

The Evidence for Contextualism about Knowledge Ascriptions

Torfinn Thomesen Huvenes



Thesis Submitted for the Master of Arts Degree,
Department of Philosophy, Classics, History of Arts and Ideas

UNIVERSITETET I OSLO

November 2006

Acknowledgements:

First and foremost I want to thank my supervisor Professor Herman Cappelen. He has provided very valuable feedback and support for which I am very grateful. I also want to thank Professor John Hawthorne for useful suggestions and comments, as well as a very rewarding conversation. Furthermore, I want to thank everyone else who has supported me in my work on this thesis. I am very grateful.

Table of Contents

1	Introduction	1
2	Contextualism about Knowledge Ascriptions.....	5
2.1	The Context Sensitivity of ‘Know’	5
2.2	Fleshing Out Contextualism.....	7
2.3	Alternatives to Contextualism 1: Classical and Subject-Sensitive Invariantism	9
2.4	Alternatives to Contextualism 2: Relativism	11
3	The Cases	15
3.1	Case 1: The Bank Case.....	15
3.2	Case 2: The Airport Case	17
4	Expanding the Range of the Data.....	19
4.1	More Data about Our Use of ‘Know’	19
4.2	Ignorant High Stakes	21
4.3	Are the Cases Psychologically Realistic?	27
5	Contextualism and Context Shifting Arguments	31
6	An Alternative to a Semantic Treatment of the Cases	37
6.1	Contextualism and Warranted Assertability	38
6.2	Criteria for a Plausible Candidate for a WAM.....	41
6.3	A More Plausible Candidate for a WAM in Defence of Invariantism.....	45
6.4	Cancellability	49
6.5	A WAM for ‘Tall’	51
6.6	The Sensitivity of Competent Language Users with Respect to the Semantic Workings of Their Own Language.....	53
6.7	Appendix: A Positive Argument for Contextualism Based on the Knowledge Account of Assertion.....	55
7	Contextualism and Semantic Blindness	63
7.1	Propositional Attitude Reports and ‘Know’	63
7.2	Comparative Judgements of Content and Metalinguistic Claims	65
7.3	Semantic Blindness and the Positive Evidence for Contextualism.....	68
8	Summary and Conclusions.....	73
9	Bibliography.....	77

1 Introduction

There are many expressions in English the content of which is sensitive to, or depend on, the context in which they are uttered. Different utterances of the expression ‘here’ may refer to different locations depending on the location of the speaker. Similarly, the expression ‘that’ may refer to different objects in different contexts. In one context, ‘that’ may refer to a book, in another context it may refer to a tree. The expressions ‘here’ and ‘that’ are prime examples of what we may call ‘context sensitive expressions’. Context sensitivity is a phenomenon that has been studied by philosophers of language who are interested in the workings of natural languages such as English. However, the phenomenon of context sensitivity has also attracted the interest of philosophers of other disciplines such as epistemology and metaphysics. By postulating that a certain expression, or perhaps even a class of expressions, is context sensitive, these theorists hope to provide solutions to epistemological or metaphysical problems.

Contextualism about knowledge ascriptions is a view that has received a lot of attention and has been at the centre of much debate in recent years. Simply put, contextualism about knowledge ascriptions (henceforth just ‘contextualism’) is the view that the English expression ‘know’ as it occurs in ascriptions of propositional knowledge is context sensitive. One can ascribe knowledge to a subject in one context, while denying knowledge to the same or a similarly positioned subject in another context, but still be speaking truly in both cases. Prominent defenders of contextualism in the past decade include Stewart Cohen, Keith DeRose and David Lewis. The contributions of Stewart Cohen and Keith DeRose will be at the centre of the discussion of contextualism which I will undertake here. In particular, I will put a lot of emphasis on the contributions of Keith DeRose as he has addressed many of the questions that I will be concerned with. Additionally, I will draw on the work of non-contextualists such as, among others, John Hawthorne, Jason Stanley and Timothy Williamson. In the case of both contextualists and non-contextualists, I will focus on the contributions to the debate in its most recent incarnation to the exclusion of the historical antecedents of the positions and arguments under discussion.

As a semantic thesis about the expression ‘know’ in English, it seems appropriate that a discussion of contextualism should belong to the philosophy of language, and more specifically to semantics. However, contextualism promises to have significant implications

for epistemology, especially if one takes knowledge to be as central to epistemology as Timothy Williamson (2000) has argued that it is. Contextualist solutions have been offered for important epistemological problems such as the challenge posed by sceptical arguments.¹ According to the picture offered by the contextualists with respect to providing a solution to the problem of scepticism, an utterance of the sentence ‘I know that I have hands’ will be true in most everyday contexts where we are not considering any radical sceptical possibilities, assuming of course that the subject has hands.² However, when discussing epistemology and considering the sceptical scenario that we are handless brains in vats, an utterance of the sentence ‘I know that I have hands’ might be false. Indeed, the prospect of providing a response to the problem of epistemological scepticism has probably been one of the main motivating factors behind contextualism. Keith DeRose emphasizes this when he says, ‘Contextualist theories of knowledge attributions have almost invariably been developed with an eye toward providing some kind of answer to philosophical skepticism’ (1995, p. 4).

However, contextualism as I have presented it is a semantic thesis and it appears problematic to accept a semantic thesis solely on the basis of its application to epistemological problems such as the problem of scepticism. Herman Cappelen and Ernie Lepore make the following point:

There can’t be metaphysical, epistemological, or moral arguments for an expression being context sensitive. If an English expression *e* is context sensitive, then that’s a fact about English and it can only be established by whatever procedures semanticists employ to establish such facts. (Cappelen and Lepore 2003, p. 43)

As a result of this, the contextualists have felt the need to find evidence for contextualism. In attempting to do this, contextualists have looked to our ordinary use of sentences containing the expression ‘know’. Typically this has involved presenting a pair of cases with the subjects in both cases being in a similar epistemic position with respect to a true proposition. The contextualists argue that while it seems correct to ascribe knowledge to the subject in one of the cases, it also seems correct to deny knowledge to the subject in the other case. However, the cases only differ with respect to features which are not generally considered to be

¹ David Lewis (1996) argues that contextualism is also in a position to provide a solution to other central epistemological problems such as the lottery problem and the Gettier problem. Stewart Cohen (1998) discusses Lewis’ proposal, but while he agrees that contextualism provides a solution to problem of scepticism and the lottery, he finds a contextualist solution more problematic in the case of the Gettier problem.

² The precise mechanics behind this change from the everyday context to the epistemological context can vary between different ways of working out the contextualist position. I will not get into those details since I will not be discussing the application of contextualism to the problem of scepticism.

epistemically relevant. From this the contextualists conclude that the expression 'know' is context sensitive.

It is cases such as these, and the evidence that they allegedly provide for contextualism, that I will discuss in this thesis. In doing so, I will focus on two cases presented by Stewart Cohen (1999) and Keith DeRose (1992) respectively. I do not intend to undertake a comprehensive discussion of the merits and flaws of contextualism. For the most part I will ignore the objections that have been raised directly against contextualism insofar as they do not touch on the evidence for contextualism. Rather, I want to focus on the cases that contextualists have advanced as evidence for their position. Nor will I undertake a further discussion of the epistemological implications of contextualism, including the potential application to the problem of scepticism that I briefly mentioned above.

It should be noted that I do not wish to rule out the possibility that there are other ways of providing evidence for contextualism than the cases I will be discussing. For instance, if one takes the position that all vague expressions are context sensitive, one could find support for the view that 'know' is context sensitive by arguing that it is vague. However, even leaving the merits of such an argument aside, it is an open question whether this is what the contextualists have in mind when they argue that 'know' is context sensitive. Among other things, it is unclear how this version of contextualism could be used to provide a response to the sceptical arguments.³ I propose to leave such considerations aside and focus on the cases presented by Cohen and DeRose.

My overall aim is to present contextualism along with some of the cases which are thought to motivate this position, and then undertake a critical discussion of these cases with respect to the question of whether they provide evidence for contextualism. This overall aim will be reflected in the structure of my thesis. I will begin by giving a more detailed and precise formulation of contextualism (section 2). This will involve providing a more precise formulation of what it means for the expression 'know' to be context sensitive. Moreover, I will mention some of the other positions in the debate and how they differ from contextualism in order to cast further light on contextualism and to place it in the theoretical landscape. Thus having stated the contextualist position, I give a more thorough presentation of a pair of typical cases that contextualists claim provide evidence for their position (section 3). This paves the way for the main part of the thesis which will be a discussion of some problems with taking the cases in question to provide evidence for contextualism. I begin this part of the

³ This seems to be quite close to Stephen Schiffer's (1996, p. 327-328) thoughts on these matters.

thesis by raising some initial complications (Section 4). In this section I look at some data that contextualists have tended not to focus on, and discuss the worry that the cases appealed to by the contextualists are not psychologically realistic. Then I turn to a discussion of the type of argument that contextualists rely on in order to draw the conclusion that ‘know’ is context sensitive from the cases in question (section 5). The main part of the discussion will centre on the question of whether the cases really require a treatment in terms of semantic context sensitivity, or whether there is an alternative explanation in terms of warranted assertability (section 6). Finally, I will connect the discussion of the evidence for contextualism with the topic of semantic blindness (section 7). At the end I make a summary of the discussion and offer some concluding remarks (section 8).

2 Contextualism about Knowledge Ascriptions

In order to undertake a proper discussion of the evidence for contextualism I will begin by providing a more precise formulation of the contextualist position. On the understanding of contextualism that I will be concerned with here, contextualism is the view that the expression ‘know’ (including its different inflectional forms, such as third person singular and past tense) is context sensitive. My way of putting the issue differs slightly from for instance Jason Stanley (2005) and Keith DeRose (1992), who prefer to say that contextualism is the view that knowledge ascriptions (i.e. knowledge-ascribing and knowledge-denying sentences) are context sensitive. However, I take this point to be primarily a matter of terminology.

In section 2.1 I provide a more precise and detailed formulation of what it means for ‘know’ to be a context sensitive expression, and make a few clarifying remarks regarding what the contextualist thesis amounts to. In section 2.2 I look at different ways of further developing contextualism. The last two sections deal with alternatives to contextualism. I present invariantism as an alternative to contextualism in section 2.3 and relativism in section 2.4. The aim in these two sections is to further clarify the contextualist position by looking at how it differs from alternative positions.

2.1 The Context Sensitivity of ‘Know’

Since contextualism is the thesis that ‘know’ is a context sensitive expression, it is useful to have a more precise formulation of what this amounts to. Following the general schema for stating what it means for an expression to be context sensitive used by Herman Cappelen and Ernie Lepore (2005b, p. 146), we may formulate the context sensitivity of ‘know’ in two different ways:

- (CS') To say that ‘know’ is context sensitive is to say that its contribution to the truth conditions of utterances u of a sentence S containing ‘know’ references various aspects of the context of u .
- (CS'') To say that ‘know’ is context sensitive is to say that its contribution to the propositions expressed by utterances of sentences containing ‘know’ varies from context to context.

This means that different utterances of a sentence containing ‘know’, can express different propositions, or alternatively, have different truth conditions, in different contexts.⁴ In the following I will mostly put the context sensitivity of ‘know’ in terms of a variation in truth conditions, following (CS'). However, I do not intend the use of one rather the other to convey any substantial theoretical commitments, but will use the definitions as convenient, and so I do not see any reason to make an explicit choice between (CS') and (CS").⁵ This reflects my intention to remain as neutral as possible with respect to questions about the nature of semantic content, such as the question of whether the semantic content of an utterance of a sentence is exhausted by its truth conditions, or whether the semantic content is a proposition which incorporates the truth conditions of the utterance.

We can further illustrate what context sensitivity is by means of an example. Certain expressions are obviously context sensitive. Take for instance the indexical ‘I’ as it occurs in sentence (1).

- (1) I am a philosophy student.

The contribution of the expression ‘I’ in sentence (1) to the proposition expressed or the truth conditions of an utterance of the sentence, depends on who is the speaker in the context in which (1) is uttered. Since different utterances of (1) may thus express different propositions and have different truth conditions, it may be true as uttered by one speaker, for instance myself, and false as uttered in a context with a different speaker (i.e. someone who is not a philosophy student).⁶ The contextualists argue that similar considerations apply in the case of ‘know’.⁷

- (2) Al Gore knows that George W. Bush is the President of the United States.

⁴ I will talk about utterances rather than sentences of expression-in-context. While this may seem to suggest a focus on spoken uses of language, this is not my intention. The discussion is meant to cover written inscriptions as well.

⁵ It will for instance be convenient to appeal to the difference between truth conditions and warranted assertability conditions.

⁶ For ease of exposition, I will not be careful about whether I attribute truth and falsity to utterances of sentences, the sentences themselves, ascriptions and denials of knowledge, assertions (i.e. the speech act) or the proposition expressed by an utterance of a sentence. My intended use does therefore not reflect any view about what the ultimate bearers of truth and falsity are. For instance, if one finds it problematic to attribute truth to utterances one might view the truth or falsity of an utterance as derived from the truth or falsity of the proposition it expresses.

⁷ One should not conclude from this that contextualism requires ‘know’ to be an indexical like ‘I’. There may be other ways of implementing the context sensitivity of ‘know’ than indexicality. See section 2.2.

According to contextualism the expression ‘know’ in sentence (2) is similar to the expression ‘I’ in that its contribution to the proposition expressed by an utterance of the sentence depends on certain features of the context, for the moment leaving it an open question exactly which features of the context are important and how the context sensitivity is implemented. In an ordinary context the truth conditions of an utterance of (2) may be such that it comes out true. On the other hand, an utterance of (2) may be false in a more demanding context, perhaps one in which sceptical possibilities are salient or a context in which the possibilities of a coup d’état has been raised.

Note that the discussion here only concerns ‘know’ as it occurs in ascriptions of propositional knowledge, that is sentences of the form ‘S knows that *p*’, in which ‘know’ denotes a relation between an individual and a proposition.⁸ I will not discuss ‘know’ as it occurs in sentences like (3), in which ‘know’ denotes a relation between individuals rather than a relation between an individual, or individuals, and a proposition.

(3) Tony Blair knows George W. Bush.

It is worth bringing up a small point of clarification regarding tense and context sensitivity. It is clear that different utterances of (2) may differ in truth value if uttered at different times. For instance, (2) would be false as uttered prior to Bush becoming President. However, this should not be taken as vindicating the contextualist claim that ‘know’ is context sensitive, and throughout the thesis I will ignore any context sensitivity due to the tense of the verb ‘know’, or indeed any other context sensitivity due to tense.⁹ We may therefore follow Stanley in taking contextualism to require ‘know’ to be ‘context-sensitive in a distinctly epistemological way’ (2005, p. 16).

2.2 Fleshing Out Contextualism

Having thus stated the contextualist position, it should be noted that there are different ways of fleshing it out. While I do not want to focus on different ways of developing and implementing contextualism in this thesis, I still think it is a good idea to give a brief outline

⁸ At least I will assume throughout the discussion that ‘know’ denotes a relation between an individual (or individuals) and a proposition.

⁹ To be more precise, we could treat an explicit knowledge ascription as having the form ‘S knows at *t* in *w* that *p*’, where *S* is the subject, *p* is a proposition, *t* is a time and *w* is a possible world (Williamson 2005, p. 214). In the following I will ignore the time-parameter and world-parameters and assume that these are fixed.

of some different directions that contextualists may take in developing their position, as fleshing out the contextualist position will make the following discussion easier to understand.

David Lewis (1996) develops his version of contextualism by appealing to a shift in the possibilities that the evidence possessed by a subject *S* must rule out in order for it to be true to say ‘*S* knows that *p*’. According to his version of contextualism ‘*S* knows that *p*’ is true if and only if *S*’s evidence eliminates every possibility in which *p* does not hold (Lewis 1996, p. 551). However, the domain of the quantifier is restricted, meaning that some possibilities are left outside the domain and can be ignored. Which possibilities the subject’s evidence must rule out, and which can be ignored, depends on the context. In certain contexts it might be necessary to rule out the possibility of being a brain in a vat, while in most everyday contexts this is not required because that possibility is being properly ignored. This is also the framework originally suggested by Stewart Cohen (1988). However, according to a later version of Cohen’s position he defends a view according to which justification is required for knowledge and what shifts is ‘how justified a belief must be in order to be justified *simpliciter*’ (1999, p. 60). DeRose (1995) on the other hand prefers to talk about a shift in the epistemic standards associated with ‘know’. Sometimes the contextually determined standards for knowledge are demanding, requiring the subject to be in a strong epistemic position. In other contexts, a weaker epistemic position may be sufficient.

Contextualist accounts may also differ with respect to which features of the context are viewed as relevant to the variation in the truth conditions of utterances of sentences containing ‘know’. Different versions of contextualism may place emphasis on different features such as the practical stakes involved or the attentions of the speaker. Lewis’ (1979; 1996) account seems to focus on the latter while DeRose (1992, p. 916) prefers to stress the former. This does not mean that a contextualist is forced to choose between one or the other, only that this is an area in which contextualist accounts may differ, and it is worth emphasizing that though their perspectives are slightly different, both Lewis and DeRose appeal to both set of features.

Moreover, there is the question of how the context sensitivity of ‘know’ is to be implemented. A very straightforward suggestion is to treat ‘know’ as an indexical like ‘I’ or ‘here’. According to this model, ‘know’ would denote different relations in different contexts. I will follow Stanley (2005, p. 48) in understanding Cohen, DeRose and Lewis as sharing the view that the alleged context sensitivity of ‘know’ is to be implemented by treating ‘know’ as an indexical. However, this is not only way of implementing the context sensitivity of ‘know’. Another way of implementing the context sensitivity of ‘know’ is to argue that there is a

hidden indexical associated with 'know' in the logical form of sentences containing 'know'. For instance, Peter Ludlow (2005) has presented a version of contextualism according to which there is an implicit L-marked position for standards of knowledge associated with 'know'. He takes his view to be motivated by the fact that 'know' can take modifiers such as 'by scientific standards' and 'with some reliability' as in (4).

- (4) Stephen Hawking knows that there are black holes by scientific standards.

Whether this is adequate evidence for postulating that there is an implicit L-marked position for standards of knowledge associated with 'know' is not a question that I will pursue at any great lengths. However, it is worth noting, as Stanley (2005, p. 70) points out, that modifiers such as 'by scientific standards' are not only used with respect to sentences containing 'know' or other epistemic terms, and that this suggests that we are dealing with a more general phenomenon.

It should be noted that according to the version of contextualism proposed by Lewis (1996), 'know' involves restricted quantification over possibilities. His version of contextualism therefore allows the context sensitivity of 'know' to be understood in terms of the context sensitivity of quantifiers with contextually restricted domains. This means that the details of how the context sensitivity of 'know' is ultimately implemented, depends on one's views on how the context sensitivity of quantifiers is to be implemented.

For the purpose of this thesis I propose mostly to gloss over the differences between different approaches to contextualism. In the following I will focus on contextualism as the view that 'know' is context sensitive as defined by (CS') and (CS'') rather than the different ways of developing and implementing that view. Moreover, I assume that the discussion which follows is relevant for different versions of contextualism. However, there are points in the following discussion where what one has to say about the more detailed workings of the alleged context sensitivity of 'know' becomes more important, and I will do my best to point them out.

2.3 Alternatives to Contextualism 1: Classical and Subject-Sensitive Invariantism

So far, I have presented contextualism as the view that the expression 'know' is context sensitive. My present concern is to briefly sketch some of the alternative positions in the

debate about contextualism. This is in order to place contextualism in the theoretical landscape, and to help clarifying further what the contextualist position amounts to.

One alternative to contextualism is the view which is often referred to as ‘invariantism’. The invariantists deny that ‘know’ is context sensitive. According to the invariantist account, the contribution of the expression ‘know’ to the proposition expressed by an utterance of a sentence which contains it does not depend on the context. Consequently, assuming that there are no other context sensitive features of the sentence in question, different utterances of a sentence containing ‘know’ express the same proposition. It is normal to make a distinction between sceptical invariantism on the one hand, and moderate or non-sceptical invariantism on the other. According to the sceptical invariantists, the truth conditions of knowledge-ascriptions are such that few, if any, of our ordinary knowledge ascriptions are true. The non-sceptical invariantists on the other hand, maintain that many of our knowledge ascriptions are true. This is not an attempt to provide a precise definition of sceptical and non-sceptical invariantism, but it is all that is needed for the present purpose. Most of the invariantist accounts that will be discussed in this thesis are versions of non-sceptical invariantism, but I will not put a lot of emphasis on this distinction.

Furthermore, there is a position which has emerged and attracted a lot of attention in recent years. This position is often called ‘subject-sensitive invariantism’ (just ‘sensitive invariantism’ in the following) and is distinguished from classical or insensitive invariantism. A prominent defender of this view is Jason Stanley (2005). John Hawthorne (2004, ch. 4) has also presented a version of this view, but while his treatment is largely positive, he does not explicitly endorse it.¹⁰ Everyone in the debate can agree that whether a subject knows a given proposition depends on certain features of the situation of the subject. Examples of relevant features of the subject’s situation may include, among other things, whether the subject believes the proposition in question and whether she possesses good evidence for it.¹¹ The sensitive invariantist offers a special account in that they permit features of the subject’s situation which are not traditionally seen as epistemically relevant to play a role in determining whether the subject knows the proposition in question. Typically, this includes, but need not be restricted to, the practical concerns and interests of the subject (i.e. what is at stake for the subject). Timothy Williamson states the matter accordingly, ‘For them [the

¹⁰ It should be noted that neither Hawthorne nor Stanley use the term ‘subject-sensitive invariantism’. Hawthorne (2004) uses the term ‘sensitive moderate invariantism’ and Stanley (2005) refers to his view as ‘interest-relative invariantism’. The name, ‘subject-sensitive invariantism’ is nevertheless quite common in the literature on the subject, and is used among others by DeRose (2004b; 2005).

¹¹ These examples are not meant to be uncontroversial. Lewis (1996, p. 556) for instance, wants to allow for cases of knowledge without belief.

sensitive invariantists], what epistemic standard the subject must meet in order to know is sensitive to non-epistemic features of the subject's circumstances' (2005, p. 217).¹²

Thus the features of the situation of the subject which the sensitive invariantists emphasize are similar to the features of the context which affect a shift in truth conditions according to the contextualist view. However, it is important to stress the difference between these positions. It is the context of utterance, the context in which the utterance is made, which is important according to contextualism. The sensitive invariantist account on the other hand, concerns the situation of the subject to whom knowledge is ascribed or denied. Consequently, they make somewhat different predictions about cases in which the context which the subject occupies does not coincide with the context of utterance. That is cases in which the speaker and the subject are not the same person.

Stanley further stresses the extent to which the two positions differ in that contextualism is a semantic thesis while sensitive invariantism is not. 'In contrast with contextualism, the advocate of IRI [interest-relative invariantism] has no semantic burden to discharge. IRI is not a semantic thesis at all; it is rather a metaphysical thesis about the nature of the knowledge relation' (Stanley 2005, p. 120). It may be argued that Stanley overplays this distinction. Sensitive invariantism is not neutral on the truth conditions of utterances of sentences containing 'know' because it allows the practical situation of the subject to play a role in determining whether an ascription of knowledge is true. Still, sensitive invariantism differs from contextualism in not positing the semantic phenomenon of context sensitivity. All the same, contextualism and sensitive invariantism are motivated by a similar sort of cases. As a result, sensitive invariantism makes the task of providing evidence for contextualism somewhat more complicated.

2.4 Alternatives to Contextualism 2: Relativism

Another recently proposed alternative to contextualism is relativism about truth. This is a position which has been developed and defended by John MacFarlane (2005a). On MacFarlane's version of relativism the truth of a sentence or proposition depends on the context of assessment and not just the context of utterance ('context of use' in MacFarlane's terminology). The context of assessment is the context in which a speech act is being assessed or evaluated as opposed to the context of utterance or use in which the sentence is uttered, and

¹² Williamson's statement of the sensitive invariantist position turns on the non-epistemic features of the subject's situation being understood as those features which are not truth-conducive, or at least those features which are traditionally taken to be non-epistemic features. See Stanley (2005, p. 2).

the relativists argue that the truth value of certain sentences or propositions can vary with the context of assessment. A sentence or proposition, the truth value of which varies with the context of assessment, is assessment sensitive. Relativism, in its weakest form, is the view that there is at least one assessment sensitive sentence, either in a natural language or in some conceivable language (2005a, p. 328).

MacFarlane has not put forth relativism as a global view along the lines that all sentences are assessment sensitive, but has rather argued for a relativist treatment of certain classes of sentences. Most interestingly for the purpose of the present thesis, MacFarlane (2005b) has argued that a relativist semantics might be plausible in the case of 'know'. This would amount to the expression 'know' contributing to the assessment sensitivity of sentences that contains it. Like contextualism, relativism about knowledge ascriptions is a semantic thesis about the expression 'know' and sentences that contains it. But while both contextualists and relativists about knowledge ascriptions hold that the truth value of sentences containing the expression 'know' depends on context, relativists about knowledge ascriptions differ from contextualists in taking the context of assessment rather than the context of utterance to be important.

In a sense one could say that relativism is a somewhat more radical semantic thesis than contextualism. While there are uncontroversial examples of context sensitivity in the sense suggested by the contextualists, as was indeed illustrated by the expression 'I' in (1), there does not appear to be similarly uncontroversial examples of assessment sensitivity.¹³

- (1) I am a philosophy student.

While I will not discuss the relative merits of relativism versus contextualism in this thesis it is important to note that, like sensitive invariantism, relativism draws support from the same cases that allegedly provide evidence for contextualism, together with dissatisfaction with the contextualist position as well as the classical and sensitive invariantist positions (2005b). However, since relativism appears to be a more radical semantic thesis than contextualism, it is unclear whether the possibility of a relativist account of the cases that will be the focus of this thesis in itself undermines the evidence for contextualism. Still, the discussion of these cases is not irrelevant for the prospects of a relativist treatment of 'know'. On the one hand

¹³ A relativist treatment has also been suggested in the case of epistemic modals (Egan, Hawthorne and Weatherson 2005) and future contingents (MacFarlane 2003) to name two examples. However, a relativist treatment is far from uncontroversial in either case.

there is the possibility that a relativist position might be an alternative if the contextualist account is found to be inadequate. On the other hand it means that relativism relies on much of the same evidence as contextualism. Thus, if this evidence is undermined it would presumably undermine the case for a relativist treatment of knowledge ascriptions as well.

3 The Cases

In order to provide evidence for the thesis that ‘know’ is a context sensitive expression, contextualists appeal to intuitions about different utterances of sentences containing ‘know’. One might consider our intuitions concerning sceptical arguments as a possible source of evidence for contextualism. In contexts where we are considering sceptical scenarios we are apt to find utterances of knowledge-ascribing sentences to be false that we would judge to be true in most ordinary contexts. The contextualists can account for this in terms of different utterances sentences containing ‘know’ having different truth conditions. However, there are certain limitations involved with this strategy for providing evidence for contextualism. First, there is the highly controversial nature of cases involving sceptical scenarios. Second, as DeRose (2005, p. 178) notes, our intuitions about these cases are not very strong or stable. Many would probably resist the intuition that we do not know that we have hands in contexts where sceptical possibilities are salient. Therefore, contextualists have presented other cases that they argue provide more reliable evidence that ‘know’ is context sensitive. My goal here is to present two such cases. In section 3.1 I present DeRose’s Bank Case and in section 3.2 I present Cohen’s Airport Case.

3.1 Case 1: The Bank Case

Since the nature of cases involving sceptical possibilities remain controversial, other cases which do not involve such radical sceptical possibilities, but focus on more mundane concerns, have been presented by contextualists seeking to provide support for the view that ‘know’ is context sensitive. I will begin with a pair of cases originally presented by DeRose:

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

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

a *very* bad situation. And, of course, the bank is not open Sunday. My wife reminds me of the facts. She then says, “Banks do change their hours. Do you know the bank will be open tomorrow?” Remaining as confident as I was before that the bank will be open then, still, I reply, “Well, no. I’d better go in and make sure.” (1992, p. 913)

I will refer to this pair of cases collectively as ‘the Bank Case’. In Bank Case A, where relatively little is at stake and serious possibilities of error have not been raised, DeRose’s ascription of knowledge to himself appears to be true, assuming as DeRose (1992, p. 914) stipulates in his presentation, that the bank is in fact open Saturday. However, in Bank Case B, where the stakes are higher and DeRose’s wife has raised the possibility of the bank having changed its hours, the contextualists argue that we have equally strong intuitions that DeRose’s denial of knowledge is true. However, DeRose’s epistemic position with respect to the proposition that the bank will open is the same in the two cases. In both cases it is based on the evidence that he was at bank two weeks ago on a Saturday. Thus DeRose’s utterance of (5a) in Case A appears to be true, but so does the utterance of (5b) in Bank Case B.

(5a) I know that it’ll be open.

(5b) Well, no, I don’t know that it’ll be open.¹⁴

Contextualism provides a way of removing the apparent contradiction by appealing to the alleged context sensitivity of ‘know’. If different utterances of (5a) and (5b) have different truth conditions due to the context sensitivity of the expression ‘know’ it is easy to see how both could be true.¹⁵ In the context described in Bank Case A the contribution of ‘know’ to the truth conditions of the utterance of (5a) is such that it comes out true. On the other hand, if DeRose were to utter (5a) in the context described in Bank Case B the contextualists would argue his utterance would be false. The contextualists can point to the higher stakes involved and the possibilities of error raised in Bank Case B as affecting a shift in truth conditions, making (5b) come out true and (5a) come out false. As a result of this contextualism allows us to retain our intuitions that the utterances made by DeRose in both Bank Case A and B are true.¹⁶

¹⁴ I have taken the liberty of slightly modifying the sentence from DeRose’s presentation of the case in order to make it more explicit.

¹⁵ It is of course possible to have utterances of (5a) and (5b) that are both true because they take place in contexts with different speakers or at different times. However, this does not affect the present cases where the time and the speaker are fixed, and I therefore ignore this complication both in this case, and in the following cases.

¹⁶ I do not mean to suggest that contextualism as the semantic thesis that ‘know’ is a context sensitive expression is committed to these being the correct truth values for the utterances of (5a) and (5b). The point is rather that contextualism is in a position to account for our intuitions concerning the truth values of DeRose’s utterances in

Though this version of the Bank Case seems to pose a problem for classical invariantists who are seemingly forced to reject our intuitions about the truth values of the utterances in either Bank Case A or B, it is worth noting that an important feature of this version of the Bank Case is that it is a first-person case. That is to say that in this case the subject and the speaker are the same person. This is important in evaluating the effectiveness of this case as a means of providing evidence for contextualism because such first-person cases can be accommodated by sensitive invariantism as well as contextualism. The sensitive invariantists can argue that the higher stakes involved in Bank Case B prevent DeRose from knowing that the bank will be open Saturday while maintaining that the ascription of knowledge to him in Bank Case A is true. The upshot of this is that sensitive invariantism can deliver the same verdict as contextualism with respect to these cases, namely that DeRose's utterance of (5a) in Bank Case A, and his utterance of (5b) in Bank Case B, are both true.

3.2 Case 2: The Airport Case

In response to sensitive invariantism being able to accommodate first-person cases such as the Bank Case, contextualists have attempted to come up with cases which are less easily accommodated by sensitive invariantism. Rather than giving a different version of the Bank Case, I turn to Stewart Cohen's Airport Case.¹⁷ My main reason for doing this is that I want to avoid any unnecessary confusion as a result of mixing the first-person and third-person cases.

Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes I know – it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really *know* that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen 1999, p. 58)

The Airport Case is very similar to the Bank Case. In both cases the alleged shift in truth conditions is tied to the practical stakes involved and the possibilities of error that are salient. When Smith makes an utterance of (6a) it comes out true, assuming again that the plane stops

Bank Case A and B, and I will confine my discussion to a version of contextualism which seeks to accommodate these intuitions.

¹⁷ For a third-person version of the Bank Case presented with the purpose of providing evidence against sensitive invariantism see DeRose (2005, p. 184) as well as Stanley (2005, p. 3-6).

in Chicago, since the practical stakes are quite low in his context and there are no serious and salient error possibilities.

(6a) Yes I know – it does stop in Chicago.

Though Mary and John has the same evidence as Smith in the itinerary, more is at stake for them and they are considering the possibility that the itinerary contains a misprint or that the schedule has been changed. Consequently, the contextualists can accommodate the intuition that if either Mary or John were utter (6b) their utterance would seem to be true, by appealing to a shift in truth conditions due to the context sensitivity of ‘know’ similar to the one described as taking place in the Bank Case.

(6b) Smith doesn’t know that the plane will stop in Chicago.

On the other hand it is less clear how this case can be accommodated by sensitive invariantism. Since it is the situation of the subject rather than the context of the speaker (i.e. the context of utterance) which is important according to sensitive invariantism, it is unclear why Mary and John’s denial of knowledge to Smith should come out true at the same time as Smith’s self-ascription of knowledge appears to be true. Even though much is at stake for Mary and John, little is at stake for Smith, and according to sensitive invariantism it is his practical concerns that are important for the purpose of determining whether (6b) is true. As a result, the Airport Case appears to pose problems both for classical and sensitive invariantism.¹⁸

¹⁸ One should be careful about ruling out the possibility that there are strategies for accommodating third-person cases such as the Airport Case that are available to proponents of sensitive invariantism, which do not extend as far as giving classical invariantists an account of the original Bank Case. Such strategies are discussed by both Hawthorne (2004, ch. 4) and Stanley (2005, p. 98-104), but I will not discuss them here as my concern is not with sensitive invariantism, but with contextualism.

4 Expanding the Range of the Data

The empirical question of whether the contextualists are correct in their claims about what intuitions competent language users have about the Bank Case and the Airport Case is beyond the scope of this thesis. Still, I will raise a number of complicating factors and issues with respect to the data presented by the contextualists and this is my main goal in this section. In section 4.1 I look at more data about the our use of ‘know’ in order to illustrate that there is more to our linguistic practices involving ‘know’ than what is revealed by just focusing on the Bank Case and the Airport Case. I then turn to a different sort of cases in section 4.2. These cases resemble the original cases, except that the conversational participants are ignorant of the high stakes involved. Such cases have been thought to present a problem for contextualism and I will discuss whether and how contextualists can accommodate them. Section 4.3 deals with the worry that the cases presented in section 3.1 and 3.2 are not psychologically realistic because they ignore the effect of the more pressing practical concerns on the confidence of the conversational participants, and that they therefore fail to provide secure evidence for contextualism.

4.1 More Data about Our Use of ‘Know’¹⁹

In presenting their cases, contextualists tend to focus on a limited set of features of our linguistic practices involving ‘know’. The main focus of contextualists is typically with straightforward ascriptions like (5a) and denials like (5b).

(5a) I know that it’ll be open.

(5b) Well, no, I don’t know that it’ll be open.

However, I think it is useful to look beyond straightforward sentences like these when examining the evidence for contextualism. My goal in this section is fairly modest insofar as I will only be presenting a very limited range of examples, and the examples are also fairly simple. Moreover, the primary intention is neither to strengthen, nor to undermine, the case for contextualism. Rather, I hope that this will serve to bring out some of the complexities

¹⁹ Thanks to John Hawthorne for very useful suggestions and discussion concerning the examples in this section.

involved in our linguistic practices involving ‘know’ that might be ignored if one focuses solely on the cases presented by Cohen and DeRose.

Taking the Bank Case as a starting point, and focusing on Bank Case B where the practical concerns are more pressing, I think it is interesting to note that there are other sentences than (5a) and (5b) that it would be plausible to think of as uttered by DeRose given his situation. For instance, consider (7a) as uttered by DeRose in the context described in Bank Case B.

(7a) I think I know that it’ll be open, but I’d better go in and make sure.

There does not appear to be anything obviously wrong with uttering (7a) in a context like the one described in Bank Case B, and though intuitions about the truth value of (7a) may differ, it seems clearly better than (5a). But (7a) is not the only alternative to (5b) in the conversation in Bank Case B. For instance we may also consider (7b) as uttered by DeRose in Bank Case B.

(7b) I think it’ll be open, but I’d better go in and make sure.

In the case of (7b) the subject avoids using ‘know’ altogether, but it may still be acceptable given that the main question DeRose and his wife are concerning with is whether they should deposit their paycheques right away or whether they can wait until Saturday.

Furthermore, an interesting type of case is provided by hedged sentences. For instance, we may follow Stanley (2005, p. 37 n. 3) in treating ‘really’ as it occurs in (7c) as a hedge.

(7c) I don’t really know that it’ll be open.

According to Paul Kay (1997, p. 134) hedges such as ‘technically’, ‘loosely speaking’, and in this case ‘really’, are typically used to make comments on the utterances or sentences in which they occur.²⁰ Applying this to the case of ‘know’, Stanley comments that ‘in so using ‘really’, one concedes the infelicity of asserting that one knows the proposition in question. Note that this is consistent with it being perfectly true throughout that one knows that proposition’ (2005, p. 37 n. 3). According to this treatment of ‘really’ as it occurs in sentences

²⁰ Kay (1997) does not mention ‘really’ as his discussion focuses on the hedges ‘technically’ and ‘loosely speaking’, but note that ‘really’ is included in George Lakoff’s (1973, p. 472) list of English hedges.

like (7c), an utterance of the hedged sentence (7c) contains a comment on the inappropriateness of asserting the unhedged (5a). Again, (7c) seems to be acceptable given the circumstances in Bank Case B.

It is worth noting that Cohen, given the treatment of ‘really’ adopted here, describe Mary and John as making a hedged claim when discussing Smith’s situation in the Airport Case. Cohen has them agreeing that ‘Smith doesn’t really *know* that the plane will stop in Chicago’ (1999, p. 58). It is difficult to say whether this has any effect on our intuitions about the cases, and more precisely, whether it makes us more likely to accept Mary and John’s claim as true. All the same, I will continue to treat Mary and John as uttering the unhedged sentence (6b).

(6b) Smith doesn’t know that the plane will stop in Chicago.

These considerations do not directly undermine the intuitions that contextualists claim that we have about (5a) and (5b). The point is rather to reveal some of the complexities involved in our use of sentences containing ‘know’, and given the alternatives, it is unclear what is the most natural candidate in the context described in Bank Case B. In his own description of Bank Case B, DeRose has himself making a strong claim by going as far as uttering (5b). However, it is not obvious that this is more, or even as natural, as the weaker alternatives (7a) to (7c). But unlike (5b), the alternatives (7a) to (7c) are consistent with (5a) being true as uttered in the context in Bank Case B. The contextualist can still maintain that our intuitions tell us that (5a) is false, and (5b) true, as uttered in Bank Case B, but I think that a theory which tries to draw support from our ordinary use of ‘know’ should take a broader range of data into consideration.

4.2 Ignorant High Stakes

There is another type of case that contextualists have tended not to focus on. In these cases there is a lot at stake, but the conversational participants are not aware of this. Stanley (2005) discusses such cases under the name ‘Ignorant High Stakes’. Using the original Bank Case as a starting point, one might consider a version of Bank Case B where there is still a lot at stakes, but where DeRose and his wife are unaware of the more pressing practical concerns and consequently the need to deposit their paycheques before Monday. In such a case DeRose might proceed to utter (5a), as he is not aware of the higher stakes.

(5a) I know that it'll be open.

However, Stanley argues that given the higher stakes involved in the context described in Bank Case B, our intuition is that DeRose's utterance is false despite him being ignorant of the increased costs of being wrong.²¹ Similarly in the Airport Case, where Mary or John might utter (8) because they are unaware of the importance of the plane stopping in Chicago.

(8) Smith knows that the plane will stop in Chicago.

Again, (8) is intuitively false. But it is not obvious how cases like this fit into a plausible contextualist picture and whether they can be effectively used as a part of the contextualists' argumentative strategy. Stanley (2005, p. 25) argues that since the conversational participants in this case are ignorant of the increased stakes, the contextualists would have to hold that the contribution of the expression 'know' to the truth conditions of the utterance of (5a) must be determined independently of the beliefs and intentions of the conversational participants. If the contribution of 'know' to truth conditions of the utterance of (5a) is to be determined by the intentions of the conversational participants, this would make the case in which the conversational participants are ignorant of the higher stakes similar in this respect to the original version of Bank Case A in which stakes are in fact low. We should then expect the same verdict with respect to the truth value of both DeRose's utterance of (5a) in the original version of Bank Case A, and the utterance of (5a) in the Ignorant High Stakes case where there is a lot at stake, but the conversational participants are unaware of this. However, Stanley maintains that while the former utterance is intuitively true, the latter utterance is intuitively false.

The problem that Stanley raises with respect to contextualism and these cases is that it is problematic for the contextualists to take the contribution of the expression 'know' to truth conditions to be fixed independently of the intentions of the conversational participants in this way. According to Stanley, this would be a departure from the standard view of how context sensitive expressions work. He claims that 'the contextualist could accommodate Ignorant High Stakes, but only at the cost of advancing a rather dramatic claim about the potential

²¹ Stanley has a slightly more detailed presentation of a version of the Bank Case in which the participants are ignorant of the high stakes (2005, p. 5).

semantic effects of non-psychological facts about extralinguistic content' (Stanley 2005, p. 26).

Stanley admits that there are context sensitive expressions, the contribution of which to the truth conditions of utterances of sentences that contain them is arguably determined independently of the intentions of the conversational participants, but argues that there are only a few plausible candidates. 'But even those who accept the existence of such a class agree that the list is very small, perhaps restricted to 'I', 'yesterday', 'tomorrow', and a few other terms' (Stanley 2005, p. 26 n. 3). Though Stanley does not pursue this further, there might be room for dispute even in the case of expressions like 'today', 'tomorrow' and 'yesterday'. John Perry (2001, p. 61) notes that the introduction of time zones and daylight savings time makes even these expressions a matter of debate if one views time zones and daylight savings time as political matters as opposed to viewing the reference of these expressions as being determined solely by the time and place of the utterance.

To this we may add that the indexical 'I' is associated with a very simple rule which determines the referent in a particular context.²² An utterance of the expression 'I' refers to the speaker in the context of utterance. In the case of 'know' on the other hand, matters seem much more complex. The features of the context on which the exact contribution of 'know' to the truth conditions of utterances of sentences that contain it depends, such as the practical stakes involved in the context, seem less amenable to being captured by such a simple rule. This suggests that the case of 'know' is different from that of 'I' at least in some respects which might discourage someone who wants to model the alleged context sensitivity of 'know' on the context sensitivity of 'I' with respect to the referent of the expression 'I' in a particular context being determined independently of the intentions of the conversational participants.

Of course, this point should not be pushed too far. It rests on the assumption that the high degree of complexity involved in the case of 'know' makes an account in terms of the intentions of the conversational participants more natural, and this may not be obvious. Furthermore, these issues may be more complicated than Stanley allows in his argument against contextualism. There are other expressions that contextualists may point to in defending a view according to which the contribution of 'know' to the truth conditions of utterances of sentences that contain it is fixed independently of the intentions of the conversational participants. For instance, David Kaplan originally defended the view that the

²² We may think of this as related to what is called 'character' in David Kaplan's (1989a) terminology. According to Kaplan, character is a function from contexts to contents.

referents of demonstratives, ‘that’ being a prime example thereof, are determined by a demonstration, that is, ‘typically, though not invariably, a (visual) presentation of a local object discriminated by a pointing’ (1989a, p. 490). The role of pointing gestures in determining the referent of a demonstrative is emphasized by Colin McGinn who suggests the following rule: ‘the referent of a token of ‘that F’ is to be the first F to intersect the line projected from the pointing finger, i.e. the F at the place indicated – on might almost say geometrically – by the accompanying gesture’ (1981, p. 163). Though Kaplan later changed his mind, arguing that the role he previously assigned to a demonstration in determining the referent should instead be played by a special sort of intention on the part of the speaker he called a ‘directing intention’ (Kaplan 1989b, p 582), this does not mean that his earlier view on the matter was altogether implausible. The fact that these matters are quite complex can be illustrated by considering an example, modified from Kaplan’s (1978, p. 335).²³ Consider (9) as uttered by Kaplan while pointing at a picture on the wall behind him.

(9) That is a picture of one of the greatest philosophers of the twentieth century.

But whereas Kaplan thinks that the picture is a picture of Rudolf Carnap, the philosopher, the picture of Carnap has been replaced by a picture of Spiro Agnew, the Vice President of the United States under Richard Nixon. While Kaplan intends to refer to a picture of Carnap, he is pointing at a picture of Agnew. According to Kaplan, he is intuitively talking about a picture of Agnew and not a picture of Carnap. Kaplan originally took as this as suggesting that it is the demonstration, and not the intention, that determines the referent. It is debatable whether this provides any sort of conclusive evidence in favour of Kaplan’s original view, but at least it illustrates how complex these issues are. A view where the intentions of the conversational participants play a less significant role should therefore not be dismissed out of hand, at least not in the case of demonstrative expressions. With this in mind it is worth asking whether Stanley is warranted taking the contextualists to be ‘advancing a rather *dramatic* claim about the potential semantic effects of non-psychological facts about extralinguistic content’ (my emphasis) if they adopt the view that the contribution of ‘know’ to the truth conditions of utterances of sentences that contain it is determined independently of the intentions of the conversational participants.

²³ Kaplan’s (1978) original example involved a new demonstrative, ‘Dthat’, that he introduced in the same paper.

Matters become even more complex when considering the possibility that contextualists may adopt an intermediate position. If an intermediate position is available, contextualists may allow that the intentions of the conversational participant to play a role in determining the contribution of 'know' to the truth conditions of utterances of sentences that contain it, and still be able to accommodate the Ignorant High Stakes cases. According to such a view, the role of intentions is constrained by other features of the context such as the actual, as opposed to the perceived, practical stakes involved. This gives the intentions of the conversational participants some, but limited, control over the contribution of 'know' to the truth conditions of utterances of sentences that contain it, thus preventing the intentions of the conversational participants from determining truth conditions that are inappropriate given the practical stakes involved in that context. According to this view, utterances of (5a) would have the same, or at least similar, truth conditions in the Ignorant High Stakes version of Bank Case B and in the original version of Bank Case B. Consequently, (5a) would be false as uttered in both cases. I have only begun to sketch an outline of such a view. Whether it is ultimately a plausible view, and whether and how it can be worked out in more detail, are further questions that I will not pursue here. It is sufficient for present purposes to note that the Ignorant High Stakes cases do not force contextualists to adopt the stronger view that the intentions of the conversational participants do not play any role in determining the contribution of 'know' to the truth conditions of utterances of sentences that contain it.

DeRose takes a rather cautious attitude towards cases in which the conversational participants are ignorant of the high stakes involved. He says, 'In fact, I think that speakers are free to use standards even wildly inappropriate to the practical situation they face – for instance, to use low standards when they face an extremely high stakes situation in which it would be much wiser for them to employ much higher standards' (DeRose 2005, pp.177-178). On the one hand these remarks suggest that he does not take the contribution of 'know' to the truth conditions of utterances of sentences that contain it to be determined independently of the intentions of the conversational participants. However, DeRose (2005, p. 178) goes on to warn that our intuitions about such cases might be distorted as a result of mistaking the impropriety of using inappropriate standards for the utterance being actually false. But there seems to be a danger associated with this sort of response to the Ignorant High Stakes cases, namely that similar considerations will undermine the evidence for contextualism provided by our intuitions about the cases that favour contextualism, like DeRose's Bank Case

However, it should be noted that the Ignorant High Stakes version of Bank Case B differs from the original version in that the latter involves the intuition that DeRose's utterance of (5b) is true and not just that his utterance of (5a) is false.

(5b) Well, no, I don't know that it'll be open.

It may be argued that this difference makes our intuitions about DeRose's utterance in the original version of Bank Case B more reliable than our intuitions about his utterance in the Ignorant High Stakes version. This derives from DeRose's (1999, p. 199; 2002, pp. 192-193; 2005, pp. 174-175) general preference for relying on intuitions about the truth of utterances rather than intuitions about the falsity of utterances, as he takes the former to be more secure than the latter. But the intuitions appealed to with respect to the Ignorant High Stakes cases are intuitions about the falsity of utterance, rather than intuitions about truth as in the regular versions of the Bank Case and Airport Case. For instance, in Bank Case B, the contextualists argue that not only would an utterance of (5a) by DeRose be intuitively false, but also that DeRose's utterance of (5b) is intuitively true. But as long as DeRose is ignorant of the higher stakes it would appear unnatural to have him utter (5b), making it difficult to appeal to intuitions about truth when it comes to the Ignorant High Stakes cases. I will return to the issue of this alleged asymmetry between intuitions about truth and intuitions about falsity in section 6.2 where I will undertake a more thorough discussion of this methodological assumption, and DeRose's attempt to motivate it. For now, I will be content with noting that it is not clear that this distinction should be made to carry as much weight as DeRose wants it to carry.

While I do not take these considerations to be conclusive with respect to how the Ignorant High Stakes cases fit into the contextualist story, I think that these considerations illustrate how general assumptions about the nature of context sensitivity in natural languages may impinge on the debate about the alleged context sensitivity of 'know', and potentially constrain how contextualism is defended and developed. More specifically, these issues raise some interesting questions about the extent to which the workings of the alleged context sensitivity of 'know' must be similar to the workings of other context sensitive expressions. If one is willing to accept that the workings of the context sensitivity of 'know' are different from the workings of most other context sensitive expressions, and that the case of 'know' is somehow special in this respect, one might find it less problematic to take the contribution of

'know' to the truth conditions of utterances of sentences that contain it to be determined independently of the beliefs and intentions of the conversational participants.²⁴

4.3 Are the Cases Psychologically Realistic?

One matter which deserves some attention at this point is the worry that the Bank and Airport Cases are not psychologically realistic as presented by the contextualists. The issue is the stipulation, made explicit by DeRose in his presentation of the Bank Case, that the subject's level of confidence remains the same when moving from a context where there is not much at stake to a context where the stakes are higher. But this stipulation may not be psychologically realistic. As Kent Bach remarks about the Bank Case, 'It seems to me that unless he's trying to placate his wife, his belief would have to be shaken somewhat' (2005, p. 76). One way of formulating this worry considered, but ultimately rejected by Stanley (2005, pp. 6-7), is that the more pressing practical concern causes the subject's degree of belief in the proposition in question to fall to a point below what is needed for knowledge, or perhaps that the loss of confidence somehow defeats the subject's evidence. Applied to the first-person Bank Case this provides an account of the intuition that it is true for DeRose to deny knowledge to himself in Bank Case B. The higher practical stakes involved in the context described in Bank Case B compared to the context described in Bank Case A causes his confidence to be shaken, thus ensuring that it is no longer true to ascribe knowledge to him with respect to the proposition in question.

But as Stanley is quick to point out, this account cannot be extended to third-person cases such as the Airport Case. In the Airport Case there is no reason to suppose that there is any change in the subject's level of confidence, simply because Smith is not in a situation in which there are any particularly pressing practical concerns. While there is more at stake in Mary and John's context there is no reason to think that this would undermine Smith's confidence. Thus this way of formulating the worry would not undermine the effectiveness of the Airport Case as a means of providing evidence for contextualism.

However, Bach (2005, pp. 76-80) has argued that there is a way of spelling out the initial worry that would apply to third-person cases as well as first-person cases, but which is not explicitly discussed by Stanley. Rather than focusing on the subject's degree of belief,

²⁴ Indeed, one might take precisely the sort of Ignorant High Stakes cases under discussion to motivate a version of contextualism according to which the contribution of the expression 'know' to the truth conditions of utterances of sentences that contain it is determined independently of the beliefs and intentions of the conversational participants.

Bach proposes to start from the assumption that one must confidently believe the proposition that p in order to ascribe to someone else knowledge that p , and that one must have some doubts about p in order to deny knowledge that p to someone who is in a similar epistemic position as oneself. This starting point is based on the relatively uncontroversial further assumption that knowledge implies truth. As Bach argues, ‘In general, you can’t coherently assert that someone else knows that p if you are not confident that p and think that it still needs to be verified’ (2005, p. 77).

Applied to the Airport Case this explains why Mary and John deny knowledge to Smith. This is because Smith is relying on the same evidence as Mary and John (i.e. the itinerary), and Mary and John are not confident that it has been firmly established that the flight has a layover in Chicago. According to Bach, the higher stakes and the more salient possibilities of error have the effect of raising the ascriber’s threshold for confidently believing. This threshold marks how strong reasons one needs in order to confidently believe the proposition in question. Sometimes this involves having more evidence than is required of the subject in order for it to be true to ascribe knowledge to him or her. Bach argues that this is what happens in the Airport Case and that this offers a plausible account of our intuitions about Mary and John’s denial of knowledge to Smith.

However it is worth noting that these considerations would not apply to Ignorant High Stakes cases considered in section 4.2. The account offered by Bach cannot account for the intuition that (5a) would be false as uttered in the context in the Ignorant High Stakes version of Bank Case B or the intuition that (8) would be false as uttered by Mary or John in their context in the Ignorant High Stakes version of the Airport Case.

- (5a) I know that it’ll be open.
- (8) Smith knows that the plane will stop in Chicago.

This is because the ascriber is ignorant of the higher stakes and thus there is no reason to suppose that the ascriber’s confidence is shaken and that his or her threshold for confidently believing is raised. This suggests that more than Bach’s account is needed in order for us to properly account for our intuitions about the variation in our use of ‘know’ revealed by the various cases under discussion. But as the discussion in section 4.2 demonstrates, it is not clear whether and how these cases fit into a plausible contextualist picture of how the contribution of ‘know’ to the truth conditions of utterances of sentences that contain it is determined. Therefore it is also unclear whether the Ignorant High Stakes cases can be made

to fit into a contextualist argumentative strategy. If contextualists adopt the methodological assumption defended by DeRose that intuitions about truth are more reliable than intuitions about utterances, there would also be a methodological obstacle facing contextualists wanting to use the Ignorant High Stakes cases to alleviate the worry that their cases are not psychologically realistic. This makes the question about how the Ignorant High Stakes cases fit into a contextualist picture even more significant.

5 Contextualism and Context Shifting Arguments

So far I have presented DeRose's Bank Case and Cohen's Airport Case along with some initial complications. In the present section I will focus on the argument that DeRose and Cohen rely on in taking our intuitions about these cases to provide evidence for contextualism. While I have already presented the Bank Case and the Airport Case along with the reasoning employed by the contextualists in arguing that these cases provide evidence for contextualism in section 3.1 and 3.2, my goal in the present section is to identify some of the assumptions behind their argument, and to look at how the argument for the context sensitivity of 'know' relates to some questions concerning context sensitivity in general.

DeRose makes a confident claim about the extent to which cases such as the Bank Case provide evidence for the thesis that the expression 'know' is context sensitive:

The best grounds for accepting contextualism concerning knowledge attributions come from how knowledge-attributing (and knowledge-denying) sentences are used in ordinary, non-philosophical talk [...] This type of basis in ordinary language provides not only the best grounds we have for accepting contextualism concerning knowledge attributions, but, I believe, is evidence of the very best type one can have for concluding that any piece of language is context sensitive. (2005, p. 172).

If DeRose is correct then Cohen and DeRose's argument for contextualism based on the cases presented in section 3.1 and 3.2 provides strong evidence that 'know' is a context sensitive expression. In my discussion of the argument that Cohen and DeRose rely on in taking our intuitions about the Bank Case and the Airport Case to provide evidence for the context sensitivity of 'know' I will focus on Cappelen and Lepore's (2003; 2005b) discussion of what they refer to as 'context shifting arguments'. If correct, these considerations might raise doubts about DeRose's statement.

A context shifting argument involves describing different utterances, which are imagined as taking place in different contexts, of a sentence containing a suspected context sensitive expression. If one has intuitions about these utterances having different truth conditions or truth values, or different things being said or expressed by these utterances, that is taken as showing that the expression is context sensitive, assuming that our intuitions cannot be attributed to ambiguity, ellipsis, vagueness or any other context sensitive features of the sentence in question such as tense or other context sensitive expressions.

This description of context shifting arguments fits very well with what is going on in the case of Cohen and DeRose's argument for contextualism based on the Bank Case and the Airport Case. While Cohen and DeRose do not describe different utterances of the same sentence containing 'know', but rather pairs of sentences, the basic argument is similar. Cohen and DeRose appeal to intuitions about variations in the truth values of utterances of sentences containing 'know', and based on these intuitions they conclude that 'know' is a context sensitive expression.

However, Cappelen and Lepore raise an important problem concerning how context shifting arguments are used. The problem is that the theorists employing context shifting arguments ignore the role of their own context, the context in which the context shifting argument is presented. 'To not make the context of the thought experiment an essential variable of the experiment is like trying to measure the speed of objects around you while ignoring your own speed' (Cappelen and Lepore 2005b, p. 10).

Cappelen and Lepore distinguish between the storytelling context in which the context shifting argument is presented and the target context which is described as a part of presenting the context shifting argument. Based on this distinction they draw a further distinction between two different kinds of context shifting arguments depending on whether the suspected context sensitive expression is used or only mentioned in the storytelling context. On the one hand there are what they refer to as 'Impoverished Context Shifting Arguments'. When presenting an Impoverished Context Shifting Argument, the suspected context sensitive expression is mentioned, but not used in the storytelling context. In this case the suspected context sensitive expression is only described as being used in the target context (Cappelen and Lepore 2005b, p. 107). It is the use of Impoverished Context Shifting Arguments that Cappelen and Lepore find problematic, because such context shifting arguments ignore the role of the storytelling context and only rely on intuitions about the use of the suspected context sensitive expression in the target context. Consequently, these Impoverished Context Shifting Arguments do not provide reliable evidence that an expression is context sensitive according to Cappelen and Lepore.

On the other hand there are what they call 'Real Context Shifting Arguments'. When giving a Real Context Shifting Arguments for a suspected context sensitive expression, the expression is both used and mentioned in the storytelling context. The suspected context sensitive expression is first used in the storytelling context, and is then mentioned in the storytelling context when it is described as being used in the target context (Cappelen and Lepore 2005, p. 107). According to Cappelen and Lepore, Real Context Shifting Arguments

are reliable indicators of context sensitivity, and if we can provide a Real Context Shifting Argument for a suspected context sensitive expression, that is strong evidence for the context sensitivity of the expression in question.

To illustrate how a Real Context Shifting Argument for the indexical expression ‘I’ might look like, consider (10).

- (10) I am a philosophy student, but when my friend Malthe who studies medicine and not philosophy says, ‘I am a philosophy student’ he says something false.

This seems to be a good candidate for a Real Context Shifting Argument. The expression ‘I’ is used in the storytelling context as well as mentioned in describing its use in the target context. Furthermore, and crucially, (10) is intuitively true, and there does not seem to be anything strange about it.

Applying this distinction between Real and Impoverished Context Shifting Arguments to Cohen and DeRose’s arguments for contextualism, it is clear that they are offering an Impoverished Context Shifting Argument. The expression ‘know’ as it occurs in DeRose’s description of the Bank Case and Cohen’s description of the Airport Case is described as being used in the target context (i.e. the context described in the Bank Case or the context described in the Airport Case), but is only mentioned in the storytelling context. This is evident from the fact that the expression ‘know’ only occurs between quotation marks.²⁵

Furthermore, a Real Context Shifting Argument does not appear to be forthcoming in the case of ‘know’ as the following attempt by Cappelen and Lepore to construct such an argument illustrates:

Known Rupert. Right now, I’m doing philosophy and thinking about Rupert. Rupert, however, is not now doing philosophy. Instead, he’s home making tea. Rupert doesn’t know he is 30 years old. For Rupert to know he is 30 years old, he has to rule out the possibility that he is a brain in a vat. Rupert, however, is unaware of (or not thinking about) this possibility. And so he’s ignoring a possibility that must be ruled out in order for anyone to know anything at all. Still, when Rupert utters in the comfort of

²⁵ Strictly speaking, as Cappelen and Lepore are aware, Cohen’s presentation of the Airport Case raises some questions with regards to these issues. In fact, the expression ‘know’ appears to be used in the storytelling context as it occurs outside quotation marks. For instance, Cohen writes, ‘Mary and John agree that Smith doesn’t really *know* that the plane will stop in Chicago’ (1999, p. 58). But as Cappelen and Lepore (2005b, p. 120-121) point out it is difficult of making sense of Cohen’s presentation of the Airport Case if ‘know’ is to be understood as taking on the value it is has in the storytelling context as this would have Mary and John agreeing and disagreeing with utterances in the storytelling context. In fact, Cohen (1999, p. 65) admits that he prefers to speak in the object language rather than use the appropriate metalinguistic locutions, due to stylistic considerations, but emphasizes that one should not be misled by this. Therefore I think it is safe to conclude that Cohen is offering an Impoverished Context Shifting Argument for ‘know’.

his home, 'I know I am 30 years old' what he says is true, because he's ignoring this possibility, even though this possibility has to be considered for Rupert to know anything at all. (2005b, p. 110)

In this case the expression 'know' is used in the storytelling context in addition to being mentioned in describing its use in the target context, (i.e. Rupert's context). However, this does not appear to be a plausible candidate for a real context shifting argument simply because the story appears to be intuitively false. Note the difference between this story and (10) in this respect. Thus it appears that contextualists must rely on impoverished context shifting arguments when attempting to provide evidence for the thesis that 'know' is a context sensitive expression.

Cappelen and Lepore take the central, and according to their view, faulty, assumption behind the use of Impoverished Context Shifting Arguments to be what they refer to as 'The Mistaken Assumption'. The Mistaken Assumption, as formulated by Cappelen and Lepore, states that²⁶

A theory of semantic content is adequate just in case it accounts for all or most of the intuitions speakers have about speech act content, i.e., intuitions about what speakers say, assert, claim, and state by uttering sentences. (2005b, p. 53)

I have presented the argument for contextualism in terms of intuitions about the truth values, and the truth conditions, of utterances of sentences containing 'know'. While Cappelen and Lepore formulate The Mistaken Assumption in terms of intuitions about what speakers say, assert, claim, and state, it can be extended to intuitions about the truth values and truth conditions of utterances as well. Put in this terminology, though somewhat weaker than in the original statement, the assumption underlying the use of Impoverished Context Shifting Arguments is that if such arguments trigger intuitions about the truth conditions of an utterance of a sentence in English, then these are semantically significant intuitions, intuitions that a semantic theory of English must accommodate when assigning truth conditions to the utterance in question.

The following discussion (section 6) will focus on a more specific alternative to a semantic treatment of our intuitions concerning the Bank Case and the Airport Case. However, Cappelen and Lepore's claim that The Mistaken Assumption is a critical assumption behind the use of Impoverished Context Shifting Arguments demonstrates how

²⁶ Cappelen and Lepore (2005b, p. 53-54) entertain two different formulations of The Mistaken Assumption, but I only consider the weaker, less precise, formulation here.

the assumption that such intuitions reveal a variation in the truth conditions of utterances of sentences containing ‘know’ may be challenged on a more general level.

One might object that contextualists do not need to accept The Mistaken Assumption as a general assumption, but only need to argue that it is the sort of intuitions that they appeal to in giving a context shifting argument for ‘know’ that are semantically significant. However, this raises the problem of drawing a principled distinction between the sort of intuitions about what speakers say, assert, etc. that a semantic theory must account for, and the sort of intuitions that are irrelevant when doing semantics. Cappelen and Lepore’s distinction between Real and Impoverished Context Shifting Arguments can be seen as one way of doing this. But it does not appear to be any obvious way of doing this for contextualists who want to continue relying on Impoverished Context Shifting Arguments. In light of this it seems reasonable to attribute a commitment to some version of The Mistaken Assumption to contextualists relying on an Impoverished Context Shifting Argument for ‘know’.

Cappelen and Lepore are not alone in being suspicious of assumptions like The Mistaken Assumption. For instance, Nathan Salmon (1991) has argued that when a speaker utters a declarative sentence which semantically expresses the proposition that p , the speaker typically also asserts that p . But it does not follow that when a speaker asserts that p , that p is the proposition semantically expressed by the utterance of the sentence in question. Inferring semantic content from asserted content in this way would involve committing a fallacy Salmon labels ‘The Pragmatic Fallacy’. But whereas Cappelen and Lepore are concerned with question concerning context sensitivity, Salmon arrives at his conclusion based on considerations involving such issues as Keith Donnellan’s (1966) distinction between attributive and referential uses of definite descriptions, and cases involving incomplete definite descriptions.

On the other hand, it is not only in the case of ‘know’ that theorists rely on The Mistaken Assumption and Impoverished Context Shifting Arguments in order to support their claims about context sensitivity. Contextualists can therefore appeal to how Impoverished Context Shifting Arguments are used to provide evidence that other expressions are context sensitive. As DeRose points out, ‘Moving to context sensitivity, such facts about ordinary usage also provide us with our primary, most important and best evidence that clearly context sensitive terms like ‘tall’ are context sensitive in the way that we take them to be’ (2005, p. 191). DeRose is calling attention to the fact that considerations which are similar to those appealed to by contextualists in the case of ‘know’ also provide evidence that gradable adjectives such as ‘tall’, ‘rich’ and ‘flat’ are context sensitive. Since the view that gradable

adjectives are context sensitive is almost universally accepted, this would seem to strengthen the contextualists' case in arguing that 'know' is context sensitive.

But while it is widely held that gradable adjectives are context sensitive, it is worth noting that Cappelen and Lepore (2005a; 2005b) reject this precisely on the grounds that the context sensitivity of gradable adjectives is only supported by Impoverished Context Shifting Arguments.²⁷ Furthermore, they argue that Impoverished Context Shifting Arguments can in principle be provided for any expression. Their argument involves taking a wide range of different sentences and constructing context shifting arguments involving these sentences (Cappelen and Lepore 2005b, ch. 3). If Cappelen and Lepore's argument is sound, then a contextualist who relies on an Impoverished Context Shifting Argument for 'know' must accept that similar arguments can be used to show that every expression is context sensitive.

I do not presume to be able to settle questions concerning how large the set of context sensitive expressions in English is within the scope of this thesis, nor is it my intention to tackle these issues. Nevertheless I think it is important to note that the question about the extent to which cases such as the Bank Case and the Airport Case provide evidence that 'know' is context sensitive, the main question in this thesis, depend on these more general issues. The arguments in this section suggest a picture according to which contextualism goes together with there being many context sensitive expressions. This is illustrated both by Cappelen and Lepore's claim that similar arguments can be constructed in the case of any expression, and to a lesser extent by DeRose's appeal to how arguments similar to those used by contextualists to argue that 'know' is context sensitive are used to support the view that gradable adjectives are context sensitive. Whether such a picture is plausible as compared to one in which the range of context sensitive expressions is more limited, possibly excluding 'know', is a question I will not attempt to answer here.

²⁷ I am ignoring the possibility that the context sensitivity of gradable adjectives can be supported by appeal to what Cappelen and Lepore (2005b) refer to as 'Incompleteness Arguments'.

6 An Alternative to a Semantic Treatment of the Cases²⁸

Perhaps the most straightforward alternative to a semantic treatment of our intuitions about the Bank Case and the Airport Case is a treatment in terms of warranted assertability.²⁹ Warranted assertability is a matter of whether it would be warranted or appropriate to assert a given sentence in a given context.³⁰ Whereas contextualists argue that our intuitions about the Bank Case and the Airport Case track the truth conditions of the utterances in question, it may be objected that our intuitions might instead be tracking the conditions under which it would be warranted or appropriate to assert these sentences. According to this objection, contextualists are simply mistaken in thinking that our intuitions must concern the former rather than the latter. If this is correct, then the evidence for contextualism would be undermined because there would be no need to treat our intuitions about the Bank Case and the Airport Case as revealing a variation in the truth conditions of different utterances of sentences containing ‘know’. DeRose goes as far as claiming that ‘The chief bugaboo of contextualism has been the concern that the contextualist is mistaking a variability in the conditions of warranted assertability of knowledge attributions for a variability in their truth conditions’ (2002, p. 167). The attempts by DeRose (1999; 2002) to alleviate this worry for contextualism will be central to the following discussion.

My goal in this section is to show that it is harder for contextualists to rule out the possibility that they are mistaking a variation in warranted assertability conditions for a variation in truth conditions than it initially appears to be. I begin 6.1 by presenting an initial characterization of what DeRose calls a ‘warranted assertability manoeuvre’ (‘WAM’ for short) in defence of the view that ‘know’ is not context sensitive. I then present a worry raised by DeRose (1999; 2002) that such WAMs could too easily be used to explain away any putative counterexample to a semantic theory, regardless of how implausible the theory appears to be. DeRose concludes that we must distinguish plausible candidates for a WAM from implausible candidates and introduces three criteria for a plausible candidate for a WAM

²⁸ Some of the material in this section is taken from a paper I wrote in the spring of 2006 as a part of fulfilling the requirements of the class FIL4471 Referanse, kontekst og talehandling II.

²⁹ This is not to say that a treatment in terms of warranted assertability is the only alternative to a semantic treatment. Williamson (2005) has argued that our intuitions about cases such as the Bank Case and the Airport Case may be the result of bias caused by psychological salience effects. See also the discussion of Bach’s (2005) views in section 4.3.

³⁰ ‘Warrant’ is used here in a broad sense. Having warrant is not just a matter of having epistemic warrant.

designed to accomplish this. These criteria will be the focus of section 6.2. In section 6.3 I examine attempts to provide a WAM in defence of invariantism that satisfies DeRose's criteria. In section 6.4 I examine the issue of whether considerations involving cancellability might supplement DeRose's criteria in the contextualist attempt to rule out such a WAM. Based on the discussion in sections 6.2 to 6.4 I argue in section 6.5 that DeRose's criteria are inadequate in so far as the contextualists' argumentative aims are concerned. In order to underwrite this point, and to illustrate the difficulty of providing a plausible set of criteria that satisfies the contextualists' argumentative purposes, I provide a WAM that satisfies these criteria in defence of the view that 'tall' is not context sensitive. I then argue in section 6.6 that competent language users must be sensitive to the semantic workings of their own language in order for cases such as the Bank Case and Airport Case to provide evidence for contextualism. At the end of section 6 I include an appendix (6.7) on an attempt on the part of DeRose (2002) to use considerations about warranted assertability to provide a positive argument for contextualism.

6.1 Contextualism and Warranted Assertability

According to DeRose (2002, p. 171), this objection to contextualism amounts to a WAM in defence of invariantism. Following DeRose's account, a WAM involves explaining intuitions about the truth and falsity of utterances in terms of warranted assertability. The idea behind the use of a WAM is that asserting a sentence may fail to be warranted even though the utterance of the sentence is true, and that there may be contexts in which it is warranted to assert something false. This idea is combined with the further assumption that we are liable to confuse warranted assertability for truth, at least under certain circumstances.

An example might illustrate how truth and warranted assertability may fail to converge.³¹ Suppose someone wants to know whether Phil is a good philosophy student. Moreover, assume that that Phil is indeed an excellent and very talented philosophy student in addition to being very punctual and having beautiful handwriting, and that the person asked knows all of this. However, if the person asked were to answer by uttering (11), even if it is in fact true, that would not be warranted because it would convey, by means of an implicature, that Phil is not a very good philosophy student, at least on the assumption that punctuality and good handwriting are not the most important qualities to look for in a good philosophy student.

³¹ I am using a modified version of one of Paul Grice's (1991a, p. 31) examples of conversational implicature.

(11) Phil is always on time and has beautiful handwriting.

In this case we are probably not very likely to confuse the fact that one cannot warrantably assert (11) for an utterance of (11) being false, but it is a clear example of a sentence which cannot be warrantably asserted in a certain context, despite the sentence being true as uttered in that context.

Starting with how this may be applied to the Bank Case, a WAM may target the intuitions about the utterances in either Bank Case A or Bank Case B. This involves either explaining the truth of (5a), and the corresponding falsity of (5b), in Bank Case A as a result of mistaking the warranted assertability of (5a) for the utterance being true, or taking the truth of (5b), and the corresponding falsity of (5a), in Bank Case B to be explained by DeRose being warranted in asserting (5b) in that context.

(5a) I know that it'll be open.

(5b) Well, no, I don't know that it'll be open.

Similarly with respect to the Airport Case, it may either be that we are mistaking the warranted assertability of (6a) for Smith's utterance being true, or that we are making a similar mistake in the case of Mary or John uttering (6b).

(6a) Yes I know – it does stop in Chicago.

(6b) Smith doesn't know that the plane will stop in Chicago.

For the moment it is sufficient to note the possibility of confusing a variation in warranted assertability conditions for a variation in truth conditions, leaving it an open question whether it is our intuitions about (5a) and (6a) or the intuitions about (5b) and (6b) that are confused in this way.

This characterization also leaves the details of how the WAM is executed open. However, it is worth noting that DeRose focuses on WAMs involving an appeal to implicatures, at least insofar as this is what he focuses on in his examples. Therefore, I will make WAMs involving implicatures the focus of the discussion here as well. With this in mind, it seems natural to put the issue more broadly in terms of whether our intuitions about

the cases admit of a post-semantic, pragmatic treatment, rather than the semantic treatment defended by the contextualists.

DeRose objects against the use of a WAM in defence of invariantism on the grounds that such a strategy would not only provide invariantism with a defence against the cases presented by the contextualists, but could be used in defence of any semantic theory facing putative counterexamples. Stanley echoes DeRose's concern. He claims that 'By undermining the data for semantic theory, this kind of strategy threatens to undermine the semantic project' (Stanley 2005, p. 14). The worry raised by DeRose and Stanley is that unrestricted use of WAMs threatens to undermine the evidence for semantic theories in general. If it can always be claimed that our intuitions track warranted assertability conditions rather than truth conditions we would have no reliable way of providing evidence for semantic theories. DeRose (2002, p. 174) provides several examples which illustrate the general point. We can for instance imagine a semantic theory of the expression 'is a bachelor' that entails that utterances of sentences like (12) and (13) have the same truth conditions.

(12) Bill is a bachelor.

(13) Bill is a man.

According to this theory it is not a part of the truth conditions of (12) that Bill must be unmarried. A theory which omits this part of the truth conditions of (12), and entails that (12) is true even if Bill is married just as long as Bill is a man, does not appear to be very plausible. Such a theory would face a long list of counterexamples involving married men, of whom it would appear false to say that they are bachelors. But DeRose points out that the theory might be defended by offering a WAM in its defence. Thus it may be argued that uttering (12) generates an implicature to the effect that Bill is unmarried. Therefore it is only a part of the warranted assertability conditions of (12) that Bill must be unmarried, but not a part of the truth conditions of utterances of (12). This means that an assertion of (12) is unwarranted if Bill is married, but not false. On the basis of these considerations DeRose concludes that we should not adjust our original estimate of the plausibility of the theory in question, but rather question the legitimacy of the WAM made in its defence.

6.2 Criteria for a Plausible Candidate for a WAM

DeRose does not go as far as to argue that it is never legitimate to appeal to considerations involving warranted assertability when defending a semantic theory. He accepts that there are circumstances in which a WAM is both legitimate and appropriate, but argues that we should distinguish between plausible and implausible candidates for a legitimate WAM. In order to distinguish plausible candidates from implausible candidates, DeRose presents three criteria that must be satisfied for a WAM to be a plausible candidate for a legitimate WAM. These three criteria provide the focus of the discussion in this section.

DeRose (2002, p. 175) points out that one problematic feature of the illegitimate WAM used in defence of a theory that entails that (12) and (13) have the same truth conditions is that it relies on special rules associated with the expression ‘is a bachelor’.

(12) Bill is a bachelor.

(13) Bill is a man.

The contrast DeRose draws is between the sort of WAM that relies on special rules associated with the expression in question on the one hand, and WAMs that exploit general rules of conversation, exemplified by the Cooperative Principle and the conversational maxims formulated by Paul Grice (1991a), on the other. Only instances of the latter sort, those involving appeal to general conversation rules, are candidates for being an effective and legitimate WAM. This is one the criteria that must be satisfied in order to have a plausible candidate for a WAM, according to DeRose, and it effectively imposes a constraint on the sort of WAMs that are legitimate.³²

DeRose makes a connection between this point and the earlier point that a WAM could be used in defence of any semantic theory, no matter how implausible it might initially appear. He points out that ‘it’s not so easy to generate the implicatures you need to deflect the apparent counterexamples to your theory by means of *general* conversational rules that can be tested on very different sentences’ (DeRose 1999, p. 200; 2002, p. 176). By introducing a constraint on the sort of WAMs that one can legitimately appeal to, it becomes harder to

³² I am going to understand DeRose’s criteria as imposing necessary conditions that a WAM must satisfy in order to be a plausible candidate for a legitimate WAM. However, this reading may be too strong. It would for instance rule out the possibility of an explanation in terms of conventional implicature, and one might not want to rule out this in the case of for instance ‘but’ even though one does not necessarily think that ‘but’ should be understood in this way. That being said, I propose, for the purpose of the present discussion, to ignore these complications.

explain away the force of putative counterexamples with the help of WAM, thus removing the threat of such a practice undermining the evidence for semantic theories in general.

DeRose (2002, p. 197 n. 16) takes the distinction between WAMs that rely on special rules and WAMs that rely on general rules, to correspond to Grice's (1991a; 1991b) distinction between conventional implicatures and conversational implicatures.³³ Whereas conventional implicatures are determined by the meaning of the words alone, conversational implicatures depend on the assumption that speakers are observing the Cooperative Principle and the conversational maxims. When a hearer is attempting to work out what was conversationally implicated, she relies on the meaning and the referents of the words used by the speaker, information about the context and background information, the Cooperative Principle and the conversational maxims, and the assumption that the relevant information is available to both the speaker and the audience, (Grice 1991a, p. 31).

The next question is where this leaves us with respect to the case of 'know'. DeRose claims that invariantists have not offered a WAM based on general conversational rules. Thus the WAM made in defence of invariantism fails to satisfy the criterion for being a legitimate WAM. In fact, DeRose (2002, p. 176) argues that invariantists have often settled for pointing out the possibility that our intuitions are a result of mistaking a variation in warranted assertability conditions for a variation in truth conditions, without invoking any considerations involving general rules of conversation to support claim.

DeRose also states further criteria for having a plausible candidate for a WAM. According to DeRose's (1999, p. 198) second criterion, a WAM is more plausible in cases where there are conflicting intuitions. Such cases involve the intuition that both an utterance *u* of a sentence *S* and an utterance *u'* of a sentence not-*S* seems false, or at least somewhat wrong. Consequently, there is a pressure to explain away the intuition of falsity with respect to either *u* or *u'*. DeRose points out that in the case of the expression 'is a bachelor' there is no such conflict. An utterance of (12) is intuitively true if Bill is an unmarried man, and false if he is not an unmarried man. Thus there is no pressure to explain away either of the intuitions.

However, it is not clear that this criterion poses a problem for a WAM made in defence of invariantism. As Jessica Brown (2005, p. 283; 2006, p. 411-413) points out, the Bank Case and the Airport Case appear to involve conflicting intuitions. These cases reveal that an utterance of a sentence containing 'know' may appear to be true as uttered in one context, and false as uttered in another context, despite the subject being in the same, or at

³³ I am ignoring the possibility of implicatures that are neither conventional, nor conversational, throughout this discussion.

least similar, epistemic position with respect to a true proposition in both contexts. While the contextualists offer a means of resolving this apparent contradiction by arguing that ‘know’ is a context sensitive expression, Brown argues that to claim that these cases do not involve conflicting intuitions on these grounds amounts to begging the question against invariantism.

Turning to DeRose’s third criterion, DeRose (1999, p. 199; 2002, pp. 192-193; 2005, pp. 174-175) considers it more difficult to explain away intuitions about the truth of utterances by means of a WAM than it is to explain away intuitions about the falsity of utterances. He argues that asserting a falsehood while implicating something true still counts as unwarranted or inappropriate. ‘For, except where we engage in special practices of misdirection, like irony or hyperbole, don’t we want to avoid falsehood both in what we implicate and (especially!) in what we actually say?’ (DeRose 2002, pp. 192-193). There are many ways in which an assertion may be unwarranted even though it is true, for instance by being irrelevant to the on-going conversation or implicating something false. In these cases we might mistake the assertion being unwarranted for it being false. On the other hand, DeRose claims that there are not correspondingly many ways in which an assertion may be warranted even though it is false.

DeRose claims that this poses a problem for any WAM in defence of invariantism since contextualists appeal to both the intuition that DeRose’s utterance of (5a) in Bank Case A is true, and the intuition that his utterance of (5b) in Bank Case B is true, and not just the intuition that (5a) would be false as uttered in Bank Case B. Similarly, with respect to

- (5a) I know that it’ll be open.
- (5b) Well, no, I don’t know that it’ll be open.

Cohen’s Airport Case, where an invariantist WAM would have to explain away either the intuition that Smith’s utterance is true, or the intuition that Mary and John’s utterance is true. Either way, the invariantists are forced to explain away intuitions about truth, and not just intuitions about falsity (DeRose 1999, p. 201; 2002, p. 193).

One might object that it is not clear that a WAM in defence on invariantism must explain away both intuitions about truth and falsity. Someone taking this line of response may appeal to the sort of data mentioned in section 4.1 and argue that it is not natural for DeRose to go as far as to utter (5b) in Bank Case B, and similarly, with respect to the Airport Case where Mary and John goes as far as uttering (6b). The objector could claim that it would be

more natural for DeRose to utter something like (7a) in a context like the one in Bank Case B (see section 4.1 for more alternatives).

(6b) Smith doesn't know that the plane will stop in Chicago.

(7a) I think I know that it'll be open, but I'd better go in and make sure.

While this line of response may have merit, it turns on finding support for the claim that it is unnatural for DeRose to go as far as to utter (5b). Contextualists claim that DeRose's utterance of (5b) in the Bank Case is intuitively true and natural given the higher stakes in that context.

But is the asymmetry between intuitions about truth and intuitions about falsity plausible in the first place? Stanley (2005, p. 123 n. 8) takes a sceptical attitude towards this alleged asymmetry. He notes that it makes a distinction between intuitions about the truth of assertions of negated propositions on the one hand, and the falsity of assertions of non-negated propositions on the other, and he finds this distinction uncomfortable. Stanley appears to make an intuitive point, but he does not discuss DeRose's motivation for the asymmetry between intuitions about truth and intuitions about falsity.

Brown (2005, p. 284; 2006, pp. 413-415) however, argues that a false utterance may appear true when it is warranted in the sense that what it pragmatically imparts is true, either because the conversational participants focus on what is pragmatically imparted, or are mistaken about the relationship between what is pragmatically imparted and the truth conditions of the utterance. When speakers are concerned with communicating something that is true, it is not clear that they should worry about their utterance being strictly speaking false.

It is possible to take this point further. For DeRose's criterion to be plausible it seems that language users must be sensitive to the truth conditions of utterances, and whether the utterance semantically expresses something true or something false, even when they are focusing on what the utterance pragmatically imparts. Otherwise, there would be no reason not to think that a false utterance could appear to be true because it pragmatically imparts something true. But this comes dangerously close to assuming what is under discussion, namely the contextualist contention that we are not liable to confusing warranted assertability and truth in the cases at hand.

This means that the criterion that a plausible candidate for a WAM should be based on general rules is left to do most of the work in ruling out a WAM in defence of invariantism. The criterion that a WAM should be based on conflicting intuitions seem to be satisfied, and

the criterion that that a WAM should only explain away intuitions about truth does not appear to be plausible.

6.3 A More Plausible Candidate for a WAM in Defence of Invariantism

Patrick Rysiew (2001) has sought to meet DeRose's challenge by providing a WAM that is based on general rules of conversation. While DeRose (2002) raises an objection against Rysiew's account, I argue that the Brown's (2005) attempt to provide a WAM based on general conversational rules avoids this objection. Therefore, I conclude that there is a plausible candidate for a WAM in defence of invariantism that satisfies the criterion that a WAM must be based on general rules.

Rysiew (2001) defends a non-sceptical version of invariantism according to which an utterance of (5a) by DeRose comes out true in both Bank Case A and Bank Case B, and similarly with respect to the Airport Case, in which (6a) is true as uttered in both Smith's context and Mary and John's context.

- (5a) I know that it'll be open.
- (6a) Yes I know – it does stop in Chicago.

However, he argues that the intuition that (5a) is false as uttered in the context in Bank Case B and the intuition that (6a) is false as uttered in Mary and John's context, can be explained without resorting to contextualism or an illegitimate WAM of the sort discussed in section 6.1 and 6.2.

In attempting to explain our intuitions, Rysiew (2001, pp. 487-488) distinguishes between the proposition semantically expressed by an utterance of a sentence and the proposition, or propositions, pragmatically imparted by the same utterance. Whereas an utterance of a sentence containing 'know' semantically expresses only that the subject is in a good enough epistemic position to know the proposition in question, it may pragmatically impart that the subject is in a good enough epistemic position for the purpose of the context in question. It is this distinction that allows Rysiew to explain our intuitions. He claims that while the propositions semantically expressed by the relevant utterances of (5a) and (6a) are true, the propositions pragmatically imparted are false, and it is the latter that explains our intuitions.

Developing his view further, and in more detail, Rysiew (2001, p. 488) adopts an understanding of the strength of the epistemic position of the subject in terms of the alternatives that the subject can rule out. According to this account, the subject of a knowledge ascription must be able to rule out any relevant alternatives to a proposition in order for the knowledge ascription to be true. This view bears important similarities to the versions of contextualism defended by Cohen (1988) and Lewis (1996).³⁴ However, according to Rysiew's position, the range of relevant alternatives is fixed and does not vary with context.³⁵ What counts as a relevant alternative to a proposition is determined by what normal humans treat as a likely counter-possibility to that proposition (Rysiew 2001, p. 488).

What do vary with context according to Rysiew are the salient alternatives. The salient alternatives are those possibilities that the conversational participants have in mind in a particular context. According to Rysiew, 'salience is the occasion-sensitive notion: it picks out, rather indiscriminately, features of speakers' psychologies which accompany their everyday uses of "knows," however idiosyncratic, unusual, and peculiar to the conversational setting' (2001, p. 489). Thus the relevant and the salient alternatives need not be the same. But even though the salient alternatives have no bearing on the proposition semantically expressed by utterances of sentences containing 'know', they have bearing on the proposition or propositions pragmatically imparted.

In applying this framework to the Bank Case, Rysiew (2001, p. 490) argues that DeRose is in a position to rule out any relevant alternatives, and that the proposition semantically expressed by an utterance of (5a) by DeRose would be true in both Bank Case A and Bank Case B. Furthermore, he argues that since the range of relevant and salient alternatives is similar in Bank Case A, the proposition pragmatically imparted by DeRose utterance is also true. On the other hand, in Bank Case B, while the relevant alternatives stay the same, there are other salient alternatives that DeRose cannot rule out, such as the possibility of the bank having changed its hours. An utterance of (5a) by DeRose would therefore pragmatically impart that he is in a position to rule out these alternatives, and consequently the proposition pragmatically expressed comes out false. Conversely, his utterance of (5b) pragmatically imparts the he is not in a position to rule out all the salient alternatives.

³⁴ See section 2.2 for Cohen's (1988) and Lewis' (1996) views.

³⁵ Rysiew does not want to rule out the possibility of some variation in what counts as relevant. For example, he mentions the possibility of a variation associated with 'know' being a vague expression. He therefore takes himself to be proposing 'a "fairly" invariantistic semantics for "know(s)"' (Rysiew 2001, p. 489).

(5b) Well, no, I don't know that it'll be open.

Even though the proposition semantically expressed by DeRose's utterance of (5b) is false, as he is in a position to rule out all the relevant alternatives, the proposition pragmatically imparted by his utterance is true.

Rysiew (2001, p. 491) takes his account to be supported by Grice's (1991a, p. 27) maxim 'Be relevant' under the category of Relation, thereby satisfying DeRose's criterion that a legitimate WAM should appeal to general rules of conversation. What is relevant to the conversation in one context may fail to be relevant in another context and vice versa, and the maxim of relevance calls for speakers to make their contributions relevant to the conversation at hand. Since what is relevant to the conversation in Bank Case B are the salient alternatives such as the bank having changed its hours, we cannot preserve the assumption that DeRose is observing the maxim of relevance unless we take his utterance of (5b) to pragmatically impart that he cannot rule out all the salient alternatives. Rysiew states the general point,

It is, I claim, because speakers strive to conform, and are known to strive, to the maxim of Relation (and more generally, to [the Cooperative Principle]) that an utterance of the form "S knows that *p*" is naturally taken to communicate [...] that *S*'s epistemic position is 'good enough' given the epistemic standards that are operative in the context in question. (2001, p. 491)

He argues that adherence to the maxim of relevance explains why an utterance of sentence containing 'know' pragmatically imparts that the subject is in a good enough epistemic position for the purpose of the context in question, or in terms of his preferred framework, that the subject is in a position to rule out any salient alternatives to the proposition in question. Thus Rysiew makes a general point regarding the explanation of why DeRose's utterance of (5b) pragmatically imparts that he is able to rule out salient alternatives such as the bank having changed its hours. Consequently, he is able to argue that his account is supported by considerations involving general conversational rules such as Grice's maxim of relevance.

However, DeRose is not satisfied with Rysiew's account. DeRose (2002, p. 198 n. 17) understands Rysiew as postulating that utterances of sentences containing 'know' have two meanings, a semantic meaning and a pragmatic meaning (i.e. the proposition semantically expressed and the proposition pragmatically imparted), and DeRose finds this feature of Rysiew's account problematic.

I am highly suspicious of accounts that help themselves to two such meanings in the way Rysiew's does, disliking not only the loss of economy in explanation, but also worrying that we will not be able to combat all manner of absurd theories about the truth conditions of various sentences if defenders of these theories are able to posit separate pragmatic meanings that do the work of accounting for usage in the troublesome cases, allowing their account of the truth conditions to sit safely off in the corner, out of the fray. (DeRose 2002, p. 198 n. 17)

It is not entirely clear how to best understand DeRose's objection to Rysiew. The problem does not appear to be that the proposition or propositions pragmatically imparted by utterance of sentences containing 'know' may be different from the proposition semantically expressed. DeRose accepts that what an utterance semantically expresses may differ from what it is used to pragmatically impart. Rather, DeRose seems to be objecting to Rysiew on the grounds that, according to Rysiew's account, the pragmatic meaning is not explicable as generated by the semantic meaning together with general principles. He goes on to state that, 'Of course, in *some* sense, we all must hold that our assertions carry meanings beyond what goes into their truth conditions. [...] But I find such a claim unobjectionable when the second, pragmatic meaning is generated by the semantic meaning together with general principles' (DeRose 2002, p. 198, n. 17). While Rysiew's account is supported by appeal to Grice's maxim of relevance, DeRose focuses on these considerations as supporting Rysiew's explanation of why we focus on the pragmatically imparted proposition and not proposition semantically expressed, rather than as explaining the proposition pragmatically imparted as the joint outcome of these considerations and the proposition semantically expressed. According to DeRose's reading of Rysiew, the maxim of relevance explains why we focus on the proposition pragmatically expressed as that is what is relevant in the context. However, the maxim of relevance does not explain how and why the pragmatically imparted proposition is generated in the first place. Once again, DeRose raises the familiar worry that allowing explanations that postulate two meanings in the way that Rysiew's account does will undermine the evidence for semantic theories by making it too easy to explain away putative counterexamples.

All the same, it does not follow that the invariantists cannot provide a WAM that does not involve postulating two meanings. Like Rysiew, Brown (2005) tries to develop a WAM in defence of a non-sceptical version of invariantism. But while Brown does not discuss DeRose's argument against Rysiew, her account does not appear to involve postulating two meanings in the sense that DeRose finds objectionable. In developing her WAM, Brown makes use of the notion of strength of epistemic position. According to the invariantist position she defends, how strong the epistemic position of the subject must be for an

ascription of knowledge to the subject to be true does not depend on the context, as the level of strength required is invariant across different contexts.

Once again drawing on Grice's maxim of relevance, it might be that in certain contexts a stronger epistemic position is relevant to the conversation at hand. Take for instance the context in Bank Case B in which an utterance of (5a) by DeRose seems false. According to the picture defended by Brown (2005, p. 281), it is a stronger epistemic position that is relevant in this context. Even though (5a) is true as uttered by DeRose in this context because he is in a strong enough epistemic for it to come out true, it is a stronger epistemic position that is relevant in this context. Asserting (5a) would therefore implicate that DeRose is in the stronger epistemic position relevant in this context, as his assertion would not be relevant to the conversation at hand. Because DeRose is not in the stronger epistemic position the implicature is false and the assertion unwarranted, despite (5a) being true as uttered in that context.

Brown (2005, p. 281) also follows Rysiew in understanding the strength of the epistemic position in terms of the range of alternatives that the subject can rule out.³⁶ In the context in Bank Case B, where there is a lot at stake and DeRose's wife has raised the possibility of the bank having changed its hours, a strong enough epistemic position to rule out possibilities such as this one is relevant. By asserting (5a) DeRose implicates that he is in a strong enough epistemic position to rule out the possibility of the bank having changed its hours, and since he is not, he thereby implicates something false.

The account offered by Brown does not appear to involve two meanings in the problematic sense. The implicatures she appeals to in explaining our intuitions about the cases are explicable as the joint outcome of the proposition semantically expressed, the general conversational rule of relevance, and features of the context. Thus understood her account is not susceptible to the objection that DeRose raised against Rysiew on the ground that Rysiew's account was seen by DeRose as postulating two meanings.

6.4 Cancellability

It therefore seems that there is a more plausible candidate for a WAM in defence of invariantism than DeRose is willing to acknowledge. Nevertheless, contextualists may point to other problems with the WAM proposed by Brown. In this section I will look at the

³⁶ It should be stressed that this is not the only way to develop the invariantist position. Brown (2006, pp. 424-428) has also developed a WAM based on understanding the strength of the subject's epistemic position in terms of the range of possible worlds across which the subject's belief tracks the truth.

question of whether considerations involving the cancellability of conversational implicatures may supplement DeRose's original criteria in an attempt to rule out a WAM in defence of invariantism.

On the assumption that the implicature she appeals to is a conversational implicature, as indeed seems to be the case, we should be able to cancel the implicature. If an utterance of a sentence conversationally implicates that p , the implicature should be cancellable, either explicitly by adding 'but not- p ' or something similar, or contextually by finding a different context in which an utterance of the same sentence would not generate the implicature (Grice 1991a, p. 39; 1991b, p. 44). But this may pose a problem for Brown, since arguably, the putative implicature generated by an utterance of (5a) by DeRose in the context in Bank Case B cannot be explicitly cancelled without making the utterance sound wrong or at least somewhat awkward. Consider (14) as an attempt to explicitly cancel the implicature.

(5a) I know that it'll be open.

(14) I know that the bank will be open, but I can't rule out the possibility of the bank having changed its hours so I'd better go in and make sure.

Is there something wrong with (14) as uttered by DeRose in the context in Bank Case B? Intuitions may differ on this point. Cohen (1999, pp. 59-60) seems to be of the mind that an utterance like of a sentence (14) not only sounds awkward, but goes as far as to argue that 'We know, but we need to investigate further' sounds inconsistent. He therefore argues that the putative implicature is not cancellable, and that the phenomena should instead be given a contextualist treatment. Rysiew (2001, p. 495) and Brown (2006, p. 428) on the other hand, do not find utterances like this uncomfortable, and argue that the putative implicature can in fact be explicitly cancelled. We are therefore left with conflicting intuitions regarding cancellation attempts like (14) with no clear, intuitive answer as to whether it is an acceptable way of cancelling the implicature.

Furthermore, Rysiew and Brown proceed to question the underlying assumption that an implicature must be cancellable without discomfort, arguing that we should not expect all conversational implicatures to be cancellable without discomfort. Brown claims that we should not expect the cancellation to be comfortable in the cases where speakers are mistaken or confused about the relationship between what an utterance semantically expresses and what it pragmatically imparts.

As Rysiew notes, Grice's own remarks do not adequately settle the matter. According to Grice's formulation, 'a putative conversational implicature that p is explicitly cancellable if, to the form of words the utterance of which putitatively implicates that p , it is admissible to add *but not p* , or *I do not mean to imply that p* ' (1991b, p. 44). But this does not make it clear when it is inadmissible to add 'but not p '. Is it enough that the utterance is uncomfortable or does it have to be obviously inconsistent? Rysiew argues that the former condition is too strong. Consider (15b) as an attempt to cancel the implicature generated by an utterance of (15a) to the effect that they fell in love first and then got married.

(15a) They fell in love and got married.

(15b) They fell in love and got married, but not in that order.

But (15b) sounds odd though not contradictory. It therefore appears that this conversational implicature is not cancellable without discomfort, assuming, and this seems quite plausible, that this is in fact a case involving a genuine conversational implicature.³⁷ In the case of (14) it was noted that there are conflicting intuitions as to how uncomfortable it is. While Cohen claims that it sounds inconsistent, it is far from clear that this is the prevailing intuition and given the conflicting intuitions it is hard to argue that it is clearly contradictory.

It is worth adding that the focus of these considerations have been whether the putative implicature is explicitly cancellable. But there does not seem to be a similar problem in taking the putative implicature to be contextually cancellable. For instance, DeRose's original utterance of (5a) in Bank Case A does not implicate that he can rule out the possibility of the bank having changed its hours. This means that to the extent that there is a problem with respect to the putative implicature that Brown appeals to being cancellable, it is only a problem with the implicature being explicitly cancellable. But as I take it that the above discussion demonstrates it is not clear that there is a problem at all.

6.5 A WAM for 'Tail'

I take the conclusion of this discussion to be that DeRose's criteria, even when supplemented by considerations about cancellability, fail to rule out the possibility that our intuitions about the Bank Case and the Airport Case can be given a post-semantic, pragmatic treatment, rather

³⁷ Stephen Levinson (1983, p. 108) for instance, argues that cases involving 'and' and a temporal sequence of the events should be given a treatment in terms of conversational implicature, and not for instance a treatment in terms of 'and' being ambiguous with one of its senses being equivalent to that of 'and then'.

than a semantic treatment in terms of ‘know’ being context sensitive. However, it may be argued that rather than showing that our intuitions about the cases admit a pragmatic treatment, it shows that DeRose’s criteria are inadequate. In order to illustrate this point, I will provide a WAM in defence of the view that ‘tall’ is not a context sensitive expression. But rather than taking this to show that some further set of criteria is needed, I argue that this may instead be understood as illustrating the difficulty of providing a set of criteria that will serve the contextualists’ argumentative purposes.

As was noted earlier (section 5), the view that gradable adjectives such as ‘tall’ are context sensitive is almost, though not quite, universally accepted. If Tim is significantly taller than the average person, even if he is not exceptionally tall, an utterance of (16a) will be intuitively true in a context in which the conversational participants are discussing people of average height.

(16a) Tim is tall.

But (16a) is intuitively false if uttered in a context in which the conversational participants are discussing the height of professional basketball players and this is usually taken to provide evidence that ‘tall’ is context sensitive. An utterance of (16a) can be intuitively true as uttered in one context and intuitively false as uttered in another context. However, it may be argued that this is merely a result of mistaking considerations involving warranted assertability for considerations involving truth. Asserting (16a) in a context in which the topic of the conversation is the height of professional basketball players is unwarranted, even if (16a) is true as uttered in such a context.

Suppose that someone claims that an utterance of ‘S is tall’ is true if and only if S measures more than n meters where n does not vary with context, leaving the question of how the measuring is to be done and similar concerns aside. Suppose further that Tim measures more than n meters. In that case, an utterance of (16a) is true regardless of whether the topic of the conversation is basketball players or ants. Expanding the view under consideration, an utterance of ‘S is tall for a professional basketball player’ is true if and only if S measures more than n_+ meters, where n_+ is a higher number than n , and assume that while Tim measures n meters he does not measure n_+ meters. Therefore an utterance of (16b) is false.

(16b) Tim is tall for a professional basketball player.

In a context in which the conversational participants are discussing the height of professional basketball players what is relevant is whether someone measures n_+ meters, not whether someone measures n meters. Thus an assertion of (16a) would not count as relevant in this context unless the speaker was taken to communicating something else, namely that Tim measures n_+ meters. In other words, asserting (16a) implicates (16b). But since Tim does not measure n_+ meters, (16b) is false. We are therefore in a position to explain the intuition that (16a) would be false as uttered in such a context by appealing to the generation of a false implicature.

This WAM is similar to the WAM used by Brown in defence of invariantism in that both are based on considerations about relevance. Contextualists may take this to indicate that it is this reliance on a maxim of relevance that is problematic. It might thus be argued that considerations about relevance are a too powerful explanatory tool when it comes to WAMs. However, it is not clear how this worry should be spelled out in more detail. It appears ad-hoc to claim that considerations about relevance cannot play a role in supporting WAMs in the way that considerations involving other conversational maxims can. Why should a WAM based on a maxim of relevance be taken to provide a less plausible explanation of our intuitions than a WAM based on another conversational maxim?

While this does not prove that there is no plausible criteria that will serve the argumentative purposes of the contextualist, I think that it illustrates the difficulty of providing such a set of criteria. Furthermore, I am not aware of any plausible criterion that could be used to strengthen or supplement DeRose's criteria in the desired way.

6.6 The Sensitivity of Competent Language Users with Respect to the Semantic Workings of Their Own Language

Contextualists might instead take the difficulties in finding a set of criteria that serves their argumentative purposes to motivate a sceptical stance towards WAMs in general, arguing that cases of conversational implicature are typically such that it is possible for language users to distinguish what an utterance semantically expresses from what it conversationally implicates.³⁸ Such a line of argument may find support in Grice's requirement that it must be possible to work out the presence of an implicature if it is to count as a conversational implicature. In working out the implicature the hearer is relying on a grasp of the proposition semantically expressed by the utterance in question, or to use Grice own words, 'the

³⁸ This line of response may not be available to DeRose who employs a WAM in defence of his views about the truth conditions of epistemic possibility statements (DeRose 1999, pp. 196-197; 2002, p. 174-175).

conventional meaning of the words used, together with the identity of any references that may be involved' (1991a, p. 31). But if a grasp of a the proposition semantically expressed by an utterance is a condition for being able to work out the implicature, this might be difficult if language users were subject to the kind of confusion that the effectiveness of a WAM would seem to require.

In response to this line of argument it may be argued that Grice allowed for the possibility that an implicature can be intuitively grasped, and only required that it must be possible to supply the argument relevant to the working out of the implicature. Grice states that 'even if it can in fact be intuitively grasped, unless the intuition is replaceable by an argument, the implicature (if any) will not count as a conversational implicature' (1991a, p. 31).³⁹ Thus the objection may be taken to involve a too strong reading of Grice which leaves out the possibility that the implicature may be intuitively grasped.

Furthermore, it seems that this objection, like DeRose's distinction between intuitions about truth and intuitions about falsity (see section 6.2), presupposes a high degree of sensitivity in language users with respect to the distinction between the truth conditions of an utterance and what the utterance pragmatically imparts. But unlike DeRose's attempt to distinguish between plausible and implausible candidates for a legitimate WAM, this line of argument amounts to a more thorough rejection of the idea that we are liable to confuse a variation in warranted assertability conditions for a variation in truth conditions. It therefore seems that our intuitions concerning the Bank Case and the Airport Case only provide evidence for contextualism on the assumption that language users are highly sensitive to the semantic workings of their own language. If language users are not sensitive to the semantic workings of their own language, why should we take them to be able to distinguish the truth conditions of an utterance from what the utterance pragmatically imparts even when the focus is on the latter? We are therefore left with the following conditional: If our intuitions about the Bank Case and Airport Case are to provide evidence for contextualism, then competent language users must be sensitive to the semantic workings of their own language. But as I will argue in section 7.3, the view that competent language users are sensitive to the semantic workings of their own language seems to be an untenable position for contextualists.

³⁹ I am not concerned with the question of what is the correct interpretation of Grice, nor do I want to get into a general debate about how to understand this requirement on conversational implicatures. The point is merely to point to a possible response to the objection in question.

6.7 Appendix: A Positive Argument for Contextualism Based on the Knowledge Account of Assertion

In this section I will look at an attempt by DeRose (2002) to construct a positive argument for contextualism based on a contextual variation in the standards for warranted assertability and the knowledge account of assertion. I will also look at Brown's (2005) response to this argument. The point of this section is to show that even though such an argument would be significant if successful, DeRose's argument is not as compelling as it initially appears to be when it is spelled out in more detail and that contextualists are left with the original problem of ruling out the possibility that they are mistaking a variation in warranted assertability conditions for a variation in truth conditions.

The knowledge account of assertion (henceforth just 'the knowledge account') has been defended and developed by among others Timothy Williamson (1996, 2000, ch. 11), and I will follow DeRose in focusing on Williamson's version of the knowledge account in the following discussion. According to Williamson (2000, p. 243), the constitutive rule of assertion is the knowledge rule.

(The knowledge rule) One must: assert p only if one knows p .⁴⁰

According to the knowledge account the knowledge rule is the unique rule specific to assertion. This means that the knowledge rule is essential to the speech act of assertion and allows us to individuate it from other speech acts.

While a comprehensive discussion of the merits of the knowledge account is beyond the scope of this thesis, I think it is fair to say that the knowledge account has some intuitive plausibility. For instance, it explains why asking the questions 'How do you know?', and more aggressively, 'Do you know?', in response to an assertion constitutes a challenge. These questions constitute a challenge because only knowledge warrants assertion, and the absence of an answer would mean the absence of warrant (Williamson 2000, pp. 252-253). The knowledge account also promises to give an explanation of the version of Moore's paradox involving 'know'. There is something wrong with asserting ' p , and I do not know that p ', even though it could very well be true. According to the knowledge account, to make this assertion warrantably one would have to know that p and that one does not know that p , but

⁴⁰ See Williamson (1996, p. 494; 2000; p. 243).

this is impossible, thereby explaining why there appears to be something wrong with such assertions (Williamson 2000, pp. 253).⁴¹

Since I will be discussing the knowledge account in connection with contextualism, the knowledge rule should be given an appropriate metalinguistic formulation.⁴² Cohen (2004, p. 486) has proposed the following metalinguistic version of the knowledge rule:

(The knowledge rule*) S may assert *p* in C only if ‘S knows *p*’ is true at C.

DeRose argues that the knowledge account, together with a contextual variation in the standards for warranted assertability, provides a strong, positive argument for taking ‘know’ to be context sensitive.

If the standards for when one is in a position to warrantably assert that P are the same as those that constitute a truth condition for “I know that P,” then if the former vary with context so do the latter. In short: The knowledge account of assertion together with the context sensitivity of assertibility yields contextualism about knowledge. (DeRose 2002, p. 187)

The argument takes as one of its premises the connection between knowledge and warranted assertability stated by the knowledge account. Furthermore, the argument is based on the premise that there is a contextual variation in the standards for warranted assertability. But if there is a variation in the standards for warranted assertability and those standards are the same as the truth conditions of utterances of sentences containing ‘know’, then DeRose concludes, there must also be a variation in the truth conditions of those utterances.

If sound, this argument would show a different route from the variation in our use of ‘know’ revealed by our intuitions about cases such as the Bank Case and the Airport Case, to the context sensitivity of ‘know’. The argument for contextualism that have been the focus of the discussion so far is that our intuitions about the cases provide evidence for the context sensitivity of ‘know’ by revealing a variation in the truth conditions of utterances of sentences containing ‘know’. DeRose’s argument takes a more indirect route by arguing via the knowledge account and the variation in the standards for warranted assertability in order to establish the context sensitivity of ‘know’. As DeRose (2002, p. 188) points out, this

⁴¹ Williamson (2000, ch. 11) argues that the knowledge account has other advantages over its rivals as well.

⁴² Hawthorne (2004, pp. 85-91) has argued that there is a tension between the knowledge account and contextualism. While ‘know’ is ascriber dependent (i.e. it depends on the context of utterance), whether someone has warrant to make an assertion is not. He argues that the problem remains even when the knowledge account has been given an appropriate metalinguistic formulation. However, I will not pursue this problem here, and I will assume that contextualists can appeal to the knowledge account.

argument, if sound, would not be vulnerable to the objection that it rests on the mistake of confusing warranted assertability for truth as the intuitions appealed to are intuitions about warranted assertability in the first place, leaving no room for a confusion of this sort.

However, it is not obvious that the argument is sound when it spelled out in more detail. The first problem involves sensitive invariantism. Even granting the truth of the premises, there seems to be a problem with taking this as an argument for contextualism. As Thomas Blackson (2004) and Brown (2005) point out, the argument is not effective as an argument against sensitive invariantism because it does not establish that the variation revealed by the argument is due to the context sensitivity of 'know' rather than a sensitivity to features of the situation of the subject. The argument as it stands is therefore invalid as an argument for contextualism. DeRose (2004b) admits the failure of the argument to establish the truth of contextualism on its own, but argues that third-person cases, such as the Airport Case, still provide evidence against sensitive invariantism.

The upshot of this is that it seems better to view DeRose's argument as an argument against classical invariantism rather than an argument design to establish that 'know' is context sensitive on its own. It would still be significant if the argument could be made to work against classical invariantism, thereby ruling it out as an alternative to contextualism. However, on closer examination of the argument it is not clear that it succeeds in establishing even this conclusion. Brown (2005) has argued that DeRose's also fails as an argument against classical invariantism.

Brown identifies DeRose's argument as having two premises. The first premise is that there is a contextual variation in how well positioned one must with respect to a proposition in order to be warranted in asserting that proposition. Brown (2005, p. 266) refers to the first premise as the 'context sensitivity of assertion', but since I prefer to reserve the label 'context sensitivity' for the phenomenon of semantic context sensitivity, I will instead refer to this premise as 'the variability claim'. However, I continue to use Brown's (2005, p. 266) formulation:

(The Variability Claim) How well positioned one must be with respect to p to assert warrantably that p depends on context.

The variability claim draws support from cases like the Bank Case and the Airport Case. While we could warrantably assert that p in a context like the one in Bank Case A, we could not warrantably assert that p in a context like the one in Bank Case B, despite being just as

epistemically well positioned with respect to p in both contexts. This point is demonstrated by the original cases presented by DeRose and Cohen, but as the case of (17) illustrates, it applies equally well to sentences not containing 'know'.

- (5a) I know that it'll be open.
- (17) The bank will be open Saturday.

In the context in Bank Case B the intuition is that DeRose cannot warrantably assert either (5a) or (17), but in the context in Bank Case A both seem intuitively acceptable.

The second premise of the argument is the knowledge account. The knowledge account as stated above is the view that the knowledge rule is the unique, constitutive rule of assertion. However, the knowledge rule and the variability claim do not yield the desired conclusion for the contextualists. It is consistent with these premises that there are other rules governing assertion, thus leaving it open whether the variation in how well positioned one must be in order to make an assertion could be accounted for by a contextual variation in these rules. To make the argument valid we need to add a further premise in order to rule out the possibility of there being such rules. Brown (2005, p. 271) therefore attributes the following uniqueness claim to DeRose:

(The Uniqueness Claim) The knowledge rule is the only rule governing how well positioned one must be to assert p warrantably.

According to the version of the knowledge account defended by Williamson the knowledge rule is the unique rule specific to assertion, thus seemingly making the uniqueness claim a natural premise if one already accepts the knowledge account. Moreover, DeRose claims to be following Williamson in accepting something like the uniqueness claim, 'I will join Williamson in holding that this [the knowledge rule] is the only rule governing assertion that has to do with asserting only what one is positioned well enough with respect to' (2002, p. 180). It appears that DeRose has something very similar to the uniqueness claim formulated by Brown in mind when making this statement. However, it is not clear that this is what Williamson has in mind (Brown 2005, p. 273). Williamson states that, 'There may be other evidential norms for assertion, if they can be derived from the knowledge rule and considerations not specific to assertion' (Williamson 2000, p. 257). Williamson appears to accept that there may be other rules governing how epistemically well positioned one must be

with respect to a proposition to assert it, as long as these rules are not specific to the speech act of assertion. But this does not rule out the possibility that the variation in how well positioned one must be in order to make an assertion could be the result of a contextual variation in a rule that is not specific to assertion, but is derived from the knowledge rule together with more general considerations not specific to assertion.

The contextualists might therefore want to go further than Williamson and claim that the knowledge rule is the only rule governing how well positioned one must be with respect to a proposition to assert it. However, this raises the question of how to motivate the stronger claim on behalf of the knowledge account. The view that the knowledge rule is the constitutive and individuating rule of assertion only requires the weaker view claim that the knowledge rule is the unique rule specific to assertion. The contextualists therefore need to find further motivation for making the stronger claim.

Furthermore, by making a stronger claim on behalf of the knowledge account, the contextualists run the risk of making the knowledge account less plausible. For instance, Williamson (2000, p. 257) argues that the knowledge account can explain why it appears warranted to assert that p when one reasonably believes that one knows that p , despite p being false and consequently that one does not know that p . Given those circumstances, it is reasonable to believe that one has warrant to make the assertion, thus explaining why the assertion is reasonable without being warranted in the strict sense that it complies with the knowledge rule. It might otherwise be seen as a problem for the knowledge account that an assertion made under those circumstances appears to be warranted. However, this defence of the knowledge account rests on the ability to derive the reasonableness of asserting that p when one believes that one knows that p from the knowledge rule and considerations not specific to assertion. By arguing that the knowledge rule is the only rule governing how well positioned one must be with respect to a proposition to assert it, the contextualists would run the risk of undermining this defence of the knowledge account.

Brown (2005, p. 274) also cites the rule of relevance, which states that one must make one's contribution to the conversation at hand relevant, as an example of a rule not specific to assertion that might have to do with the epistemic position of the subject (see also section 6.3). It might be that the subject matter of a conversation directly deals with the epistemic position of the conversational participants with respect to a certain proposition. In such a context, the rule of relevance would concern the epistemic position of the conversational participants.

In light of these problems, it is not clear why an invariantist should accept the uniqueness claim. But Brown (2005, pp. 274-278) outlines another strategy that contextualists might pursue in attempting to construct a sound argument against invariantism based on the knowledge account. According to this strategy, DeRose's argument can be understood as involving a more restricted sense of 'warrant'. Williamson says, "'Warrant' is used here as a term of art, for the evidential property (if any) which plays the role of property C in the correct simple account of assertion. This use need not correspond exactly to that of 'warrant' in everyday English" (2000, p. 243). Following these remarks by Williamson, it might be argued that DeRose's argument only concerns 'warrant' in the more restricted sense of whether an assertion complies with knowledge rule.

By understanding 'warrant' in a more restricted sense, the uniqueness claim is no longer problematic as long as one accepts the knowledge account in the first place. The uniqueness claim is then understood as the claim that the knowledge rule is the only rule governing how well positioned one must be with respect to p to assert that p warrantably, in the restricted sense of 'warrant' that one is one warrantably in asserting that p if and only if one knows that p . However, in order for the argument to be successful, 'warrant' must be used in the same way throughout the entire argument. This point applies to the variability claim as well. But the cases appealed to by the contextualists were originally seen as providing evidence for the variability claim understood as involving 'warrant' in the ordinary sense, not in the restricted sense required in order to motivate the uniqueness claim. The contextualists are therefore left with the challenge of making sure that our intuitions about the cases track the warrantably assertability conditions as involving 'warrant' in the restricted sense. Brown argues that 'there is no reason to suppose that [the context sensitivity of assertion] concerns one particular dimension of evaluation of assertion' warrantably*, rather than reflecting a number of the different dimensions of evaluation' (2005, p. 276). The challenge facing contextualists with respect to this problem is in many ways similar to the original challenge of making sure that our intuitions about the cases track the truth conditions, rather than the warrantably assertability conditions, of utterances of sentences containing 'know'.

If Brown is right, there are two ways of construing DeRose's argument depending on whether 'warrant' is understood in its ordinary sense or in a more restricted sense. If 'warrant' is understood as being used in the ordinary sense, then the contextualists can appeal to cases such as the Bank Case or the Airport Case in order to find support for the variability claim. However, it is difficult to motivate the strong version of the uniqueness claim if one understands 'warrant' in this way. On the other hand, adopting an understanding of DeRose's

argument according to which it involves 'warrant' as understood in a restricted sense undermines the evidence for the variability claim. In both cases, Brown argues that DeRose fails to provide a successful argument against invariantism, as the contextualists fail to provide sufficient support for either the uniqueness claim or the variability claim. I take it that Brown's line of argumentation succeeds in undermining DeRose's argument, and consequently, that we are left with the original problem of whether our intuitions about the cases reveal a variation in the truth conditions or the warranted assertability conditions of utterances of sentences containing 'know', as there is no available argument from the latter to the former.

7 Contextualism and Semantic Blindness

Contextualism offers a straightforward way of accommodating our intuitions about the Bank Case and the Airport Case. But that does not mean that contextualism does not have any difficulties accounting for other features of our linguistic practices involving ‘know’. This raises questions about what contextualists should say about these features of our linguistic practices, and most significantly for the present purposes, how this relates to the contextualist attempt to argue that our intuitions about the Bank Case and the Airport Case provide evidence for contextualism. The main point of this section is to examine the issue of whether considerations about what Hawthorne (2004) calls ‘semantic blindness’ undermines the positive evidence for contextualism. In section 7.1 I present an argument by Hawthorne (2004) involving ‘know’ as it figures in the that-clauses of belief reports. Hawthorne argues that his argument demonstrates that contextualists are forced to acknowledge that competent language users display semantic blindness with respect to ‘know’. In section 7.2 I look at some further objections to contextualism considered by DeRose (forthcoming) that also aim to demonstrate that the contextualists must appeal to semantic blindness. I then present DeRose’s response to these objections and the claim that contextualists must posit semantic blindness. Having thus presented the objection to contextualism and DeRose’s response, I discuss whether these considerations involving semantic blindness undermine the positive evidence for contextualism in section 7.3. I connect this discussion with the discussion in section 6.5 where I argue that our intuitions concerning the Bank Case and the Airport Case only provide evidence for contextualism on the assumption that language users are highly sensitive to the semantic workings of their own language.

7.1 Propositional Attitude Reports and ‘Know’

Hawthorne (2004) raises an objection to contextualism based on our use of ‘know’ when it figures in the that-clauses of propositional attitude reports, and more specifically, belief reports. He argues that when it comes to such reports, competent language users seem to rely on a disquotational schema for ‘know’.

Disquotational Schema for ‘Knows’ (DSK). If an English speaker E sincerely utters a sentence *s* of the form ‘A knows that *p*’, and the sentence in the that-clause means that *p* and ‘A’ is a name or indexical

that refers to *a*, then E believes of *a* that *a* knows that *p*, and expresses that belief by *s*. (Hawthorne 2004, p. 101)⁴³

This schema allows us to take a sincere utterance of ‘S knows that *p*’ by a speaker A and truthfully report A’s belief by ‘A believes that S knows that *p*’. For instance, taking DeRose’s utterance of in (5a) Bank Case A, and assuming that his utterance is sincere, it appears that (18) is a correct way of reporting DeRose’s belief.

(5a) I know that it’ll be open.

(18) Keith DeRose believes that he knows that the bank will be open.

But combining the disquotational schema for ‘know’ with the True Belief Schema which states that if S believes that *p*, then S’s belief is true if and only if *p*, causes problems for contextualism (Hawthorne 2004, p. 99). Suppose that we were in a situation similar to that of DeRose and his wife in Bank Case B, that is to say that we were in a context where there is a lot at stake and the possibilities of error are salient, and that we are considering DeRose’s situation in Bank Case A. If someone were to utter (19) in our context, and assuming the correctness of the contextualist verdicts with respect to the original Bank Case, the utterance would be false.

(19) Keith DeRose knows that the bank will be open.

However, given the disquotational schema for ‘know’ we can truthfully report DeRose’s belief by (18) as above, and since we are assuming that the belief expressed by DeRose’s utterance of (5a) in Bank Case A is true, the True Belief Schema delivers the conclusion that (19) is true. But since it was already declared that an utterance of (19) would be false as uttered in our context, this conclusion must be wrong.

Hawthorne (2004, p. 103) concludes that contextualists will most likely reject the disquotational schema for ‘know’ when faced with this argument. He points out that we do not adopt a disquotational schema in the case of familiar context sensitive expressions such as ‘here’ or ‘today’. For instance, take a sincere utterance of (20) made several months ago. In this case I cannot use (21) to report the speaker’s belief since ‘today’ would refer to today not

⁴³ DeRose (forthcoming) has suggested that the consequent of the disquotational schema for ‘know’ should be given a metalinguistic formulation. According to DeRose’s version the consequent would state that it is true to say of E that he ‘believes that A knows that *p*’. Though I retain Hawthorne’s original formulation I will follow DeRose on this point.

to the day when the speaker made the original utterance. I could only use (21) to report the speaker's belief if we were speaking on the same day.

(20) It is hot today.

(21) She believes that it is hot today.

However, Hawthorne argues that giving up the disquotational schema for 'know' is not to be taken lightly. He makes the following claim about the role of the disquotational schema for 'know' in our ordinary practice of belief reporting:

By contrast, it looks very much as if we do adopt something like [the Disquotational Schema for 'Knows']. If for example, someone sincerely utters 'I know that I will never have a heart attack', we have no hesitation whatsoever in reporting the contents of his mind by claiming that he believes that he knows that he will never have a heart attack. That is how the verb 'know' seems to work. (Hawthorne 2004, p. 101)

Our ordinary practices involving belief reporting seem to reveal a commitment to the disquotational schema for 'know' whereas it does not in the case of 'today'. But where does this leave us with respect to contextualism? Hawthorne argues that the most plausible approach for contextualists to take would be to invoke semantic blindness. As he puts it, 'There is a real sense in which users of the word 'know' are blind to the semantic workings of their language' (Hawthorne 2004, p. 107). Thus the appeal to semantic blindness amounts to the claim that competent language users are even implicitly unaware of the context sensitivity of 'know' as is revealed by their apparent commitment to the disquotational schema for 'know'.

7.2 Comparative Judgements of Content and Metalinguistic Claims

DeRose (2005; forthcoming) considers Hawthorne's argument as well as two related objections to contextualism.⁴⁴ I will go through these other objections before turning to DeRose's response as he offers a similar response to both set of objections. The first objection DeRose considers is an objection from comparative judgements of content. It is argued that

⁴⁴ There are other, related objections that I will not discuss here. Stephen Schiffer (1996) has for instance argued that the contextualist solution to scepticism requires language users to be confused in a way that seems to involve semantic blindness, though it should be noted that Schiffer does not use the expression 'semantic blindness'. There are also problems concerning 'know' as it figures in the that-clauses of speech reports (see e.g. Cappelen and Lepore 2005b, ch. 7).

when someone in a context like the one in Bank Case B utters (22), this appears to contradict an utterance of (19) made in a context like the one in Bank Case A.

- (19) Keith DeRose knows that the bank will be open.
- (22) Keith DeRose does not know that the bank will be open.

But according to contextualism there is no contradiction here. Since the relevant utterances of (19) and (22) take place in contexts which differ with respect to those features which contextualists deem relevant to the context sensitivity of 'know', both utterances come out true and there is no contradiction. Again it is argued that this reveals that competent language users are blind to the context sensitivity of 'know'. The case of 'here' provides a suitable contrast. An utterance of (23a) would not be taken as contradicting an utterance of (23b) unless they happened to take place in the same location. Thus competent language users do not display a similar semantic blindness to the context sensitivity of 'here'.

- (23a) It is cold here.
- (23b) It is not cold here.

DeRose also considers a related objection from metalinguistic claims. According to this objection, if the speaker uttering (19) and the speaker uttering (22) were informed of the utterance made by the other, they would both claim that the other speaker's utterance was false. Again, this seems to reveal semantic blindness on the part of competent language users when it comes to the context sensitivity of 'know'. If the speakers were aware of the context sensitivity of 'know' they would not claim that the other's claim was false. By contrast, if the speaker of (23a) was informed of an utterance of (23b) taking place in a different location, he would not claim that the utterance was false, or would at least be clearly mistaken in doing so.

The first part of DeRose's response to these objections is to try to mitigate the force of the intuitions behind the objections. He claims that the intuitions are less powerful when the cases are presented correctly. While he admits that the intuition that an utterance of (19) contradicts an utterance of (22) is quite strong when they are presented as taking place in a dispute between the two speakers, he insists that contextualists are not committed to the verdict that both utterances are true. What he advocates is a version of contextualism according to which the contribution of 'know' to the truth conditions of both utterances is the same (DeRose 2004a; forthcoming). He refers to this view as 'single scoreboard semantics':

‘On this view, there is a single scoreboard in a given conversation; the truth-conditional content of both speakers’ uses of “knows” are given by the score registered on this single scoreboard.’ (DeRose 2004a, p. 6). While the conversational score is liable to change as conversation moves forward, DeRose does not embrace a view according to which the conversational score changes drastically between speakers. He is therefore able to claim that since the contribution of ‘know’ to the truth conditions of both utterances follows the same conversational score, both utterances cannot be true.^{45, 46}

Furthermore, DeRose (2005, p. 194; forthcoming) claims that in the cases where contextualists do not take the speakers to be contradicting each other (i.e. cases without disputes), the intuitions that they do contradict each other is significantly weaker. Thus, when considering the speaker uttering (19) as occupying a context like the one in Bank Case A, and the speaker uttering (22) as occupying a context like the one in Bank Case B, with neither disputing the other’s claim, we should expect competent language users to be less inclined to take the speakers as contradicting each other. Similarly with respect to the objection from metalinguistic claims, the intuition that if the speaker uttering (19) and the speaker uttering (22) were informed of the utterance made by the other they would claim that the other speaker’s utterance was false, is much weaker when the speakers are not involved in a dispute (DeRose 2005, pp. 194-195; forthcoming).

DeRose (2005, pp. 196-197; forthcoming) also argues that similar considerations weakens the intuitions behind the disquotational schema for ‘know’. He describes a case in which Louise is at a tavern discussing whether Jim was at work yesterday with her colleagues. Based on the testimony of a reliable informant and seeing his hat in the hall, Louise sincerely utters (24a).

(24a) I know that Jim was at the office yesterday.

Louise utterance is intuitively true. On the other hand, one of her colleagues, Thelma, is being interviewed by the police who are investigating a serious crime and wants to know whether Jim

⁴⁵ DeRose (2004a; forthcoming) favours the view that in cases of disputes like the one between the speaker uttering (19) and the speaker uttering of (22), both utterances lack a truth value. In these cases, the speakers would set the contribution of ‘know’ to the truth conditions of utterances of sentences containing it differently. DeRose (2004a, p. 18) admits that this amounts to a failure of bivalence, but does not think that this is a problem.

⁴⁶ DeRose (forthcoming) also extends this point to cases where the speaker uttering (22) is disputing an earlier utterance of (19) or vice versa. In this case only the latter utterance lacks a truth value. Note that this should make contextualists careful when presenting their cases. If Mary and John’s utterance are understood as disputing Smith’s earlier utterance of (6a) in the Airport Case by uttering (6b), their utterance could be regarded as having no truth value according to the view advocated by DeRose.

was at the office yesterday. Though Thelma has access to the same evidence as Louise, there is a lot more at stake in her context and she therefore proceeds to utter (24b).

(24b) I don't know that Jim was at the office yesterday.

Again, the utterance seems to be intuitively true. But what if Thelma was asked by the police if Louise could provide more information regarding Jim's whereabouts? Could Thelma, who is aware of Louise having sincerely uttered (24a), use (25) to correctly report Louise's belief?

(25) Louise believes that she knows that Jim was at the office yesterday.

DeRose claims that this belief report appears to be wrong. The intuition that this would be a correct belief report is at least a lot weaker in this case. He further argues that this demonstrates that Hawthorne's claims about the apparent commitment of competent language users to the disquotational schema for 'know' are based on a failure to look at the right sort of cases.

Based on these considerations, DeRose (forthcoming) argues that even if language users display semantic blindness with respect to 'know' this does not pose a problem for contextualism. Even if 'know' turns out not to be context sensitive, those language users whose intuitions match the contextualist verdicts with respect to the problematic cases would be subject to semantic blindness. Since the problematic cases for contextualism do not seem to yield clear intuitions one way or another, at least when presented correctly, there will be some language users whose intuitions turn out to be incorrect. Thus we have to accept that there is some semantic blindness whether contextualism is the correct view or not.

7.3 Semantic Blindness and the Positive Evidence for Contextualism

However, accepting semantic blindness on the part of competent language users may be more problematic for contextualists than DeRose is willing to admit. The problem is that semantic blindness threatens to undermine the positive evidence for contextualism. MacFarlane makes the following point:

If ordinary speakers have a faulty grasp of the *meaning* of "know", then we cannot confidently appeal to variability in the standards they require someone to meet in order to count as "knowing" as support for a

theory about the meaning of “know”. Yet these data are the primary evidence in favor of contextualism. (2005b, p. 215)

If language users are blind to context sensitivity of ‘know’ with respect to problematic cases such as belief reports and comparative judgements of content, then why should we take their intuitions to be a reliable when it comes to cases such as the Bank Case and the Airport Case. Even if there is semantic blindness regardless of whether ‘know’ is context sensitive or not, there is still a problem for contextualism because contextualists are relying on the intuitions of competent language users to provide evidence that ‘know’ is context sensitive.

As it stands, it may be replied that MacFarlane’s objection ignores a crucial difference between cases such as the Bank Case and the Airport Case and the problematic cases involving belief reports and comparative judgements of content. While contextualists claim that our intuitions about the former cases are relatively clear and provides uniform support for contextualism, DeRose (forthcoming) has argued that our intuitions about the latter cases are far less clear and do not uniformly undermine contextualism (see section 7.2). Therefore, the contextualists may take the intuitions of competent language users about the Bank Case and the Airport Case to provide evidence for contextualism, despite the fact that the same language users display semantic blindness with respect to the problematic cases.

However, even granting this point, there may still be a problem for contextualism in taking competent language users to be blind to the context sensitivity of ‘know’. In section 6.5 it was argued that our intuitions about the Bank Case and the Airport Case only provides evidence for contextualism on the assumption that language users are sensitive to the semantic workings of their own language. But this assumption is difficult to square with competent language users displaying semantic blindness. If competent language users are sensitive to the semantic workings of their own language, it is difficult to see how they could be affected by semantic blindness which amounts to being unaware of the semantic workings of their own language, and more specifically, the context sensitivity of ‘know’. It therefore appears that the position that the language users are sensitive to the semantic workings of their language is untenable for contextualists admitting that language users display semantic blindness. But in that case there is no reason to suppose that our intuitions about the Bank Case and the Airport Case provide evidence for contextualism.

Perhaps a case could be made that language users only display semantic blindness with respect to certain features of the language. It could thus be argued that language users are sensitive to the features of the language that are important in providing evidence for contextualism while at same time displaying semantic blindness with respect to other features

of the language. This line of response may have some merit with respect to Hawthorne's objection involving 'know' as it figures in the that-clauses of belief reports as it could be argued that language users display semantic blindness with respect to such embedded uses of 'know'. However, this response does not help with the objection from comparative judgements of content and the objection from metalinguistic claims. These objections are based on how language users think about 'know' as it is used in other contexts and comparing the contents of different utterances of sentences containing 'know'. But contextualist cannot argue that competent language users are liable to display semantic blindness when thinking about what is going on in contexts other than the one that they occupy. As was illustrated by Cappelen and Lepore's arguments discussed in section 5, contextualists are relying on intuitions about what is going on in other contexts when using Impoverished Context Shifting Arguments to provide evidence for the context sensitivity of 'know'. Therefore, if contextualists were to argue that such intuitions are liable to be the result of semantic blindness, they would undermine the evidence for contextualism. In conclusion, the problem remains. If language users display semantic blindness, it becomes difficult to argue that language users are sensitive to the semantic workings of their own language. But this sensitivity seems to be required if our intuitions about the Bank Case and the Airport Case is to provide evidence for contextualism.

Again, it may be that the best bet for the contextualists is to look for companions in guilt, and again it is gradable adjectives such as 'tall' that provide a suitable candidate. In section 6.5 I argued that a WAM is available in defence of the view that 'tall' is not a context sensitive expression. Furthermore, our practices seem to betray as much commitment to the disquotational schema for 'tall' as it does for 'know' (DeRose forthcoming). If someone sincerely utters (16a) it seems that I can use (26) to report her belief.

(16a) Tim is tall.

(26) She believes that Tim is tall.

This seems to present a dilemma if one thinks that the argument against contextualism that I have presented in this section is convincing, but also thinks, as most do, that gradable adjectives such as 'tall' are context sensitive. This is because there seems to be an analogous argument that can be used to undermine the evidence for the context sensitivity of 'tall'. However, this in itself does not solve the problem for contextualism as it does not show where the argument goes wrong, but only proposes to present a dilemma. It also raises questions

about whether the case of ‘tall’ and ‘know’ are as analogous as contextualists claim.⁴⁷ I do not propose to offer any conclusions as to the right way out of this dilemma as my concern is with the debate about the evidence for the context sensitivity of ‘know’, not with gradable adjectives. I am therefore content with having presented an argument that purports to undermine the case for contextualism based on our intuitions about the Bank Case and the Airport Case.

⁴⁷ Hawthorne (2004, p. 104) is aware of the worry that his argument could be used against gradable adjectives as well, and worries that it shows that his argument proves too much. His response is to insist that there are more clarificatory devices available in the case of gradable adjectives than there is in the case of ‘know’ and that this makes for a relevant difference. I do not wish to enter into a debate about Hawthorne’s response, except to note that his claim that we have fewer clarificatory devices in the case of ‘know’ has been challenged (DeRose forthcoming). See also Ludlow (2005) for suggestions about clarificatory devices in the case of ‘know’.

8 Summary and Conclusions

I began this discussion with a presentation of contextualism and two cases that contextualists take to provide evidence for the thesis that ‘know’ is a context sensitive expression: DeRose’s Bank Case and Cohen’s Airport Case. These cases form the central part of the contextualists attempt to provide evidence for the context sensitivity of ‘know’ that I have been discussing here. Initially these cases may seem to provide strong evidence for the context sensitivity of ‘know’. Whereas an ascription of knowledge is intuitively true in one context, a denial of knowledge with respect to the same true proposition and the same subject, or a similarly epistemically positioned subject, is intuitively true in another context. Contextualists can give a straightforward account of this variation in our intuitions by appealing to the context sensitivity of ‘know’. Different utterances of the same sentence containing ‘know’ may have different truth conditions, and a sentence containing ‘know’ may therefore be true as uttered in one context, but false as uttered in another.

However, as I take it that the discussion demonstrates, matters are more complicated than they might initially appear to be. In section 4 I looked at data involving our use of ‘know’ that goes beyond the cases that contextualists typically appeal to. In particular, I discussed Stanley’s (2005) Ignorant High Stakes cases and how these cases can be made to fit with a plausible version of contextualism. This discussion raises questions about how contextualism should be developed, and the constraints placed on contextualism by general assumptions about the nature of context sensitive expressions. More specifically it raises the question of how the contribution of ‘know’ to the truth conditions of utterances of sentences that contain it is determined. I also discussed the worry that the Bank Case and the Airport Case are not psychologically realistic because they ignore the effects of the higher stakes on the confidence of the participants and that this, rather than the context sensitivity of ‘know’, may be the source of our intuitions. This connects with the discussion about the Ignorant High Stakes cases, since these cases would provide contextualists with a way of removing this worry. While I do not feel obliged to offer any definitive conclusions with respect to these matters, I think that they illustrate the need for contextualists to be clearer on how contextualism is to be developed and implemented if these questions are to be resolved. These questions become even more pressing given the discussion in section 6, since contextualists cannot take our intuitions about these cases to be the result of semantic blindness, but need to

find another way of accommodating them that does not make language users ignorant of the semantic workings of their language.

I further pursued the connection between contextualism and more general assumptions about semantics and context sensitivity in section 5. In this section I relied heavily on Cappelen and Lepore's (2003; 2005b) arguments concerning context shifting arguments to bring out some of the assumptions behind the contextualist argumentation based on our intuitions about the Bank Case and the Airport Case. In relying on cases such as the Bank Case and the Airport Case, contextualists assume that our intuitions about cases like these are semantically significant. This has further consequences for questions about context sensitivity in general, such as the question of how large the set of context sensitive expressions is.

What I take it that discussion in section 4 and section 5 reveal, is that providing evidence for contextualism is not simply a matter of finding a pair of cases that reveal the correct sort of variation in our intuitions about utterances of sentences containing 'know'. These matters are connected with more general questions about semantics and context sensitivity, and contextualists need to make potentially controversial assumptions about these matters in order to provide evidence for the context sensitivity of 'know'. The contextualist attempts to argue our intuitions reveal a variation in the truth conditions of utterances of sentences containing 'know' may therefore be challenged by calling these assumptions into question. However, rather than pursuing these general questions about context sensitivity and semantics, the rest of my discussion have focused on a more specific alternative to treating our intuitions as revealing a variation in the truth conditions of utterances of sentences containing 'know'.

The main part of the thesis involves the question of whether our intuitions about the Bank Case and the Airport Case reveal a variation in the truth conditions of utterances of sentences containing 'know', or whether they can be given a post-semantic, pragmatic explanation in terms of a variation in warranted assertability conditions. DeRose defends contextualism on the grounds that such an explanation could be given with respect to any attempt to provide evidence for (or against) a semantic theory. He proposes a set of criteria that will restrict the use of such explanations and rule out the possibility that our intuitions about the Bank Case and the Airport Case can be thus explained. However, I argued that the attempts to provide a plausible set of criteria that rules such an explanation, fails. Based on this I further argued that in order for our intuitions about the Bank Case and the Airport Case to provide evidence for contextualism, contextualists must assume that competent language users are be sensitive to the semantic workings of their own language.

In section 7 I connected these issues with issues concerning semantic blindness. In the light of certain objections to contextualism it seems that contextualists have to admit that competent language users display semantic blindness with respect to 'know'. However, I argued that contextualists cannot hold that competent language users display semantic blindness as that would undercut the assumption that language users are sensitive to the semantic workings of their own language, and thus undermine the positive evidence for contextualism.

Where does this leave us with respect to the evidence for contextualism? If successful, the argument demonstrates that our intuitions about the Bank Case and the Airport Case need not be taken as providing evidence for contextualism. However, it should be noted that some caution is called for in drawing sweeping conclusions at this point. It was noted that an analogous argument could be used to undermine the evidence for the context sensitivity of 'tall'. This may or may not be problematic depending on one's position with respect to the context sensitivity of gradable adjectives, and may raise the suspicion that the argument proves too much.

Furthermore, it is important to stress that the argument does not purport to show that contextualism is false. Even if successful, the argument only undermines the positive evidence for contextualism. Nor does the argument rule out the possibility that there are other ways of providing evidence for contextualism that does not involve the assumption that language users are sensitive to the semantic workings of their own language.

Having said that, I think that the argument reveals a serious problem for contextualists seeking to provide evidence for the context sensitivity of 'know'. On the one hand contextualists want to claim that competent language users are sensitive to the semantic workings of their own language, but on the other hand they appeal to semantic blindness when faced with intuitions that are problematic for contextualism. But according to the argument I have presented, there is a tension here that threatens to undermine the evidence for contextualism, and in the end it is far from clear that we must treat our intuitions about the Bank Case and the Airport Case as providing evidence for the context sensitivity of 'know'.

9 Bibliography

- Bach, K. 2005, 'The Emperor's New 'Knows'', in Preyer, G. and Peter, G. (eds.), *Contextualism in Philosophy*, Oxford University Press, Oxford, pp. 51-89.
- Blackson, T. A. 2004, 'An Invalid Argument for Contextualism', *Philosophy and Phenomenological Research*, 68, pp. 344-345.
- Brown, J. 2005, 'Adapt or Die: the Death of Invariantism?', *The Philosophical Quarterly*, 55, pp. 263-285.
- 2006 'Contextualism and Warranted Assertibility Manoeuvres', *Philosophical Studies*, 130, pp. 407-435.
- Cappelen, H. and Lepore, E. 2003, 'Context Shifting Arguments', *Philosophical Perspectives*, 17, pp. 25-50.
- ——— 2005a, 'A Tell Tale: In Defense of Semantic Minimalism and Speech Act Pluralism', in Preyer, G. and Peter, G. (eds.), *Contextualism in Philosophy*, Oxford University Press, Oxford, pp. 197-219.
- ——— 2005b, *Insensitive Semantics*, Blackwell Publishing, Oxford.
- Cohen, S. 1988, 'How to be a Fallibilist', *Philosophical Perspectives*, 2, pp. 91-123.
- 1998, 'Contextualist Solutions to Epistemological Problems', *Australasian Journal of Philosophy*, 76, pp. 289-306.
- 1999, 'Contextualism, Skepticism and the Structure of Reasons', *Philosophical Perspectives*, 13, pp. 57-89.
- 2004, 'Knowledge, Assertion, and Practical Reasoning', *Philosophical Issues*, 14, pp. 482-491
- DeRose, K. 1992, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research*, 52, pp. 913-929.
- 1995, 'Solving the Skeptical Problem', *The Philosophical Review*, 104, pp. 1-52.
- 1999, 'Contextualism: An Explanation and Defense', in Greco, J. and Sosa, E. (eds.), *The Blackwell Guide to Epistemology*, Blackwell Publishing, Oxford, pp. 187-205.
- 2002, 'Assertion, Knowledge, and Context', *The Philosophical Review*, 111, pp. 167-203.
- 2004a, 'Single Scoreboard Semantics', *Philosophical Studies*, 119, pp. 1-21.
- 2004b, 'The Problem with Subject-Sensitive Invariantism', *Philosophy and*

- Phenomenological Research*, 68, pp. 346-350.
- 2005, 'The Ordinary Language Basis for Contextualism, and the New Invariantism', *The Philosophical Quarterly*, 55, pp. 172-198.
- forthcoming, 'Bamboozled by Our Own Words', *Philosophy and Phenomenological Research*.
- Donnellan, K. 1966, 'Reference and Definite Descriptions', *The Philosophical Review*, 75, pp. 281-304.
- Egan, A., Hawthorne, J. and Weatherson, B. 2005, 'Epistemic Modals in Context' in Preyer, G. and Peter, G. (eds.), *Contextualism in Philosophy*. Oxford University Press, Oxford, pp. 131-168.
- Grice, P. 1991a, 'Logic and Conversation', in *Studies in the Way of Words*, Harvard University Press, Cambridge, Mass., pp. 22-40.
- 1991b, 'Further Notes on Logic and Conversation', in *Studies in the Way of Words*, Harvard University Press Cambridge, Mass., pp. 41-57.
- Hawthorne, J. 2004, *Knowledge and Lotteries*, Oxford University Press, Oxford.
- Kaplan, D. 1978, 'Dthat', in Cole, P. (ed.), *Syntax and Semantics*, vol. 9, Academic Press, New York, pp. 221-253, in Martinich, A. (ed.) 2001, *The Philosophy of Language*, 4th Edition, Oxford University Press, Oxford, pp. 325-338.
- 1989a, 'Demonstratives', in Almog, J., Perry, J. and Wettstein, H. (eds.), *Themes from Kaplan*, Oxford University Press, Oxford, pp. 481-563.
- 1989b, 'Afterthoughts', in Almog, J., Perry, J. and Wettstein, H. (eds.), *Themes from Kaplan*, Oxford University Press, Oxford, pp. 565-614.
- Kay, P. 1997, 'Linguistic Competence and Folk Theories of Language: Two English Hedges', in *Words and the Grammar of Context*, CSLI Publications, Stanford, pp. 133-143.
- Lakoff, G. 1973, 'Hedges: A Study in Meaning Criteria and the Logic of Fuzzy Concepts', *Journal of Philosophical Logic*, 2, pp. 458-508.
- Levinson, S. 1983, *Pragmatics*, Cambridge University Press, Cambridge.
- Lewis, D. 1979, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic*, 8, pp. 339-359.
- 1996, 'Elusive Knowledge', *Australasian Journal of Philosophy*, 74, pp. 549-567.
- Ludlow, P. 2005, 'Contextualism and the New Linguistic Turn in Epistemology', in Preyer, G. and Peter, G. (eds.), *Contextualism in Philosophy*, Oxford University Press, Oxford, pp. 11-50.
- MacFarlane, J. 2003, 'Future Contingents and Relative Truth', *The Philosophical Quarterly*,

- 53, pp. 321-336.
- 2005a, ‘Making Sense of Relative Truth’, *Proceedings of the Aristotelian Society*, 105, pp. 321-339.
- 2005b, ‘The Assessment Sensitivity of Knowledge Attributions’, in Gendler, T. S. and Hawthorne, J. (eds.), *Oxford Studies in Epistemology*, vol. 1., Oxford University Press, Oxford, pp. 197-233.
- McGinn, C. 1981, ‘The Mechanisms of Reference’, *Synthese*, 49, pp. 157-186.
- Perry, J. 2001, *Reference and Reflexivity*, CSLI Publications, Stanford.
- Rysiew, P. 2001, ‘The Context Sensitivity of Knowledge Attributions’, *Noûs*, 35, pp. 477-514.
- Schiffer, S. 1996, ‘Contextualist Solutions to Scepticism’, *Proceedings of the Aristotelian Society*, 96, pp. 317-333.
- Salmon, N. 1991, ‘The Pragmatic Fallacy’, *Philosophical Studies*, 63, pp. 83-97.
- Stanley, J. 2005, *Knowledge and Practical Interest*, Oxford University Press, Oxford.
- Williamson, T. 1996, ‘Knowing and Asserting’, *The Philosophical Review*, 105, pp. 489-523.
- 2000, *Knowledge and its Limits*, Oxford University Press, Oxford.
- 2005, ‘Contextualism, Subject-Sensitive Invariantism and Knowledge of Knowledge’, *The Philosophical Quarterly*, 55, pp. 213-235.