our judgments about the high standards cases are correct. With respect to the high standards cases, however, we tend to endorse the knowledge denial. So, projective invariantists must say that these denials are correct. Correspondingly, they must say, as Schaffer and Knobe have it, that Bill doesn't know that the bank will be open in the high standards cases. Why would people find this result problematic? The reason, I take it, is this. It may be assumed that, in ordinary conversations, we generally ascribe knowledge to subjects whose epistemic position is not in the least better than Bill's in the high standards cases. So, saying that Bill fails to know that the bank will be open in the high standards cases commits one to the thesis that our ordinary practice of ascribing knowledge is generally mistaken. And this result looks very much like a potentially troubling form of skepticism.

The present objection seems serious. So, let me spend some time on stating how, I think, it can be answered. First of all, I am perfectly happy to accept the intermediate conclusion

¹⁰ Gerken (2011: 54f) also endorses an objection along these lines. This may be part of the reason why he doesn't apply his Principle of Contextual Salience (see FN 3 and 4) in the way I propose to apply it.

that Bill doesn't know that the bank will be open in the high standards cases. As indicated, this is a straightforward consequence of the idea that projective invariantism is not an error-theory. What I deny, however, is that this result leads to a problematic form of skepticism. In particular, I deny that, in ordinary conversations, we generally ascribe knowledge to subjects whose epistemic position is not better than Bill's in the high standards cases.

I reject that premise because I think that it is unfounded. To see this, note that all candidate motivations of this premise are straightforwardly bound to fail (at least in the present dialectical context). The first idea that might come to mind is the following. One could say that, given the Contextualist Data, people tend to accept the knowledge ascription in the low standards cases. Thus, they find it appropriate to ascribe knowledge to subjects whose epistemic position is not better than Bill's in the bank cases. Thus, in ordinary conversations, they would also ascribe knowledge to subjects whose epistemic position is not better than Bill's in the bank cases.

Such a reasoning clearly begs the question against projective invariantism. For, the projective invariantist holds that Bill's epistemic position differs between the low and the high standards cases. So, the only conclusion we can derive from the data about the low standards cases is that, in ordinary conversations, we would ascribe knowledge to subjects whose epistemic position is not better than Bill's *in the low standards cases*. This, however, is a conclusion the projective invariantist can happily endorse. For, the projective invariantist holds that it is perfectly correct to ascribe knowledge to such subjects. She only denies that we can truly ascribe knowledge to subjects whose epistemic position is not better than Bill's in the *high standards cases*. (And we do not accept the knowledge ascription in these cases.)

Here is another way in which one might want to substantiate the idea that we would ordinarily ascribe knowledge to subjects whose epistemic position is not better than Bill's in the high standards cases. Suppose we modify the low standards cases such that Bill's epistemic position equals the epistemic position he supposedly has in in the high standards cases according to projective invariantism. Even so, one might want to argue, we would accept the knowledge ascription in these cases. Thus, we would ordinarily ascribe

knowledge to subjects whose epistemic position is not better than Bill's in the high standards cases after all.

This reasoning is also misguided. It is based on the idea that, even if we modify the low standards bank cases such that Bill's epistemic position equals the epistemic position he supposedly has in in the high standards cases, we would still accept the knowledge ascription in these cases. This idea seems mistaken. In order to modify the low standards cases in the just proposed way, we would have to state in the scenario description that it is likely to a certain degree that the bank changes its hours. After all, this is what makes Bill's epistemic position in the high standards cases fail to amount to knowledge according to projective invariantism. However, it seems quite plausible to me that, if we modify the low standards cases in such a way, people would no longer accept the knowledge ascription. Besides, many theorists in the field should be happy to accept this verdict. For, by pointing out that it is likely to a certain degree that the bank changes its hours, we certainly make this error-possibility salient, if not to Hannah and Sarah, then at least to the reader of the scenario. And many theorists are committed to hold that making an error-possibility salient in this way does have an effect on our judgments.¹¹ So, it is an open question at best whether we would still accept the knowledge ascription in the low standards cases if we modify these cases such that Bill's epistemic position equals the epistemic position he supposedly has in in the high standards cases according to projective invariantism. Correspondingly, the above reasoning is bound to fail as well.

I can see no other way to support the premise that, in ordinary conversations, we would generally ascribe knowledge to subjects whose epistemic position is not better than Bill's in the high standards cases. Thus, I reject that premise. Correspondingly, the above skepticism objection has no force.

¹¹ Admittedly, it would at least be consistent for the contextualist to deny that verdict. For, the contextualist could consistently say that salient error-possibilities affect our judgments only if they are salient *for the speaker in question*, that is, in our case, for Hannah and Sarah, not the reader of the scenario. This response strategy, however, is not available for many other theorists. Proponents of assessment relativism, for example, according to which the truth-value we ascribe to knowledge claims depends on *our* epistemic standard, cannot resort to this response strategy. For, assuming assessment relativism, the truth-value readers of the bank cases ascribe to the knowledge claims in the bank cases should depend only on what is salient to them. Similar considerations hold for psychological error-theories. According to these theories, it is also the readers of the scenario (not the speakers in the scenario) who form mistaken judgments about the bank cases. So, again, all that should count is what is salient to them.

Remark. Even though I want to say that the epistemic position of subjects to which we ordinarily ascribe knowledge generally is better than Bill's in the high standards cases and, thus, that our ordinary knowledge claims are mostly correct, I do not want to deny that they are also oftentimes mistaken. In particular, I think that there are situations like the robin case above where we ascribe knowledge just because we fail to think of likely alternatives. In such cases, the subjects' epistemic position presumably won't be better than Bill's in the high standards cases. Thus, they will fail to know. However, it should be uncontroversial that such cases occur. Hence, it seems safe to assume that the thesis that they occur does not lead to any problematic form of skepticism.

Making the Likelihood Explicit

Here is a second, related objection to projective invariantism. Stanley alludes to this objection when he writes that "even when we have full knowledge of the likelihood of various counter-possibilities, it does not affect our judgments." (Stanley 2005b: 101) As applied to projective invariantism, the objection could be developed as follows. Projective invariantism crucially relies on the fact that the bank cases are underspecified when it comes to how likely it is that the bank changes its hours. For, this allows readers to fill in the low and the high standards cases in different ways. Thus, if projective invariantism holds, the Contextualist Data should disappear once we explicitly stipulate how likely it is that the bank changes its hours in the bank cases. To make the Contextualist Data disappear, it should suffice, for example, to end the high standards cases with a claim to the effect that the likelihood of the possibility that the bank changes its hours is zero. Stanley's quote above can be seen as the objection that this prediction is mistaken.

To respond, I think that Stanley is just wrong about this latter judgment. At the very least, the judgment is unfounded. (Note that Stanley does not back up his claim by any further considerations.) To me, it seems quite plausible that the Contextualist Data would disappear if the scenarios were specified in the mentioned respect. Suppose, for example, that we end the high standards cases with claims along the following lines.

Note that, even though Sarah mentions the possibility that the bank has changed its hours, this possibility is utterly unlikely. The bank hasn't changed its hours before. Banks in Hannah's and Sarah's environment are generally very reliable. And even if

the bank had changed its hours, this would have been properly announced at least two weeks before. So, Bill would most likely have known.

To me at least, it seems not at all clear that, once we add the above paragraph, people would still agree that Bill fails to know that the bank will be open in the high standards cases. Further empirical research may prove me wrong. But I don't think the issue can be settled from the armchair.

Far-Fetched Alternatives

Here is another closely related objection to projective invariantism. One might hold that we would obtain the Contextualist Data even if we let Sarah mention an alternative that we don't consider likely at all. For example, one might hold that the Contextualist Data would still be present if Sarah mentions the possibility, say, that a meteor destroys the bank before Saturday. This result would be problematic because projective invariantists explain the Contextualist Data at least partly on the grounds that we do consider the alternative in the high standards cases likely (and are reminded of that once the possibility is mentioned).

Again, my response is just to deny the alleged data. In fact, I think the real data favor my proposal. Gerken seconds this response in the following passage.

[A]n important consideration concerning very far-fetched salient alternatives cases must be noted. Consider, for example, the alternative that the matter in *S*'s car has spontaneously reorganized in the form of a giant lizard (MacFarlane 2005b). This salient alternative to *S*'s knowledge that the car is in the driveway does not generate strong intuitions to the effect that *S* does not know that the car is in the driveway (MacFarlane 2005b). [FN: We have at present no empirical evidence about far-fetched salient alternatives. But the intuitive judgment appears to be widely agreed upon.] The absence of an intuition in this case—or, at least, the asymmetry in strength of intuition—needs to be explained as much as the presence of intuitions in cases of less far-fetched salient alternatives. (Gerken 2012: 141f)

If Gerken is correct, it is indeed a good thing that projective invariantism does not predict contextualist effects when very far-fetched alternatives are mentioned. For, there simply are no such effects. (Admittedly, Gerken does not directly say that the Contextualist Data are absent when very far-fetched alternatives are mentioned. He leaves open the possibility that the data should only be weaker. But this is unproblematic for the projective invariantist. For, at least some of the mechanisms described above to accommodate the Contextualist Data may apply even when the mentioned error-possibility is very far-fetched. In particular, readers may still ascribe to Sarah an implicature to the effect that

the possibility she mentions is likely, or assume that the possibility is supposed to be likely in the bank cases in order to make sense of why it occurs to Sarah in these cases.)

Skeptical Arguments

The following is yet another objection in the vicinity of the previous considerations. Suppose a skeptic, instead of merely making an error-possibility salient, presents arguments to the conclusion that we don't know that the error-possibility does not obtain and then appeals to epistemic closure to reach the skeptical conclusion that we don't know what we thought we know. Suppose, for example, that the skeptic points out that we don't know that we are not brains in a vat because our sensory experiences would be exactly the same as they are if we were brains in a vat. Suppose further that she now appeals to one or the other epistemic closure principle to conclude that we don't even know we have hands. One might be inclined to think that people tend to respond to such arguments by granting that they don't know they have hands. And it might further be held that projective invariantism is unsuited to explain this datum. In support of this latter contention, Nagel, for example, argues that, when

we succumb to the temptation of rescinding knowledge ascriptions under sceptical pressure, we can at least ostensibly continue to judge the odds of being in a sceptical scenario as very low: the sceptic might move you into suspension of judgement even after granting that the odds of being a BIV are one in a quadrillion, if he is skilled at getting you to appreciate that things would appear *exactly* the same to you under that condition. (Nagel 2010: 299)

I fully agree with Nagel that projective invariantism cannot account for the fact (if it is fact) that we retract knowledge claims in the face of skeptical arguments, at least when these arguments involve far-fetched alternatives such as that of being a brain in a vat. For, as indicated by Nagel, none of the mechanisms posited above to account for the Contextualist Data applies to these reactions. First of all, we will not be reminded that it is likely to a relevant degree that we are brains in vats when this possibility is pointed out to us. After all, we presumably consider this alternative to be extremely unlikely. Relatedly, the skeptic may happily cancel any implicature to the effect that this possibility is likely without undermining her case. Finally, it seems clear that the skeptical error-possibility comes to the mind of the skeptic not because it is likely but because the skeptic has been engaging in philosophy.

But how problematic is the result that projective invariantism cannot explain our assumed reactions to skeptical arguments? I don't think this is very problematic. First of all, it

seems to be very much unclear that we do indeed tend to retract knowledge claims in the face of skeptical arguments. Many, if not most people are entirely unmoved by such considerations.¹² What does seem correct is only that *some* people become convinced by skeptical arguments (some philosophers, for example) and that, for most people, it is difficult to see where exactly skeptical arguments go awry (even if most people think that these arguments just *have to* go awry at some point or other). Admittedly, projective invariantism cannot explain these residual data either. But this result is problematic only if there are no plausible accounts of these data with which projective invariantism could be supplemented. This is not the place to argue that there are such accounts. It seems fairly clear, however, that the burden of proof is on those who want to say that there are none. For example, the projective invariantist could appeal to the strategies mentioned in connection with SPI, according to which skeptical arguments involve one or the other non-obvious argumentative fallacy.¹³

Too Small a Difference

Consider also the following objection. Nagel complains that “it is unclear exactly why a slight increase in the apparent odds of danger in cases like (b) [i.e. cases like the high standards cases] would be sufficient to overturn the positive assessment of cases like (a) [i.e. cases like the low standards cases].” (2010: 299)¹⁴ The idea here seems to be that the slight increase in the likelihood of being wrong posited by projective invariantism would, at best, move us from knowledge to borderline knowledge or from borderline knowledge to non-knowledge but not from knowledge to non-knowledge.

It seems quite unclear, however, why this should be so. For, it is quite unclear why a slight increase in the likelihood of being wrong should only make a slight difference to one’s overall epistemic position. For example, the projective invariantist could hold that, particularly when we go from zero to non-zero, an increase in the likelihood of being wrong makes a large difference to one’s overall epistemic position. Given that, there is a large difference between the subject’s epistemic positions in the bank cases. Thus, there

¹² See (Feldman 2001: 77), (Bach 2010: 108f) and (Davis 2004: 274, 2007: 436) for related observations.

¹³ See p. 101.

¹⁴ For a suggestion that seems to go in a similar direction, see the response by Schaffer to the comment by May on <http://certaindoubts.com/?p=1373>.

is no reason why the subject should not know the bank will be open in one but not the other case.¹⁵

Mere Stakes Effects

The next and final worry has it that we would obtain varying evaluations of the knowledge claims in low and the high standards cases even if we only changed the stakes between the scenarios and not the mentioned error-possibilities.¹⁶ This might again seem to undermine my proposal. For, if no error-possibility is mentioned, we arguably won't be reminded of an error-possibility, the relevant implicature will also not be present and there will be no reason to think that an error-possibility occurred to any of the subjects involved in the bank cases.

Here are my responses to this objection. First, it is far from clear that mere stakes effects cannot be explained on the basis of projective invariantism. Drawing on results from

¹⁵ Full disclosure: In fact, I don't think that the best explanation of the Contextualist Data is that Bill knows the bank will be open in the low but not the high standards cases. Instead, I am inclined to think that the best explanation is that Bill is a borderline case of a knower in all bank cases who is somewhat closer to being a knower in the low than in the high standards cases. This idea would have to be motivated in more detail. But let me briefly indicate the basic thought. It is a striking feature of the empirical results about the bank cases that, at least on average, people judge all the knowledge claims, denials as well as ascriptions, in all the bank cases, high standards as well as low standards, as at least more true than false. (See (Hansen and Chemla 2013: 309) and, relatedly, Figure 2 above on p.59. (Buckwalter 2014) confirms this effect.) Why is that? Doesn't consistency require that we judge at least some of the claims as more false than true? The following seems like a promising account. It is often held that, in the face of borderline cases, "one can go either way," that is, when a is a borderline case of an F , it is allowed to classify a as an F but it is also allowed to classify a as a non- F . See e.g. (Shapiro 2006: 9f). Suppose we grant that thesis. And suppose further that we hold that Bill is a borderline case of a knower in the bank cases. If so, we get a straightforward explanation of the datum that we accept all knowledge claims in all bank cases: We accept all knowledge claims in all bank cases because it is allowed for the relevant speaker to make all these claims. Given that Bill is a borderline knower, how do we explain that there are, nevertheless, differences between our evaluations of the knowledge claims in the low and the high standards cases? Here, we can employ projective invariantism to argue that these differences are due to the fact that Bill is somewhat closer to being a knower in the low than the high standards cases. In sum, we thus arrive at the thesis that Bill is a borderline case of a knower in all bank cases who is somewhat closer to being a knower in the low than in the high standards cases. As I said, the present proposal would have to be further substantiated. In particular, one would have to rule out alternative explanations of the datum that we tend to accept all knowledge claims in the bank cases. And one would have to spell out and motivate the thesis that, in the face of borderline cases, one can go either way. These, however, are topics for another occasion.

¹⁶ That there are mere stakes effects seems to have been confirmed in (May et al. 2010), (Sripada and Stanley 2012) and (Buckwalter and Schaffer 2013). It should be noted, though, that May et al. think that their effects are too weak to underwrite stakes sensitivity. Pinillos (2012) claims to have found very strong stakes effects. But I side with (Buckwalter and Schaffer 2013: Sec. 2.2) and (Buckwalter 2014: Sec. 3.2) in holding that this claim is not well-founded because one plausible explanation of his results is independent of the putative stakes sensitivity of knowledge attributions. Addressing this last issue in more detail would go beyond the scope of this project.

psychology, Buckwalter and Schaffer (2013: 22) argue that high stakes “can be expected to trigger anxiety in the participants [of the respective studies] and thereby ‘to intensify the generation of alternative hypotheses.’” In this way, additional alternatives may come to mind when one reads the high standards cases even if they are not mentioned. Thus, readers may also be reminded that these alternatives are not unlikely. As a result of that, they may then be less prone to ascribe knowledge because the epistemic position of the relevant subject is weakened. Second, in light of the failure of the invariantist theories discussed in the previous chapters, an explanation of the salience effect alone may already be progress enough.¹⁷ So, even if the above account of the stakes effect fails, this would not be a devastating result.

5.3. Conclusion

To sum up, projective invariantism plausibly explains the Contextualist Data. It does so by adopting invariantism and the various arguments to the conclusion that the subject’s epistemic position will be assumed to be different in the low and the high standards cases. Given that all other invariantistic candidate accounts of the Contextualist Data fail, we should thus accept projective invariantism, at least insofar as we are convinced that invariantism is the right way to go.

¹⁷ And, insofar as the contextualist wants to apply her theory to respond to skeptical arguments, the salience effect is of considerably more importance than the stakes effect anyway. See the discussion of the contextualist response to skeptical arguments in the next chapter.

Appendix: The Factivity Problem

In this chapter, I will briefly return to the contextualist account of the Contextualist Data, according to which “knows” expresses differently demanding relations depending on whether the expression is used in the low or the high standards cases. I have claimed in the introduction that the basic problem for contextualism is to provide an account of the Invariantist Data, that is, the data from speech reports and disagreement. This assumption has figured in the background in my discussion of relativism. For, I have argued there that the relativist cannot accept certain kinds of responses to my objection from indirect testimony unless she is prepared to give up at least some aspects of the objection from the Invariantist Data to contextualism. This point would be weakened if the Invariantist Data turned out not to be crucial for the case against contextualism anyway. So, in this last chapter, I will try to corroborate the assumption that they are by showing that another prominent objection to contextualism fails.

It has been argued that contextualism faces the so-called factivity problem and hence cannot be stated properly.¹ The basic idea behind this charge is that contextualists supposedly have to say, on the one hand, that knowledge ascribing sentences like “*S* knows that *S* has hands” are true when used in ordinary contexts while, on the other hand, they are not true by the standard of their own context. In this chapter, I will show that the argument to the factivity problem fails because it rests on the mistaken premise that contextualists are committed to the truth of particular ordinary knowledge attributions.

Let me start by presenting the factivity problem. To begin with, let me clarify what exactly the factivity problem is supposed to show. The factivity problem, I take it, is an argument to the conclusion that anti-skeptical epistemic contextualism cannot be stated properly if it is conjoined with contextualistic versions of the principles (KNA) that one can properly assert that *p* only if one knows that *p*, (Fact) that knowledge that *p* entails that *p* and (Clos) that we know the things that follow from what we know so long as we know that they

¹ See e.g. (Williamson 2001; 2005b), (Wright 2005), (Brendel 2005; 2009), (Baumann 2008; 2010) and (Jäger 2012).

follow from what we know. (KNA) is the knowledge norm of assertion, (Fact) is the factivity principle for knowledge and (Clos) is the principle of epistemic closure.

As stated already, epistemic contextualism roughly is the semantic thesis that “knows” expresses more or less demanding epistemic relations depending on the epistemic standard in play at the context of its use. Anti-skeptical epistemic contextualism, as I will understand it here, entails the following two additional claims. First, it entails that skepticism is false. Skepticism is most commonly construed as the thesis that “[o]rdinary knowledge attributions are generally false.” (Stanley 2005b: 82)² Hence, anti-skeptical epistemic contextualism entails that it is not the case that ordinary knowledge ascriptions are generally false. And I will assume here that this, in turn, means that anti-skeptical contextualism entails that ordinary knowledge ascriptions are generally true.³ Second, anti-skeptical epistemic contextualism entails a particular response to a particular form of skeptical argument, a response that typical epistemic contextualists endorse. Let me elaborate.

The skeptical argument I have in mind is an argument to the conclusion that skepticism is true and goes, very roughly, as follows: You don’t know that you are not a (just recently envatted) handless brain in a vat. If you don’t know that, then you also don’t know that you have hands. So, you don’t know that you have hands. Analogous arguments show that, for most ordinary claims p , we don’t know that p . If we don’t know all these things, then ordinary knowledge ascriptions are generally false. So, ordinary knowledge ascriptions are generally false. Hence, skepticism is true. (Note here that ordinary claims are supposed to be those claims that are said to be “known” in ordinary contexts.)

The anti-skeptical contextualist responds to that argument not only by pointing out why the argument is not sound but also by explaining why it seems sound to begin with. The argument is not sound, according to the anti-skeptical contextualist, because it is not true that ordinary knowledge ascriptions are generally false if we don’t know that p for most

² See also (Hawthorne 2004a: 53), (Davis 2007: 427), (Rysiew 2007: 627) and (DeRose 2009: 124).

³ Freitag (2011: 277n; 2013: 135) tries to defuse the factivity problem by claiming that the skeptic says that ordinary knowledge attribution are *necessarily* false, not that they are false as a matter of contingent fact. Correspondingly, on Freitag’s view, the anti-skeptical contextualist only has to say that ordinary knowledge ascriptions *can* be true, and this claim does not give rise to the factivity problem (as will become apparent below). Freitag’s definition of skepticism, however, goes against the grain of the contemporary discussion (see FN 2). That is why, I think, his response to the factivity problem is not fully satisfactory.

ordinary propositions p . The reason is that once the skeptical argument is presented and the corresponding error-possibility is made salient, our epistemic standard rises such that, in the scope of the argument, “knows” comes to express the very demanding relation of knowledge_S. In ordinary discourse, however, “knows” expresses the far less demanding relation of knowledge_O. And even if we don’t know_S a lot, it does not follow that we also don’t know_O a lot.⁴ The argument nevertheless *seems* sound because the inference to the claim that we don’t know_S most ordinary claims *is* perfectly sound and because “we fail to realize [...] that the skeptic’s present denials that we know various things are perfectly compatible with our ordinary claims to know those very propositions.” (DeRose 1995: 5)⁵

So, to repeat, the anti-skeptical epistemic contextualist endorses epistemic contextualism, anti-skepticism and the just sketched response to the skeptical argument. The factivity problem is not an argument against epistemic contextualism *per se* but an argument against anti-skeptical epistemic contextualism conjoined with contextualistic version of (KNA), (Fact) and (Clos).

It should be noted that, even so, the factivity problem puts considerable pressure on epistemic contextualism *per se*: The factivity problem supposedly shows that, in order to have a view that can be stated properly, an epistemic contextualist must either reject anti-skepticism or her response to the skeptical argument or give up at least one of the principles (KNA), (Fact) and (Clos). But each of these disjuncts comes at a price. First, epistemic contextualists typically hold that one major merit of their theory is that it provides a response to skepticism. Correspondingly, they can give up anti-skepticism and their response to the skeptical argument only to the extent that they are prepared to give up what they take to be a major merit of their theory.⁶ Similarly, second, (KNA) plays a crucial role in an important argument for epistemic contextualism. This argument has to be given up if (KNA) is given up.⁷ Third, (Fact) seems intuitively very plausible. So, quite

⁴ To simplify the discussion, I pretend that there are only two relations “knows” can express, the relation of knowledge_O and the relation of knowledge_S. The contextualist will actually say that there are many more such relations.

⁵ It is, of course, debatable whether the contextualist response to the skeptical argument can ultimately be sustained. See e.g. (Schiffer 1996) for further discussion.

⁶ This point is emphasized most strongly in Jäger’s (2012: 494) recent presentation of the factivity problem. See also (Ashfield 2013: 121).

⁷ See, again, (Jäger 2012: 494) for this point. See (DeRose 2002) for the argument in question.

generally, it would seem to be a drawback of a theory if it can be maintained only if this principle is rejected. Finally, even if it may be disputable whether (Clos) holds in full generality, there should still be suitably restricted versions of that principle that are, on the one hand, true and, on the other, strong enough to underwrite the subsequent arguments. So, again, the contextualist would be ill-advised to reject this principle. Discussing in detail the issues mentioned in this paragraph is a topic for another occasion.

To save some breath, I will, in what follows, use “contextualism” to refer to anti-skeptical epistemic contextualism conjoined with contextualistic versions of (KNA), (Fact) and (Clos) (unless indicated otherwise). The factivity problem, then, is an argument to the conclusion that contextualism, in this sense, cannot be stated properly.⁸

Let me now show how the factivity problem is supposed to arise. As a first step, we have to spell out the contextualistic versions of (KNA), (Fact) and (Clos). Let’s start with (KNA). A contextualist will not say that one can properly assert that p only if one knows

⁸ Setting up the factivity problem in the way I did indicates two drawbacks of at least the presentation of the response to the factivity problem given in (Brueckner and Buford 2009; 2010). First, Brueckner and Buford seem to endorse what I take to be the conclusion of this argument. They write that a contextualist “cannot ‘knowledgably’ state the contextualist thesis that ‘[S knows that h at t]’ is true in O .” (2009: 436) We need not bother here about the specific thesis contextualists supposedly cannot “knowledgably” state. What matters for our purposes is only that this thesis is claimed to be a “contextualist” thesis. It is not entirely clear what this is supposed to mean. One plausible interpretation, though, is that a contextualist thesis is a thesis that is, in some sense, part of contextualism. If that is so, Brueckner and Buford are committed to hold that contextualists cannot “knowledgably” state a claim that is part of contextualism and, as a consequence, cannot “knowledgably” state contextualism itself. Given (KNA), this just means that contextualism cannot be stated properly and hence entails the conclusion of the factivity problem. (I take this to be the bottom line of why Jäger (2012: 496) thinks that Brueckner and Buford’s response to the factivity problem can be dismissed.) Brueckner and Buford may object that I am misconstruing their idea of a “contextualist thesis.” A contextualist thesis need not be a thesis that is part of contextualism; it could also be a thesis that, say, just involves contextualistic vocabulary. Similarly, they may object that the claim quoted above is immaterial to their overall argument and that retracting it wouldn’t do harm to their response to the factivity problem. Even so, the above considerations show that, *as it stands*, Brueckner and Buford’s response to the factivity problem is not fully satisfactory because it has an obvious reading according to which it entails the conclusion of the factivity problem. Second, Brueckner and Buford mainly argue that the contextualist can accept the “contextualist dogma” that it “is possible for there to be two contexts $C1$ and $C2$ such that a sentence of the form ‘[X knows that p at t]’ is true in $C1$ and *not* true in $C2$ ” (2009: 434, 437) without being committed to the factivity problem. But this dogma seems to be just a statement of contextualism *per se*. If that is so, their argument is beside the point. Even proponents of the factivity problem do not (or at least should not) claim that contextualism *per se* cannot be stated properly. Only anti-skeptical contextualism (conjoined with the above general principles) is supposed to be unstable. (I take this to be the bottom line of why Baumann (2010: 85, 87) thinks that their argument can be dismissed.) Now, it may be that Brueckner and Buford’s dogma is supposed to already entail a certain form of anti-skepticism. Similarly, it may be that, despite the fact that they do not explicitly consider this issue, their arguments show that the contextualist can properly accept not only contextualism *per se* but also anti-skepticism. But be that as it may, a fully satisfactory response to the factivity problem should make these issues explicit.

that p . For this would mean that assertions are governed by whatever it is that “knows” expresses in the context of this article. This certainly seems arbitrary. The contextualist will correspondingly say that one can properly assert the sentence “ p ” in a context C only if “ S knows that p ” is true in C . Descending to the object language, the contextualist will hold that for every C :

(KNA_C) An assertion that p by a speaker S is proper in a context C only if S knows_C that p ,

where “knows_C” is the relation that, according to contextualism, “knows” expresses in a context C . For analogous reasons, the contextualist will spell out (Fact) and (Clos) as the claims that for every C :

(Fact_C) If S knows_C that p , then p .

(Clos_C) If S knows_C that p and S knows_C that if p , then q , then S knows_C that q .

Proponents of the factivity problem use these principles to show that contextualism cannot be stated properly in the following way. By definition, contextualism entails anti-skepticism, that is, that ordinary knowledge ascriptions are generally true. Let’s introduce “ S ” as a name for a contextualist, say, Keith DeRose, and “ h ” as an abbreviation for “ S has hands.” According to proponents of the factivity problem, the contextualist commitment to anti-skepticism then entails that “ S knows that h ” is true when uttered in an ordinary context. Correspondingly, the first premise of the factivity problem is that contextualists are committed to the claim that “ S knows that h ” is true in an ordinary context.⁹ Let’s use “knows_O” as a term for the relation that, according to contextualism, “knows” expresses in an ordinary context (as we already did above). We can then restate the first premise of the factivity problem as follows: According to contextualism,

⁹ The above reasoning to the first premise of the factivity problem seems to be present in most presentations of the factivity problem. Baumann (2008: 583; 2010: 87), for example, claims that contextualism without the claim that “knowledge attributions made in a lower context are in fact true” would be “not very attractive [...]” Wright (2005: 243) concurs that the contextualist “has no point to make” unless she can accept that “common sense’s claims to knowledge are, in their proper context, perfectly good [...]” Jäger (2012: 491f) quotes DeRose in saying that we do not speak falsely “when we claim to know [...] in [ordinary] contexts in which no skeptics are calling our attention to [skeptical] hypotheses” (DeRose 1999: 18). And Williamson states that according to contextualism, “in everyday contexts many assertions of ‘knowledge’ are [...] correct.” (2005b: 689) All four authors use these claims to straightforwardly justify the premise that contextualists must concede that “ S knows that h ” is true in an ordinary context.

(LOW_C) *S* knows_O that *h*.¹⁰

I will argue later in this chapter that this first premise cannot be derived from the contextualist commitment to anti-skepticism in the way just presented and that, as a result, we should reject the factivity problem. For now, though, I will take it for granted.

The second premise of the factivity problem straightforwardly follows from the contextualist response to the skeptical argument. In response to this argument, our contextualist conceded that we don't know_S most ordinary propositions. In particular, she conceded that we don't know_S that we have hands. Correspondingly, the second premise of the factivity problem is that

(HIGH_C) *S* doesn't know_S that *h*.

The factivity problem can now be derived as follows: According to the first premise, proponents of contextualism have to endorse (LOW_C). So, whenever a contextualist properly asserts her view, she also properly asserts (LOW_C). Assume *S* is a contextualist who properly asserts her position in a skeptical context. Then the following holds true:

(1) *S* properly asserts in a skeptical context that (LOW_C) i.e. that *S* knows_O that *h*.

This claim, together with (KNA_C), straightforwardly entails that

(2) *S* knows_S that *S* knows_O that *h*.

(KNA_C) analogously yields that

(3) *S* knows_S that (Fact_C) i.e. that if *S* knows_C that *p*, then *p*,

¹⁰ Note that it is not entirely obvious that (LOW_C) follows from the thesis that “*S* knows that *h*” is true in an ordinary context. For contexts are often construed as *possible* situations of the utterance of a sentence. And even if there are *possible* ordinary situations in which “*S* knows that *h*” is true, it does not follow that *S* knows_O that *h*. All that would follow is that *possibly S* knows_O that *h*. So, contexts must be seen here as (something like) *actual* situations of the utterance of a sentence. I mention these points because confusion about these two notions of a context may easily make it seem that (LOW_C) is just a trivial consequence of contextualism: Contextualists clearly accept that there are ordinary (possible) contexts such that “*S* knows that *h*” is true. It is also trivial that if there are ordinary (actual) contexts such that “*S* knows that *h*” is true, then *S* knows_O that *h*. But, contrary to what may seem to be the case if the two notions of a context aren't kept apart, these two claims cannot be put together to yield that contextualists trivially accept that *S* knows_O that *h*. I don't want to claim that proponents of the factivity problem did confuse these notions of a context. Certainly, though, a lax use of “context” can easily make their arguments seem more attractive than they are.

because, just as they endorse (LOW_C), contextualists endorse ($Fact_C$). But if we now combine (2), (3) and ($Clos_C$), we arrive at the conclusion that

(4) *S* knows_S that *h*.

And this claim contradicts ($HIGH_C$). So, contextualists cannot properly state their theory in skeptical contexts. However, whenever they do state their theory they find themselves in “paradigmatic epistemology or ‘philosophy classroom’ contexts in which skeptical hypotheses are salient and taken seriously.” (Jäger 2012: 492) These contexts arguably are skeptical contexts. So, contextualism cannot be stated properly at all.¹¹

In what follows, I will show that the argument to this conclusion is not sound because its first premise, that contextualists are committed to (LOW_C), is false. Note, first, that this does not mean that (LOW_C) itself false and, second, that whether (LOW_C) is true or false is an issue that will not be addressed in this chapter.

To show that the contextualist is not committed to (LOW_C), we have to show that none of the claims that make up contextualism carries with it a commitment to this claim. This is what I will be doing in what follows. Let’s start by considering the commitment to anti-skepticism which, as argued above, is standardly taken to entail the commitment to (LOW_C). Here is why it doesn’t. It seems clear that even if anti-skepticism is true, and thus that ordinary knowledge attributions are generally true, it simply does not follow that (LOW_C) is true. For it is certainly possible that all everyday utterances of the specific sentence “*S* knows that *h*” are false even if ordinary knowledge ascriptions are generally true. But if anti-skepticism doesn’t entail (LOW_C), then it is perfectly consistent for the contextualist to accept anti-skepticism without accepting (LOW_C). Hence, the contextualist is not committed to (LOW_C) just in virtue of her commitment to anti-skepticism. Of course, the contextualist would be committed to (LOW_C) if she held that

¹¹ There is room for doubt as to whether philosophical contexts (or at least those philosophical contexts where contextualism is stated) need to be skeptical contexts and, indeed, I am not unsympathetic to the idea of rejecting the factivity problem on the basis of the claim that they are not. In this chapter, however, I want to attack the factivity problem from a different, so far un(der)explored angle and hence take this point for granted. For arguments to the conclusion that philosophical contexts aren’t skeptical, see e.g. (Montminy 2008: 2–6) and (Ashfield 2013: Sec. 8). For arguments to the conclusion that they are, see (Jäger 2012: 496f). Note also that Baumann (2008: 582f) argues that the factivity problem arises already if the contextualist tries to state her position in a demanding but not necessarily skeptical context. If that is so, the factivity problem requires only the weaker (and presumably indisputable) assumption that philosophical contexts are demanding contexts. All these issues though go beyond the scope of the present discussion.

ordinary knowledge attributions are *one and all* true. But that would be absurd. We do make mistakes. What the contextualist accepts is that ordinary knowledge ascriptions are *generally* true. And this claim does not entail (LOW_C) and thus does not yield a commitment to this claim. Note that this result does not rely on my particular choice of *h* as the thesis that *S* has hands. “*h*” can stand for whatever ordinary claim we like. The contextualist will not be committed to the truth of this particular claim just in virtue of her commitment to anti-skepticism. For the claim that ordinary knowledge attributions are generally true does not entail the truth of any particular ordinary knowledge attribution.

One might want to respond to these considerations by arguing that the contextualist is committed to (LOW_C) not because anti-skepticism entails (LOW_C) but because *knowledges* of anti-skepticism entails *knowledges* of (LOW_C). Such an argument, however, would be bound to fail because it is far from clear that *knowledges* of anti-skepticism should entail *knowledges* of (LOW_C) (or, for that matter, of the truth of any particular ordinary knowledge attribution). For example, I may perfectly well know that the tickets of a yet to be drawn lottery will generally lose even if I don’t know of any particular ticket that it will lose. In just the same way, it should be perfectly possible to know_S that anti-skepticism is true, that is, that ordinary knowledge attributions are generally true, even if one does not know_S of any particular ordinary knowledge attribution that it is true.

These considerations may leave one wonder what exactly the contextualist motivation for anti-skepticism is going to be if it is not the truth of particular ordinary knowledge claims. After all, whereas, in the lottery case, we normally have a guarantee that there is only one winning ticket, we certainly have no guarantee that there is only one false ordinary knowledge claim. This question goes beyond the scope of the present discussion, but let me briefly indicate one such candidate motivation. DeRose, for example, accepts a principle of charity according to which it is “a strike against a theory of a common term of a natural language that it involves the speakers of that language in systematic and widespread falsehood in their use of that term.” (1995: 46) This principle, maybe together with the contextualist response to the skeptical argument above, may well be used in an argument for anti-skepticism that is entirely independent of the truth of particular ordinary

knowledge claims. It might be debatable whether such appeals to charity are any good.¹² But in light of the fact that leading proponents of contextualism do invoke such principles, they can certainly not be rejected out of hand in an argument supposed to show that contextualism cannot be stated properly. So, the present discussion may show that one way to corroborate the factivity problem would be to demonstrate that every possible justification of anti-skepticism appeals to the truth of particular ordinary knowledge attributions. Pending such demonstration, however, the factivity problem can be rejected. (In the foregoing paragraphs, I presupposed that knowledges of anti-skepticism is possible. The question under discussion only was whether such knowledges is possible *independently of knowledges of the truth of particular ordinary knowledge claims*. However, this presupposition might be doubted. This, in turn, might then be thought to make contextualism unstatable. I will address this worry in the final sections of this chapter.)

The analogy to the lottery case may also leave one wonder how exactly the contextualist will motivate that she doesn't know_S that (LOW_C) obtains. After all, it would be question begging just to respond that this assumption leads to the factivity problem. What the contextualist should say here is this: I fail to know_S that (LOW_C) obtains, that is, that *S* know_O that *h*, simply because knowledge_O is factive and because I don't know_S that *h*. This reasoning would not at all be question begging. For it only relies on the factivity principle and the assumption that the contextualist doesn't know_S a lot. And the proponent of the factivity problem uses these principles in her very own argument.

So much for the contextualist commitment to anti-skepticism. Is there any other ingredient of contextualism that might lead to a commitment to (LOW_C)? Consider contextualism *per se*. It might be thought that this thesis leads to (LOW_C) for the following reason. In explaining contextualism *per se*, contextualists typically present scenarios like the bank cases and then say of some of the knowledge ascriptions and denials made in these scenarios that they are true.¹³ Maybe this is even the only way to properly explain contextualism *per se*. So, one might want to conclude that contextualists are committed to say of the particular ordinary knowledge attributions that occur in bank

¹² See e.g. (Hawthorne 2004a: 126–133), (MacFarlane 2005b: 211f) and (Davis 2007: 428–430) for further discussion.

¹³ See e.g. (DeRose 2009: 1f).

case like scenarios that they are true; and, hence, that they are committed to the truth of particular ordinary knowledge attributions after all. (LOW_C) will be just a placeholder for one of these claims.

This reasoning would be mistaken for the following reason. Saying of imagined assertions in imagined scenarios that they are true is entirely unproblematic as far as the factivity problem is concerned. For example, the contextualist in a skeptical context may perfectly happily say of Hannah's low stakes assertion of "Bill know that the bank will be open" that it is true. For this just means that the contextualist must know_S that Bill knows_O that the bank will be open *in the low stakes scenario*. Given the factivity principle, it does follow that the contextualist must also know_S that the bank will be open *in the low stakes scenario*. But since the low stakes scenario is just defined such that the bank will be open, such knowledges_S seems entirely unproblematic. So, contextualism *per se* does not entail the truth of any real life ordinary knowledge attribution and hence cannot be used to underwrite the supposed commitment to (LOW_C).¹⁴

Consider then the contextualist response to the skeptical argument. Does that response entail a commitment to (LOW_C)? It seems fairly clear that it doesn't. In response to the skeptical argument, the contextualist did say that *we cannot show that "S knows that h" is false in an ordinary context* by showing that it is false in a skeptical context. However, this is not to endorse (LOW_C), that is, that *"S knows that h" is true in an ordinary context*. As for (Clos_C), (Fact_C) and (KNA_C), these principles are obviously not involved in generating a commitment to (LOW_C). So, in sum, no ingredient of the contextualist position commits the contextualists to (LOW_C). The factivity problem should thus be rejected.

To clarify the dialectical situation here, let me emphasize that I am perfectly happy with the conclusion that, in a skeptical context, there is hardly any ordinary knowledge attribution such that the contextualist could properly say of that knowledge attribution that it is true. We can easily establish this result along the lines of the reasoning subsequent to (1) above. What I am saying is that this just does nothing to establish the conclusion of the factivity problem that contextualism cannot be stated properly. For the

¹⁴ See FN 10 for a related point.

contextualist just isn't committed to say of any particular ordinary knowledge attribution that it is true.

Three further clarificatory remarks: First, when I say that in a skeptical context there is hardly any ordinary knowledge attribution such that the contextualist could properly say of that knowledge attribution that it is true, this is not to say that the contextualist should *deny* the truth of most particular ordinary knowledge attributions in these contexts. This would certainly be incompatible with accepting anti-skepticism. When asked in a skeptical context whether a particular ordinary knowledge attribution is true, the contextualist's response should generally be "I don't know (because I don't know that the putatively "known" content is true)." (Remember the discussion of the lottery example to see that this does not preclude the contextualist from properly saying that anti-skepticism is true.) Second, the previous comment is not to be understood such that the contextualist in a skeptical context should say things like "I don't know whether those ordinary people know that I have hands." This would implausibly suggest that the contextualist should generally consider it an open question whether ordinary speakers meet the excessive epistemic demands of the present skeptical context. All that the contextualist should say is "I don't know whether this particular ordinary standards usage of 'I know that the contextualist has hands' is true." This claim is perfectly compatible with the assumption that the ordinary speaker in question does not meet the skeptical epistemic demands. It only suggests that it is an open question whether she meets the far less demanding epistemic standard that governed her own assertion. And this latter contention seems entirely unproblematic. Third, even if there is hardly any ordinary knowledge attribution such that the contextualist in a skeptical context could properly say of that knowledge attribution that it is true, the same is not true when the contextualist finds herself in an ordinary context. In such a context the contextualist can properly say of an ordinary knowledge claim that it is true as long as she knows_O that the relevant speaker knows_O the propositions she professes to "know" (which will often be the case). Note though that the contextualist will have to retract such claims as soon as she finds herself in a skeptical context. For if she were to retain a commitment to them, she could no longer properly state her views, as the factivity problem shows.

In the remainder of this chapter, I will discuss and reject two revised versions of the factivity problem that might have come to mind already in the course of my discussion of the original factivity problem. Here is the first. Given that our contextualist accepts that ordinary knowledge ascriptions are generally true, she will also have to accept that, for some ordinary sentence p , there are ordinary contexts in which “ S knows that p ” is true. Descending to the object language, the contextualist has to say that there are ordinary propositions p such that S knows_O that p . Given the contextualist response to the skeptical argument, the contextualist will even have to accept that there are ordinary propositions p such that S knows_O that p and S doesn’t know_S that p . Now, let’s introduce “ h ” as an arbitrary name for one of the propositions that S knows_O but doesn’t know_S. We then arrive at the conclusion—which strongly resembles (LOW_C) and (HIGH_C)—that, according to the contextualist, S knows_O that h while S doesn’t know_S that h . Suppose now that the contextualist finds herself in a skeptical context. Then, given (KNA_C), she can properly assert the just derived commitment only if, among other things, she knows_S that S knows_O that h . Given (Fact_C) and (Clos_C), she must thus know_S that h . But, by assumption, she doesn’t. Hence, contextualism cannot be stated properly (in a skeptical context).

Here is my response to this revised factivity problem. I have nothing to object to the reasoning that leads to the conclusion that the contextualist has to accept that there are ordinary propositions p such that S knows_O that p and S doesn’t know_S that p . The rest of the argument, however, rests on a logical mistake as the following analogy brings out. Consider the arguably plausible view that sentences express different propositions depending on the time at which they are uttered. One way of specifying this view would be to assign different properties to each predicate depending on the time of its utterance. We would then presumably say that “knows” expresses different relations depending on the time at which it is used. It could express e.g. the relations of knowledge_{t1}, knowledge_{t2}, etc. Assume that this view is true. Now suppose that I want to say that there are things that I have forgotten. I will then be committed to the claim that there are proposition p such that I know_{tpast} that p but don’t know_{tnow} that p . Certainly, though, it would be absurd to present the following argument to the conclusion that I cannot properly assert that I have forgotten something: Let’s introduce “ h ” as an arbitrary name for one of the propositions that I know_{tpast} but don’t know_{tnow}. We then arrive at the claim that I am

committed to hold that I know_{tpast} that *h* and don't know_{tnow} that *h*. Given the knowledge norm of assertion, I must thus know_{tnow} that I know_{tpast} that *h* (given that I *now* want to say that I have forgotten something). Given the factivity of knowledge and epistemic closure, I must thus also know_{tnow} that *h*. But this contradicts the assumption that I don't know_{tnow} that *h*. So, I cannot properly say that I have forgotten something. Clearly, we cannot accept that reasoning if just because its conclusion is entirely absurd: We *can* properly say that we have forgotten something. But whatever mistake it is that afflicts the argument, this mistake will carry over to the completely analogous argument to the factivity problem described above. So, both these arguments should be rejected.¹⁵

Here is a second, very much neglected variation on the factivity problem that is taken from (Williamson 2001: 26).¹⁶ It is based on the already mentioned idea that the contextualist might not be able to properly state her view because she doesn't know_S that anti-skepticism is true. Let me elaborate. Instead of starting with the assumption that the contextualist is committed to (LOW_C), Williamson starts with the premise that the contextualist is committed to the following claim:

(LOW_C*) Ordinary knowledge attributions are generally true.

In the present context, there is certainly no doubt that this premise is true. For (LOW_C*) is just a statement of anti-skepticism and, on the present understanding, contextualism entails anti-skepticism. Next, Williamson introduces the following premise instead of (HIGH_C):

(HIGH_C*) *S* doesn't know_S that ordinary claims are generally true.

Williamson does not provide a reason to accept this premise and I will argue later that there is no such reason. But if we accept it, we can derive that the contextualist cannot properly state her view in the following way. Given (KNA_C), the contextualist must

¹⁵ I do not want to commit myself to saying where exactly the arguments go awry because this seems to be a tricky logical issue. One possible source of the wrongness of the arguments could be that “*h*,” as introduced, is ambiguous between what we could call a referential and a descriptive reading. On the referential reading, I don't know_{tnow/S} that *h* because I do not stand in the knowledge_{tnow/S} relation to the proposition denoted by “*h*.” On the descriptive reading I do know_{tnow/S} that *h* because I do know_{tnow/S} that whatever proposition “*h*” picks out, this proposition is true because it is supposed to be a proposition I know_{tpast/O} and hence, given factivity, a true proposition. Thanks to Bruno Jacinto for helpful discussion.

¹⁶ Kompa (2005: 18f) endorses the argument but doesn't add any substance.

knows that (LOW_C^*), that is, that ordinary knowledge ascriptions are generally true in order to properly state her position in a skeptical context. Given ($Fact_C$) and ($Clos_C$), she must then also know that ordinary claims are generally true in order to properly state her position in a skeptical context. But given ($HIGH_C^*$), she doesn't have that knowledge. So, given that philosophical contexts are skeptical contexts, it follows that contextualism cannot be stated properly.

Why should we accept ($HIGH_C^*$)? Why should the contextualist be forced to construe knowledges such that it turns out impossible for her to know that ordinary claims are generally true (setting aside the already discussed issue that the contextualist cannot appeal to the truth of particular ordinary knowledge claims)? In what follows, I will lay out what I take to be the most promising answer to these questions and show that it is mistaken. On that basis, I conclude that ($HIGH_C^*$) should be rejected. Here is the answer I have in mind.

In her response to the skeptical argument, the contextualist conceded the soundness in a skeptical context of arguments like the following, call them "Conceded:"

You don't know that you are not a (just recently envatted) handless brain in a vat.
If you don't know that, then you also don't know that you have hands.
So, you don't know that you have hands.

Given that, it might be held, the contextualist should also grant the soundness in a skeptical context of the following seemingly very similar argument, call it "General:"

You don't know that it is not the case that we all are and have always been brains in vats.
If you don't know that, then you also don't know that ordinary claims are generally true (because if we all are and have always been brains in vats, then ordinary claims are generally false).
So, you don't know that ordinary claims are generally true.

But if that argument is sound in a skeptical context, then the contextualist doesn't know that ordinary claims are generally true, and that is precisely what ($HIGH_C^*$) says.

The premise in the foregoing paragraph that the contextualist must grant the soundness in a skeptical context of General if she grants the soundness in a skeptical context of the Conceded arguments is far from trivial. After all, even if these arguments are similar, this does not entail that they should be treated alike in every respect. So, it will have to be supported by a further argument. The argument adduced here would presumably go as follows: If the contextualist wants to hold that General is not sound in a skeptical context, she needs to provide a reason why the argument is not sound in such a context. But whatever reason the contextualist puts forward here, this reason will also show that the Conceded arguments are not sound in a skeptical context either. So, the contextualist will have to grant that General is sound in a skeptical context if she grants that the Conceded arguments are.

The just stated argument for (HIGH_C*) fails because it far from clear that every reason the contextualist could have for rejecting the soundness in a skeptical context of General will also be a reason to reject the soundness in a skeptical context of the Conceded arguments. For example, the contextualist could reject General on the basis of semantic externalism. The idea here would be that if we all are and have always been brains in vats, the meanings of our words adapt such that ordinary claims turn out to be true after all.¹⁷ Such an approach would not carry over to the Conceded arguments. For these arguments turn on scenarios like that of being an *only recently envatted* handless brain in a vat. And such more local scenarios certainly do not give rise to meaning adaptation.¹⁸ To be clear, I am not endorsing the semantic externalist response here. I just want to bring out that it would be unfounded to assume that the contextualist cannot argue that General is not sound in a skeptical context without undermining her response to the Conceded arguments. Further considerations would be required to make that contention plausible, but, as far as I can see, such considerations have yet to be delivered.

Admittedly, even if the contextualist manages to show that General is not sound in a skeptical context without thereby undermining her response to the Conceded arguments, she will still make contextualism obsolete as far as responding to General itself is concerned. For a contextualist response to that argument would start from the assumption

¹⁷ This line of argument is, of course, inspired by Putnam. For helpful discussion, see (Brueckner 2012).

¹⁸ See e.g. (DeRose 1999) for this last observation.

that the argument *is* sound in a skeptical context. This, however, doesn't seem to be much of a drawback. For, it simply might well be that different skeptical arguments require different solutions and that contextualism is only part of the overall response to skepticism.

Let me emphasize once more that I am not endorsing the above response to General. I only want to show that, so far, we have no compelling argument for (HIGH_C*). Again, even if the contextualist grants the soundness in a skeptical context of the Conceded arguments, it just doesn't follow that she also has to grant the soundness in a skeptical context of General. To vindicate this transition, the above mentioned response to General would have to be ruled out, and at least a case would have to be made that there is no possible argument that could plausibly play its part. So far, however, this just hasn't been done.

This is not to say that it can't be done. So, let me conclude this chapter by showing that there still is room for maneuvering even if it turns out that the contextualist has to accept General if she accepts the Conceded arguments. This latter conditional leads to trouble only if its antecedent is true, that is, only if the contextualist has to accept the Conceded arguments. In the present chapter, contextualism has just been defined such that this assumption is unproblematic. However, as I will show in what follows, this definition might not be the only sensible option.

Note, to begin with, that it can easily seem that the skeptic, in presenting the above mentioned Conceded argument, uses the sledgehammer to break a walnut. As has frequently been observed, we need not appeal to alternatives as far-fetched as that of being a brain in a vat to challenge most ordinary knowledge claims. So-called "lottery propositions" will generally suffice.¹⁹ For example, the skeptic could argue that I don't know that my bike is parked outside because I don't know that it hasn't been stolen; I don't know that the bank will be open tomorrow because I don't know that it hasn't changed its hours; I don't know that I will never be rich because I don't know that I will not win the lottery; etc. It would seem that such consideration could be leveled against almost all ordinary knowledge claims. So, in this way, the skeptic could argue for her

¹⁹ See e.g. (Vogel 1990: 20f) and, following Vogel, (Hawthorne 2004a: 5f).

position independently of any appeal to far-fetched alternatives such as that of being a brain in a vat.

Now, contextualism could be construed as a theory designed to respond to precisely this kind of what we may call “lottery skepticism,” that is, skepticism supported by the just mentioned lottery considerations. As before, the contextualist would admit that the individual arguments the lottery skeptic presents are sound but object that they only show that we don’t know_S a lot but not that we don’t know_O a lot. The contextualist, however, would no longer apply this strategy to what we may call “brain in a vat skepticism,” that is, skepticism supported by considerations involving brains in vats and similarly far-fetched possibilities. Correspondingly, she would no longer apply her strategy to the above mentioned Conceded argument. Assuming this conception of contextualism, there is no longer any pressure for the contextualist to concede the soundness in a skeptical context of General. For she could simply claim that this argument does not fall within the purview of her theory because it involves brains in vats.²⁰

Would this differential treatment of lottery and brain in a vat skepticism be *ad hoc*? Would it be too much of a concession to admit that contextualism does not provide a response to brain in a vat skepticism? Plausibly not. Hawthorne:

Indeed, lottery-style considerations are arguably a more dialectically effective tool for the skeptic than standard brain in a vat or deceiving demon thought experiments. While many contemporary philosophers are inclined to resist the skeptic by claiming that they can, after all, know that they are not brains in vats, [...] they are not nearly so eager to embrace the claim that they know they will lose a lottery for which they hold a lottery ticket. And once this has been conceded, it is extremely hard to justify a different attitude to the other ‘lottery propositions’ that figure in the above examples. (Hawthorne 2004a: 6)

To the extent that Hawthorne is correct, the just outlined response would neither be *ad hoc* nor, arguably, involve too much of a concession. It wouldn’t be *ad hoc* because the differential treatment of lottery and brain in a vat skepticism would just line up with the differential reactions to these arguments of “many contemporary philosophers.” It would arguably also not involve too much of a concession. For, Hawthorne plausibly suggests that the strongest case for skepticism is not made by the brain in a vat skeptic but by the

²⁰ This raises the question of whether there is a lottery proposition for the thesis that ordinary claims are generally true. I take this to be at least unclear. But fully answering this question would require a more complete discussion of what lottery propositions are (one that goes beyond merely citing examples). This is a topic for another occasion. See e.g. (Douven 2007: Sec. 2) for further discussion (that tends to be favorable to my cause).

lottery skeptic. And the present version of contextualism still purports to respond to the latter.

So, to sum up, the original factivity problem fails because even though the contextualist cannot properly state (LOW_C), she is not committed to this claim and thus can still properly state her view. The first revised factivity problem fails because it rests on a logical mistake. The second revised factivity problem fails because even though the contextualist is committed to (LOW_C^*), it couldn't be shown that she cannot properly state that claim. In particular, the premise ($HIGH_C^*$) that was supposed to establish this result turns out to be unfounded at least insofar as the contextualist only professes to respond to lottery but not to brain in a vat skepticism. We can thus conclude that there is no factivity problem for contextualism. Contextualism can be stated properly.

References

- Ashfield, M. (2013), 'Against the Minimalistic Reading of Epistemic Contextualism: A Reply to Wolfgang Freitag', *Acta Analytica*, 28/1: 111–125.
- Bach, K. (1987), *Thought and Reference* (Oxford: Oxford University Press).
- (1994), 'Conversational implicature', *Mind & Language*, 9/2: 124–162.
- (2001), 'Speaking Loosely: Sentence Nonliterality', *Midwest Studies in Philosophy*, 25/1: 249–263.
- (2002), 'Seemingly Semantic Intuitions', in J. K. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth. Investigations in Philosophical Semantics* (New York: Seven Bridges Press).
- (2005), 'The Emperor's New "Knows"', in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 51–90.
- (2006a), 'Implicature vs. explicature: What's the difference?'
<<http://userwww.sfsu.edu/kbach/Bach.ImplExpl.pdf>>, accessed 25 Jun 2015.
- (2006b), 'The Excluded Middle: Semantic Minimalism without Minimal Propositions', *Philosophy and Phenomenological Research*, 73/2: 435–442.
- (2008), 'Applying Pragmatics to Epistemology', *Philosophical Issues*, 18/1: 68–88.
- (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.
- Baker, C. (2012), 'Indexical Contextualism and the Challenges from Disagreement', *Philosophical Studies*, 157/1: 107–123.
- Baumann, P. (2008), 'Contextualism and the Factivity Problem', *Philosophy and Phenomenological Research*, 76/3: 580–602.
- (2010), 'Factivity and Contextualism', *Analysis*, 70/1: 82–89.
- (2011), 'WAMs: Why Worry?', *Philosophical Papers*, 40/2: 155–177.
- Black, T. (2005), 'Classic Invariantism, Relevance and Warranted Assertability Manœuvres', *The Philosophical Quarterly*, 55/219: 328–336.

- (2008a), ‘A Warranted-Assertability Defense of a Moorean Response to Skepticism’, *Acta Analytica*, 23/3: 187–205.
- (2008b), ‘Defending a Sensitive Neo-Moorean Invariantism’, in V. F. Hendricks and D. Pritchard (eds.), *New Waves in Epistemology* (Basingstoke: Palgrave Macmillan), 8–27.
- Blome-Tillmann, M. (2008), ‘The Indexicality of “Knowledge”’, *Philosophical Studies*, 138/1: 29–53.
- (2013a), ‘Conversational implicatures (and how to spot them)’, *Philosophy Compass*, 8/2: 170–185.
- (2013b), ‘Knowledge and implicatures’, *Synthese*, 190/18: 4293–4319.
- Brendel, E. (2005), ‘Why Contextualists Cannot Know they are Right: Self-Refuting Implications of Contextualism’, *Acta Analytica*, 20/2: 38–55.
- (2009), ‘Contextualism, Relativism, and Factivity. Analyzing “Knowledge” After the New Linguistic Turn in Epistemology’, in A. Hieke and H. Leitgeb (eds.), *Reduction, Abstraction, Analysis* (Publications of the Austrian Wittgenstein Society. New Series, 11, Frankfurt a. M.: Ontos Verlag), 403–416.
- Brogaard, B. (2008), ‘In Defense of a Perspectival Semantics of “Know”’, *Australasian Journal of Philosophy*, 86/3: 439–459.
- Brown, J. (2006), ‘Contextualism and Warranted Assertibility Manoeuvres’, *Philosophical Studies*, 130/3: 407–435.
- Brueckner, A. (2012), ‘Skepticism and Content Externalism’, in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (Stanford, CA: The Metaphysics Research Lab: Center for the Study of Language and Information: Stanford University).
- Brueckner, A., and Buford, C. T. (2009), ‘Contextualism, SSI and the Factivity Problem’, *Analysis*, 69/3: 431–438.
- (2010), ‘Reply to Baumann on Factivity and Contextualism’, *Analysis*, 70/3: 486–489.
- Buckwalter, W. (2014), ‘The Mystery of Stakes and Error in Ascriber Intuitions’, in J. Beebe (ed.), *Advances in Experimental Epistemology* (Advances in Experimental Philosophy, New York, NY: Bloomsbury Publishing), 145–173.
- Buckwalter, W., and Schaffer, J. (2013), ‘Knowledge, Stakes, and Mistakes’, *Noûs*: 1–34.

- Cappelen, H. (2008a), 'Content Relativism and Semantic Blindness', in M. García-Carpintero and M. Kölbel (eds.), *Relative Truth* (Oxford: Oxford University Press), 265–286.
- (2008b), 'The Creative Interpreter: Content Relativism and Assertion', *Philosophical Perspectives*, 22/1: 23–46.
- Cappelen, H., and Hawthorne, J. (2009), *Relativism and Monadic Truth* (Oxford: Oxford University Press).
- Cappelen, H., and Lepore, E. (2005), *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* (Malden, MA: Blackwell Publishing Ltd).
- Cohen, S. (2005a), 'Contextualism Defended', in M. Steup and E. Sosa (eds.), *Contemporary Debates in Epistemology* (Contemporary Debates in Philosophy, 3, Malden, MA: Blackwell Pub.), 56–62.
- (2005b), 'Knowledge, Speaker and Subject', *The Philosophical Quarterly*, 55/219: 199–212.
- 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.
- (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.
- (2007), 'Knowledge Claims and Context: Loose Use', *Philosophical Studies*, 132/3: 395–438.
- (2010), 'The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1, by Keith DeRose', *Mind*, 119/476: 1152–1157.
- DeRose, K. (1992), 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52/4: 913–929.
- (1995), 'Solving the Skeptical Problem', *The Philosophical Review*, 104/1: 1–52.
- (1999), 'Introduction: Responding to Skepticism', in K. DeRose and T. A. Warfield (eds.), *Skepticism. A Contemporary Reader* (Oxford: Oxford University Press).

- (2002), ‘Assertion, Knowledge, and Context’, *The Philosophical Review*, 111/2: 167–203.
- (2009), *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1* (Oxford: Oxford University Press).
- (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.
- Dodd, D. (2010), ‘Confusion about Concessive Knowledge Attributions’, *Synthese*, 172/3: 381–396.
- Dougherty, T., and Rysiew, P. (2009), ‘Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions’, *Philosophy and Phenomenological Research*, 78/1: 123–132.
- (2011), ‘Clarity about Concessive Knowledge Attributions: Reply to Dodd’, *Synthese*, 181/3: 395–403.
- Douven, I. (2007), ‘A Pragmatic Dissolution of Harman’s Paradox’, *Philosophy and Phenomenological Research*, 74/2: 326–345.
- Egan, A., Hawthorne, J., and Weatherson, B. (2005), ‘Epistemic Modals in Context’, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 131–168.
- Fantl, J., and McGrath, M. (2007), ‘On Pragmatic Encroachment in Epistemology’, *Philosophy and Phenomenological Research*, 75/3: 558–589.
- (2009), *Knowledge in an Uncertain World* (Oxford: Oxford University Press).
- (2012), ‘Contextualism and Subject-Sensitivity’, *Philosophy and Phenomenological Research*, 84/3: 693–702.
- Feldman, R. (2001), ‘Skeptical Problems, Contextualist Solutions’, *Philosophical Studies*, 103/1: 61–85.
- Freitag, W. (2011), ‘Epistemic Contextualism and the Knowability Problem’, *Acta Analytica*, 26/3: 273–284.
- (2013), ‘In Defence of a Minimal Conception of Epistemic Contextualism: A Reply to M. D. Ashfield’s Response’, *Acta Analytica*, 28/1: 127–137.

- Gerken, M. (2011), 'Epistemic Focal Bias', *Australasian Journal of Philosophy*, 91/1: 41–61.
- (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.
- Harman, G. (2007), 'Epistemic Contextualism as a Theory of Primary Speaker Meaning', *Philosophy and Phenomenological Research*, 75/1: 173–179.
- Hawthorne, J. (2004a), *Knowledge and Lotteries* (Oxford: Oxford University Press).
- (2004b), 'Replies', *Philosophical Issues*, 14/1: 510–523.
- Hazlett, A. (2009), 'Knowledge and Conversation', *Philosophy and Phenomenological Research*, 78/3: 591–620.
- Huang, Y. (2007), *Pragmatics* (Oxford Textbooks in Linguistics, Oxford: Oxford University Press).
- Hugly, P., and Sayward, C. (1979), 'A problem about conversational implicature', *Linguistics and Philosophy*, 3/1: 19–25.
- Ichikawa, J. (2011), 'Quantifiers and Epistemic Contextualism', *Philosophical Studies*, 155/3: 383–398.
- Jäger, C. (2012), 'Contextualism and the Knowledge Norm of Assertion', *Analysis*, 72/3: 491–498.
- Keysar, B. (1989), 'On the Functional Equivalence of Literal and Metaphorical Interpretations in Discourse', *Journal of Memory and Language*, 28/4: 375–385.
- Kölbel, M. (2003), 'Faultless Disagreement', *Proceedings of the Aristotelian Society*, 104: 53–73.
- (2004), 'Indexical Relativism versus Genuine Relativism', *International Journal of Philosophical Studies*, 12/3: 297–313.
- Kompa, N. (2005), 'The Semantics of Knowledge Attributions', *Acta Analytica*, 20/1: 16–28.

- (2011), ‘Nonindexical Contextualism—An Explication and Defense’, in S. Tolksdorf (ed.), *Conceptions of Knowledge* (Berlin Studies in Knowledge Research, 4, Berlin: De Gruyter), 329–349.
- Lackey, J. (1999), ‘Testimonial Knowledge and Transmission’, *The Philosophical Quarterly*, 49/197: 471–490.
- Larson, M., Doran, R., McNabb, Y. et al. (2009a), ‘Distinguishing the Said from the Implicated Using a Novel Experimental Paradigm’, in U. Sauerland and K. Yatsushiro (eds.), *Semantics and Pragmatics. From Experiment to Theory* (Palgrave Studies in Pragmatics, Language and Cognition, Basingstoke: Palgrave Macmillan), 74–93.
- Larson, M., Ward, G., McNabb, Y. et al. (2009b), ‘On the Non-Unified Nature of Scalar Implicature: An Empirical Investigation’, *International Review of Pragmatics*, 1/2: 211–248.
- Lasersohn, P. (1999), ‘Pragmatic Halos’, *Language*, 75/3: 522–551.
- Leite, A. (2005), ‘Some Worries for Would-Be WAMmers’, *Grazer Philosophische Studien*, 69/1: 101–126.
- Levinson, S. C. (1983), *Pragmatics* (Cambridge Textbooks in Linguistics, Cambridge: Cambridge University Press).
- Lewis, D. (1996), ‘Elusive Knowledge’, *Australasian Journal of Philosophy*, 74/4: 549–567.
- López de Sa, D. (2008), ‘Presuppositions of Commonality. An Indexical Relativist Account of Disagreement’, in M. García-Carpintero and M. Kölbel (eds.), *Relative Truth* (Oxford: Oxford University Press), 297–310.
- Ludlow, P. (2005), ‘Contextualism and the New Linguistic Turn in Epistemology’, in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy. Knowledge, Meaning, and Truth* (Oxford: Oxford University Press), 11–50.
- MacFarlane, J. (2005a), ‘Knowledge Laundering: Testimony and Sensitive Invariantism’, *Analysis*, 65/2: 132–138.
- (2005b), ‘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.

- (2011), ‘Relativism and Knowledge Attributions’, in S. Bernecker and D. Pritchard (eds.), *The Routledge Companion to Epistemology* (New York: Routledge), 536–544.
- (2014), *Assessment Sensitivity: Relative Truth and its Applications* (Oxford: Oxford University Press).
- May, J., Sinnott-Armstrong, W., Hull, J. et al. (2010), ‘Practical Interests, Relevant Alternatives, and Knowledge Attributions: An Empirical Study’, *Review of Philosophy and Psychology*, 1/2: 265–273.
- Meibauer, J. (2009), ‘Implicature’, in J. Mey (ed.), *Concise Encyclopedia of Pragmatics* (Amsterdam: Elsevier), 365–378.
- Montminy, M. (2007), ‘Epistemic Contextualism and the Semantics-Pragmatics Distinction’, *Synthese*, 155/1: 99–125.
- (2008), ‘Can Contextualists Maintain Neutrality?’, *Philosophers’ Imprint*, 8/7: 1–13.
- (2009), ‘Contextualism, Invariantism and Semantic Blindness’, *Australasian Journal of Philosophy*, 87/4: 639–657.
- (2013), ‘The Role of Context in Contextualism’, *Synthese*, 190/12: 2341–2366.
- Nagel, J. (2010), ‘Knowledge Ascriptions and the Psychological Consequences of Thinking about Error’, *The Philosophical Quarterly*, 60/239: 286–306.
- (2011), ‘The Psychological Basis of the Harman-Vogel Paradox’, *Philosophers’ Imprint*, 11/5: 1–28.
- Petersen, E. N. (2014), ‘Denying knowledge’, *Canadian Journal of Philosophy*: 1–20.
- Pinillos, N. Á. (2012), ‘Knowledge, Experiments and Practical Interests’, in J. Brown and M. Gerken (eds.), *Knowledge Ascriptions* (Oxford: Oxford University Press), 192–219.
- Pritchard, D. (2010), ‘Contextualism, Scepticism and Warranted Assertibility Manoeuvres’, in J. K. Campbell, M. O’Rourke, and H. Silverstein (eds.), *Knowledge and Skepticism* (Topics in Contemporary Philosophy, Cambridge, MA: MIT Press), 85–103.
- Richard, M. (2004), ‘Contextualism and Relativism’, *Philosophical Studies*, 119/1/2: 215–242.
- (2008), *When Truth Gives Out* (Oxford: Oxford University Press).

- Rysiew, P. (2001), 'The Context-Sensitivity of Knowledge Attributions', *Noûs*, 35/4: 477–514.
- (2005), 'Contesting Contextualism', *Grazer Philosophische Studien*, 69/1: 51–70.
- (2007), 'Speaking of Knowing', *Noûs*, 41/4: 627–662.
- (2009), 'Epistemic Contextualism', in E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy (Spring 2009 Edition)* (Stanford, CA: The Metaphysics Research Lab: Center for the Study of Language and Information: Stanford University).
- Saul, J. M. (2010), 'Speaker-Meaning, Conversational Implicature and Calculability', in K. Petrus (ed.), *Meaning and Analysis. New Essays on Grice* (Basingstoke: Palgrave Macmillan), 170–183.
- Schaffer, J. (2004a), 'From Contextualism to Contrastivism', *Philosophical Studies*, 119/1/2: 73–103.
- (2004b), 'Skepticism, Contextualism, and Discrimination', *Philosophy and Phenomenological Research*, 69/1: 138–155.
- Schaffer, J., and Knobe, J. (2012), 'Contrastive Knowledge Surveyed', *Noûs*, 46/4: 675–708.
- Schiffer, S. (1996), 'Contextualist Solutions to Scepticism', *Proceedings of the Aristotelian Society*, 96: 317–333.
- Shapiro, S. (2006), *Vagueness in Context* (Oxford: Oxford University Press).
- Sosa, E. (1999), 'How to Defeat Opposition to Moore', *Philosophical Perspectives*, 33/13: 141–153.
- Sripada, C. S., and Stanley, J. (2012), 'Empirical Tests of Interest-Relative Invariantism', *Episteme: A Journal of Social Epistemology*: 3–26.
- Stanley, J. (2005a), 'Fallibilism and Concessive Knowledge Attributions', *Analysis*, 65/2: 126–131.
- (2005b), *Knowledge and Practical Interests* (Short Philosophical Books, Oxford: Oxford University Press).
- (2007), *Language in Context: Selected Essays* (Oxford: Oxford University Press).
- Turri, J. (2014), 'Linguistic Intuitions in Context: A Defense of Nonskeptical Pure Invariantism', in A. Booth and D. Rowbottom (eds.), *Intuitions* (Oxford: Oxford University Press), 165–185.

- Vogel, J. (1990), 'Are there Counterexamples to the Closure Principle?', in M. D. Roth and G. Ross (eds.), *Doubting. Contemporary Perspectives on Skepticism* (Dordrecht: Kluwer Academic Publishers), 13–27.
- Williamson, T. (2001), 'Comments on Michael Williams' Contextualism, Externalism and Epistemic Standards', *Philosophical Studies*, 103/1: 25–33.
- (2005a), 'Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge', *The Philosophical Quarterly*, 55/219: 213–235.
- (2005b), 'Knowledge and Skepticism', in F. Jackson and M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy* (Oxford: Oxford University Press), 681–700.
- Wilson, D. (2006), 'The pragmatics of verbal irony: Echo or pretence?', *Language in Mind: A Tribute to Neil Smith on the Occasion of his Retirement*, 116/10: 1722–1743.
- Wright, C. (2005), 'Contextualism and Scepticism: Even-Handedness, Factivity and Surreptitiously Raising Standards', *The Philosophical Quarterly*, 55/219: 236–262.