

Relativism, Sceptical Paradox, and Semantic Blindness*

Dirk Kindermann

ABSTRACT. Relativism about knowledge attributions is the view that a single occurrence of ‘S knows [does not know] that p’ may be true as assessed in one context and false as assessed in another context. It has been argued that relativism is equipped to accommodate all the data from speakers’ use of ‘know’ without recourse to an error theory. This is supposed to be relativism’s main advantage over contextualist and invariantist views. This paper argues that relativism does require the attribution of semantic blindness to speakers, viz. to account for sceptical paradoxes and epistemic closure puzzles. To that end, the notion of semantic blindness is clarified by distinguishing between content-blindness and index-blindness, and it is argued that the attribution of index-blindness required by the relativist account is implausible. Along the way, it is shown that error-theoretic objections from speakers’ inter-contextual judgments fail against relativism.

KEYWORDS. Knowledge Attributions, Relativism, Error Theory, Semantic Blindness, Sceptical Paradox, Epistemic Closure Puzzles

I Introduction

The case for relativism about knowledge attributions, as made by John MacFarlane (2005a, 2011b), rests on the following line of thought. Epistemic contextualism – the view that the content of sentences of the form ‘S knows that p’, and ‘S does not know that p’ may vary with the epistemic standards salient in the context of use – is supported by appeal to linguistic data from ordinary speakers’ use of knowledge sentences: the same speakers tend to accept ‘S knows that p’ as true when uttered in one context but not when uttered in another context. But contextualism makes incorrect predictions about, for instance, speakers’ inter-contextual truth ascriptions and their retraction of knowledge claims. The best explanation of this data available to contextualists involves the attribution of semantic blindness to speakers. Similarly, classical and subject-sensitive invariantists face troubling data, whose explanation commits them to some error theory or other. In contrast, MacFarlane claims, relativism can account for all of the data. Relativism is the only view that avoids the ‘double-edged sword’ of attributing systematic error to speakers (MacFarlane, 2005a, 215). That is its main virtue.

The relativist has a *prima facie* strong empirical case. *Pace* MacFarlane, however, I argue in this paper that relativism cannot explain all of the data from speakers’ use of knowledge sentences without the attribution of semantic blindness to speakers.

*Penultimate draft, forthcoming in *Philosophical Studies*

In section 2, I introduce the relativist's case against contextualism and in favour of relativism.¹ In section 3, I present a recent error-theoretic objection against relativism that appeals to speakers' inter-contextual judgments of relativized truth ascriptions to knowledge claims. In section 4, I show why the strategy behind this objection fails. However, there is a simpler semantic blindness objection to relativism that is immune to the given replies, as I argue in section 5. Speakers tend to be puzzled by sceptical paradoxes, and relativist explanations of this phenomenon are bound to appeal to speakers' semantic blindness. I identify two different kinds of semantic blindness involved in the debate: content-blindness and index-blindness. Finally, I show that the objection generalizes to ordinary cases of epistemic closure from a variety of data on which relativists have rested their case.

2 Contextualism and the case for relativism

On contextualist semantics, the content expressed by an occurrence of 'S knows that p', and 'S does not know that p', depends in part on the epistemic standards salient in the conversational context, where this dependence can be traced to the occurrence of 'know' (Cohen (1987); DeRose (1995); Lewis (1996); cf. Schaffer (2004)). When John uses the sentence 'Bill knows that he has hands' in an everyday context with low epistemic standards (LOW) he expresses, very roughly, the content *Bill knows relative to low standards that he has hands*. When Mary is in a context with high epistemic standards (HIGH), for instance in an epistemology class on scepticism, she uses the sentence to express the content *Bill knows relative to high standards that he has hands*.² The sentence is true as used by John in LOW, but false as used by Mary in HIGH. The truth value of knowledge sentences can vary across contexts of use even when the facts about the knowing subject's situation do not change. Contextualists motivate the variability of content by appeal to ordinary speakers' use of 'know'. While speakers accept many knowledge attributions as true in mundane contexts of use, they tend to give in to, e.g., sceptical considerations that raise the epistemic standards and reject these attributions as false in such contexts.

Trouble for contextualism comes, among other things, from speakers' inter-contextual truth ascriptions to and retraction of knowledge claims. When Mary is in HIGH, say in a conversation about brains-in-vats, she will judge knowledge attributions expressed by the sentence 'Bill knows that he has hands' false, even when the attribution is made by John in LOW. Likewise, speakers in LOW will judge knowledge denials, e.g. 'Bill does not know that he has hands', false even when they are made in

¹In this paper, I will not be concerned with MacFarlane's arguments against traditional and subject-sensitive versions of invariantism that complete his case for relativism. The presentation of his argument against contextualism is intended to exemplify his argument-by-elimination strategy and to introduce the semantic framework.

²This characterisation is simplified in two respects. First, epistemic standards – the strength of epistemic position one must be in to count as knowing – come in degrees. There are more different standards and corresponding contexts than HIGH and LOW. Second, the propositional form *S knows relative to high/low standards that p* is supposed to be neutral between different contextualist proposals for the structure of the propositions expressed.

HIGH. But contextualism predicts that in judging knowledge attributions and denials uttered in other contexts, speakers are sensitive to the epistemic standards at play in the context of use.

What is more, speakers will retract earlier knowledge claims when their context changes in relevant respects. Suppose John sincerely asserts ‘I know that my car is parked in the driveway.’ Mary points out that car thieves could be roaming John’s neighbourhood and that his car might have been stolen. She thereby raises the standards, and it is natural for John to retract his earlier assertion by saying ‘I guess I was wrong. I did not know that my car was parked in the driveway.’ But according to contextualism, John is mistaken in his retraction. After all, his earlier assertion is true in the context of its use. Contextualism predicts that it would be appropriate for John to reply: ‘I did not say that I *know* that my car is parked in the driveway. I only meant that I know by low epistemic standards that my car is parked in the driveway. And that is still true.’ This, however, is not a natural reply for John.³

Contextualists’ best response to the recalcitrant data is to adopt an error theory: Speakers are systematically mistaken in their inter-contextual truth ascriptions and retraction of knowledge claims. These mistakes are explained by speakers’ *semantic blindness*: ‘users of the word ‘know’ are blind to the semantic workings of their language.’ (Hawthorne (2004, 107); cf. DeRose (2006, 321)) The kind of semantic blindness contextualists need to ascribe is what I propose to call content-blindness:

CONTENT-BLINDNESS

Speakers are blind to the fact that particular sentences (sentences of the form ‘S knows that p’/‘S does not know that p’) can express different contents in different contexts (bar indexical expressions in the substitution instances of ‘S’ or ‘p’).

Contextualists take the content of knowledge sentences to depend on the context of use, so the kind of semantic blindness they must attribute can be further specified as use-content-blindness.

Let me settle some terminology. An expression is *indexical* iff its content at a context depends on features of the context (after disambiguation). A sentence is (semantically) *context-sensitive* iff either its content or its truth value (or both) at a context depends on features of the context.⁴ Thus, contextualism about knowledge is the view that sentences of the form ‘S knows that p’ are indexical – specifically, that the content of ‘S knows that p’ depends on epistemologically significant features of the context. It is one among other views that hold that sentences of the form ‘S knows that p’ are

³See MacFarlane (2005a) and Williamson (2005). Further objections to epistemic contextualist semantics concern, e.g., disagreement judgments (MacFarlane, 2007), belief reports (Cappelen & Lepore, 2005; Hawthorne, 2004), the analogy between ‘know’ and gradable adjectives (Stanley, 2005), the analogy with indexicals and with quantificational determiners like ‘all’ and ‘every’ (Schaffer & Szabó, forthcoming).

⁴For the above notion of indexicality and a different notion of context-sensitivity, see MacFarlane (2009).

context-sensitive.⁵ Content-blindness is ignorance of an expression's indexicality.^{6 7}

According to MacFarlane, the attribution of semantic blindness undermines the contextualist's cause. Contextualists appeal to the variability of speakers' intra-contextual truth ascriptions in support of their view. But the more data they explain away by appeal to speakers' ignorance of the semantic workings of 'know', the weaker is their appeal to speakers' usage of 'know' as evidence in favour of their semantics.⁸ For MacFarlane, the point generalises:

[A] *general* problem with positing speaker error to explain away facts about use is that such explanations tend to undermine the evidential basis for the semantic theories they are intended to support. All of these semantic theories are justified indirectly on the basis of facts about speakers' use of sentences, and the more error we attribute to speakers, the less we can conclude from these facts. (MacFarlane, 2005a, 215)

Relativists piggyback on the case against contextualism. On MacFarlane's relativist semantics, the truth value of knowledge sentences can vary from one context of assessment to another: 'S knows that p' may be true at (context of use) C_U and (context of assessment) C_{A1} and false at C_U and C_{A2} .⁹ Underlying the relativist semantics is a roughly Kaplanian picture (Kaplan, 1989). Sentences in contexts of use express contents. Contents are evaluated at an index (Lewis, 1980), what Kaplan called a circumstance of evaluation, to yield truth values.¹⁰ Kaplan took the index (circumstances of evaluation) to include at least a world and a time coordinate, determined by the context of use. Relativism can be located in this picture by the following three theses:

I. INVARIANT CONTENT

'S knows that p' expresses a content that is invariant across contexts of use (bar indexical expressions in the substitution instances of 'S' and 'p').

⁵These views include nonindexical contextualism (Brogaard, 2008; Kompa, 2002; MacFarlane, 2009) and relativism (MacFarlane, 2005a, 2011b).

⁶On the above notion of context-sensitivity, contingent sentences and tensed sentences count as context-sensitive. This broad notion of context-sensitivity is useful as an umbrella term for specific kinds of context-sensitivity, among which is the sensitivity to the epistemic standards salient at a context.

⁷Akerman & Greenough (2010) propose an alternative way of distinguishing between kinds of context-sensitivity and corresponding kinds of blindness. They provide a helpful discussion and comparison of these kinds of blindness for the case of vague expressions.

⁸For a defence of contextualism from semantic blindness objections see Blome-Tillmann (2008), Cohen (2004), DeRose (2006), and Schaffer & Szabó (forthcoming).

⁹In this paper, I will focus on MacFarlane's version of relativism. Similar relativist semantics for knowledge attributions have been proposed by Richard (2004, 2008) and Kölbel (2009). All of the arguments in favour and against relativism discussed in this paper apply, *mutatis mutandis*, to all of these versions.

¹⁰I will here ignore the differences between Lewisian index and Kaplanian circumstance of evaluation. They have no bearing on any of the arguments discussed in this paper. I will henceforth use the shorter 'index'.

2. INDEX SENSITIVITY

The index contains an epistemic standards coordinate, to which the truth value of ‘S knows that p’ is sensitive. More generally, call a sentence *index-sensitive* iff its truth value depends on (some coordinate in) the index.

3. ASSESSMENT SENSITIVITY

The truth value of ‘S knows that p’ depends on the epistemic standards salient at the context of assessment, which may be different from the context in which the sentence is uttered. Put together with 2, the epistemic standards-coordinate in the index is determined by the context in which a knowledge claim is assessed.

Relativism about knowledge attributions can thus be characterised as the view that while the semantic contribution of ‘knows’ to a sentence ‘S knows that p’ does not make the sentence indexical, the sentence is assessment-index-sensitive: Its truth value depends on the epistemic standards coordinate in the index, which is determined by the context of assessment.¹¹

Relativism yields an elegant account of the data that troubles contextualism. An assessor in HIGH is correct in ascribing falsity to ‘Bill knows that he has hands’ as uttered in LOW, since it is the standards at play in the assessor’s context that matter to her truth value ascription. Moreover, relativism predicts that John will retract his earlier assertion ‘I know that my car is parked in the driveway’ when the context is shifted to HIGH by Mary in conversation. When John in HIGH assesses his past assertion, it is appropriate for him to use the standards salient in his present context to judge his assertion false and to correct himself. Finally, relativism also explains intra-contextual truth ascriptions. When I judge knowledge sentences uttered in my own context as true (false), context of assessment and context of use are identical. Relativism’s predictions in these cases coincide with those of contextualism.

MacFarlane completes the case for relativism by making similar points against traditional invariantism and subject-sensitive invariantism. These views run into trouble with data, e.g., from intra-contextual truth ascriptions and temporal and modal embeddings of ‘know’, respectively. And again, relativism makes correct empirical predictions where invariantist views must wield error theories to their defence. MacFarlane concludes that since relativism is the only view that respects all of the data without requiring the attribution of systematic speaker error, it is superior to all other views.

3 An error-theoretic objection against relativism

Contextualism and all forms of invariantism each face trouble from cases in which their view does not predict the assessor’s egocentric sensitivity to epistemic standards

¹¹Contextualism can also be situated in this picture of semantics: It locates the relevant epistemic standards in the context of use that combines with a sentence to deliver a content; contents then are evaluated at an index, which does not include a standards coordinate.

(or stakes) in their inter-contextual judgments of knowledge attributions. One error-theoretic strategy against relativism, exemplified by Montminy (2009), presents cases for which relativism supposedly cannot predict this egocentric focus either. As we will see in section 4, this particular strategy is not promising.

According to the first variant of the strategy that Montminy pursues, relativism is committed to the attribution of systematic speaker error when speakers make judgments about explicitly relativized truth claims, made in other contexts, concerning knowledge attributions. Montminy has us consider, first, the following dialogue in Low:

- (1) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true.

Suppose Mary is in HIGH and is presented with (1). It is natural for her to judge false not only John's knowledge attribution but also Bob's assessment of John's claim. This is what relativism predicts. Mary is right in taking the standards salient in her context as relevant for the assessment of John's knowledge attribution as well as Bob's truth ascription to John's claim. Mary thinks that John and Bob do not know that Neil Armstrong was the first man to set foot on the moon, and that it is false that they know this, so Bob's claim that it is true that they know is false in Mary's context.

Trouble for relativism is supposed to arise when speakers explicitly relativize their truth ascriptions. Consider a similar dialogue in Low:

- (2) John: We both know that Neil Armstrong was the first man to set foot on the moon.
 Bob: That's true relative to this context.

According to Montminy, Mary in HIGH will judge John's knowledge attribution false and she will also judge Bob's truth ascription false, *regardless of the relativization*. This runs against relativism's prediction. On relativist semantics, the explicitly relativized truth ascription made by Bob in Low is true at every context of assessment, hence also at Mary's.

Relativists hold that there are (at least) two different truth predicates. The first is the ordinary English monadic predicate 'true'. Its application to propositions yields sentences that are themselves assessment-sensitive.¹² That is, sentences of the form 'It is true that p' or 'The proposition that p is true' have truth values that vary with the context of use and context of assessment. Importantly, monadic 'true' is disquotational in the following sense: Whenever, in a context of assessment, we correctly judge that S knows that p, we can also judge correctly that it is true that S knows that p, and vice versa. The equivalence schema *It is true that p iff p* is true at any C_U and C_A .

¹²Egan et al. (2005) take the bearers of monadic truth to be utterances. I shall here stick with MacFarlane (2007, 2011a) and take ordinary 'true' as applying to propositions.

Since ‘true’ is assessment-sensitive, the left-hand side will be true (false) at a C_U and C_A just in case the right-hand side is true (false) at that C_U and C_A .

In dialogue (1), Bob is ascribing ordinary monadic ‘true’ to the proposition expressed by John’s knowledge attribution. In the context of assessment of the dialogue, Low, this truth ascription is correct; ‘That’s true’ is true at $C_{U(J\&B)}$ and $C_{A(Low)}$. However, ‘That’s true’ is false at $C_{U(J\&B)}$ and $C_{A(High)}$. Relativists predict correctly that it is appropriate for Mary in her context of assessment to reject Bob’s assertion as false.

The second truth predicate is relational. ‘True relative to context C’ relates truth bearers to contexts in which they are evaluated.¹³ It is assessment-insensitive. That is, the truth value of ascriptions of this truth predicate to sentences or propositions, as in ‘It is true relative to this context that p’, does not vary with contexts of assessment. Thus, the proposition expressed by the following assertion is true relative to any context of assessment: ‘It is true relative to Low that John and Bob know that Neil Armstrong was the first man to set foot on the moon.’ Since this is plausibly what Bob is saying in (2), relativism predicts that Mary in HIGH will judge the proposition expressed by Bob’s utterance as true.

However, speakers in HIGH such as Mary seem to judge explicitly relativized truth ascriptions to knowledge attributions made in Low as false. So on the relativist picture, they are systematically mistaken in these judgments. The conclusion of this objection is that relativism, too, is committed to an error theory; the best argument for relativism fails.¹⁴

4 Relativist replies

There are at least two replies available to the relativist in response to this error-theoretic objection.

(1) *Discarding the data.* Relativists may remind us that while ‘true’ is a natural language English truth predicate that speakers frequently and competently use, ‘true relative to context C’, or ‘true at C_U and C_A ’, is a technical, metalinguistic truth predicate used to formulate relativist semantics *for* English sentences. To ordinary

¹³Egan et al. (2005) treat the binary predicate ‘true relative to context C’ as applying to utterances. MacFarlane prefers the metalinguistic predicate ‘true at C_U and C_A ’ used in the formulation of the relativist semantics, which relates either sentences or propositions with contexts of use and contexts of assessment. Whatever the exact details, the common denominator here is the explicit relativization of truth to a context in which the truth bearer is evaluated. I will use ‘true relative to context C’ as a predicate that applies to truth bearers, be they utterances, sentences(-in-context), or propositions.

¹⁴Montminy (2009) does not claim that this first objection goes through against the relativist, for roughly the reasons I give at the beginning of the next section. He does hold, however, that another variant of the strategy is successful, one on which the claims involving the explicitly relativized truth predicate are replaced by supposedly extensionally equivalent counterfactual conditionals. I argue in the next section that that objection does not succeed against the relativist either.

speakers, Bob's assertion of 'That's true relative to this context' will sound stilted and odd, and contrary to Montminy's Mary, they are most likely to have no clear truth value judgments.¹⁵ But even if speakers had clear judgments, we could not just assume that 'true relative to context C' in their mouths expressed the relativist's technical concept of truth. In fact, trying to accommodate speakers' intuitions about this expression would be bad methodology. By analogy, we would be ill-advised to take into account speakers' reactions to 'That's a context-sensitive sentence' in constructing a semantics for English. As a result, speakers' reactions to expressions involving 'true relative to this context' do not constitute data against relativism.

Relativists could further argue that there is no ordinary, non-technical English expression that correctly captures the theoretical concept of relative truth. This would forestall all attempts to run the error-theoretic objection with everyday expressions that purport to stand a better chance of expressing relative truth in the mouths of English speakers. Relativists could give this reply without compromising their claim that ordinary speakers can make sense of relative truth as a concept of *truth*. According to MacFarlane, we understand the meaning of 'true at C_U and C_A ', or 'true relative to context C', if we grasp 'the role [this predicate] plays in a broader theory of language use: specifically, an account of the speech act of assertion' (MacFarlane, 2005b, 329). In brief, an assertion is a commitment to the truth of what is asserted. This commitment is honoured by providing adequate grounds for the truth of what is asserted relative to the context of assessment in which the assertion is challenged. The commitment also requires that one withdraws the assertion in any future C_A in which what is asserted is shown to be untrue relative to C_A .¹⁶ Retraction data, the relativist can conclude, gives evidence that speakers take themselves to be bound by such a relativized assertoric commitment. Thus, even if speakers lack ordinary expressions to express the technical concept of relative truth, they have an implicit grasp of

¹⁵ Results from Google searches should not serve as conclusive evidence, but the fact that a Google search for 'true relative to this context' (in July 2011) resulted in only four hits – three papers in philosophy (including Montminy (2009)) and one in linguistics – strongly suggests that English speakers are not well-acquainted with the expression. 'True relative to context' offered 29 hits, all of which were philosophy or linguistics papers. A feeling of oddity was also the reaction of most philosophers and non-philosophers when I presented them with dialogue (2).

¹⁶ See for instance MacFarlane (2005a, 329) and (2005b, 336-7) for further details of the commitment undertaken by assertions.

it.¹⁷

But perhaps the relativist's blank dismissal of any sort of data involving the use of vocabulary expressing relative truth is too quick. Montminy agrees that the relativist's reply successfully rebuts the objection based on an explicitly relativized truth predicate in cases like (2). But he maintains that there are perfectly ordinary English locutions that do express relativized truth claims. The following counterfactual conditionals are

¹⁷Montminy argues that even this connection between the notion of truth at C_U and C_A and language use requires the attribution of systematic error to speakers. As he points out, the commitment to withdrawing an assertion in any future context of assessment in which what is asserted is shown to be untrue relative to that context of assessment entails (*):

- (*) In asserting 'S knows that P' ('S does not know that Q'), one commits oneself to withdrawing the assertion in any future context $C_{A(High)}$ ($C_{A(Low)}$) in which what is asserted by 'S knows that P' ('S does not know that Q') is shown to be untrue relative to $C_{A(High)}$ ($C_{A(Low)}$).

Montminy grants for the sake of argument that speakers do in fact withdraw their assertions of 'S knows that P' ('S does not know that Q') made in Low (HIGH) when challenged in HIGH (Low). But he denies that speakers 'take themselves to be bound by [(*)]': 'the fact that a speaker in Low would withdraw her previous knowledge denial made in HIGH does not entail that *when she is in HIGH*, the speaker takes herself to be committed to withdrawing her current knowledge denial, if this denial is challenged in some future Low. As a matter of fact, a speaker in HIGH would reject this commitment, that is, such a speaker would hold that it would be *incorrect* to withdraw her current knowledge denial in some future low-standards context' (Montminy, 2009, 354). Montminy concludes that relativism implies that speakers are systematically mistaken about their commitments to withdraw knowledge claims.

Relativists can resist this argument in several ways. First, it is not clear that speakers would in fact reject the commitment to withdraw knowledge attributions (denials) once the standards have risen (fallen). What could the evidence be for speakers' rejection of the commitment if it is not their linguistic behaviour (which, Montminy grants, honours this commitment)? Presumably Montminy has in mind speakers' explicit judgments about what they take their commitments to be (what they 'hold'). Pending empirical evidence, relativists may simply doubt that speakers would make judgments that stand in contrast to their actual behaviour – after all, they do seem to withdraw. (It has been widely noted that raising standards is easier than lowering them, which might in part explain why speakers are more reluctant to explicitly withdraw knowledge denials (because they do not accept switching to Low). But this does not threaten MacFarlane's claim about the commitment to withdraw, which only says that when speakers *are* in Low, they will withdraw a knowledge denial that is untrue relative to Low.)

Second, relativists may even grant that speakers do make judgments to the effect that it is incorrect to withdraw a knowledge attribution in HIGH, but deny that speakers' explicit judgments are relevant for an account that links the concept of relative truth to speakers' language use. By analogy, syntacticians would have to predict widespread error if part of their evidence was speakers' acceptance or rejection of explicitly stated grammatical rules which speakers, as a matter of fact, employ in their use of language.

Finally, it is worth pointing out that MacFarlane's 2005 commitment account of assertion is by far not the only way of linking the concept of relative truth to language use. MacFarlane (2010) argues that many of the common accounts of assertion corroborate the concept of relative truth as the link between the compositional semantics and language use as long as a corresponding norm of retraction is added. And Egan et al. (2005), Egan (2007), and Stephenson (2007) offer alternative accounts of assertion for relativists.

supposed to express in ordinary English the relativist's thesis that the proposition that John and Bob know that Neil Armstrong was the first man to set foot on the moon is true relative to John and Bob's context Low:

- (3) If it were to be assessed in John and Bob's context, the proposition that they know that Neil Armstrong was the first man to set foot on the moon would be true.
- (4) If I were in John and Bob's context, then they would know that Neil Armstrong was the first man to set foot on the moon.
- (5) If the stakes were low and no error possibilities had been mentioned, then it would be true that John and Bob know that Neil Armstrong was the first man to set foot on the moon.

Montminy claims that ordinary speakers like Mary in HIGH would understand (3) – (5) and reject them as false. But the relativist must hold that (3) – (5) are true, since they express the relativist's thesis that the attribution of knowledge to John and Bob is true relative to John and Bob's context. The relativist, Montminy concludes, is after all committed to the claim that Mary's judgments are systematically mistaken (Montminy, 2009, 350-2).

I think relativists have good reasons to reject both that (3) – (5) express relativized truth claims and that Mary's falsity judgments concerning these conditionals (if she has any clear judgments) run against relativism's predictions. On mainstream semantics of counterfactuals, the antecedent has an effect on the *worlds* at which the consequent is evaluated for truth and falsity. For instance, on the Lewis-Stalnaker account, a conditional $A > B$ is true iff B is true at the closest world(s) at which A is true. But it is an open question whether, when combined with relativism, the antecedent also affects the epistemic standards coordinate in the index. Relativism has Montminy's problem only if the antecedent does affect the epistemic standards coordinate. However, there seem to be no independent theoretical reasons why this must be the case. And the alleged data from Mary's falsity judgments give us pre-theoretical reasons to think that the antecedent has no effect on the epistemic standards coordinate with respect to which the consequent is evaluated. So relativists can make the right predictions by simply adopting the unamended Lewis-Stalnaker analysis of conditionals: When Mary assesses the knowledge claim in the consequent at the closest world(s) in which she is in John's context, or in which the stakes are low and no error possibilities have been mentioned, she will still judge this knowledge claim as false relative to her actual context of assessment HIGH. As a result, (3) – (5) are not extensionally equivalent to, and do not express, relativized truth ascriptions, and Lewis-Stalnaker-cum-relativist semantics of (3) – (5) make the intuitively right predictions concerning Mary's (putative) rejection.

(2) *Override*. It is worth noting that relativists can avoid the attribution of speaker error even if they take the data from (2) at face value and accept that 'true relative to

context C , as used by ordinary speakers, is the relativist's metalinguistic truth predicate. On this assumption, what the data shows is that speakers take their own current contexts of assessment to override all other contexts of assessment. A context C_{A1} *overrides* a context C_{A2} when the assessment in C_{A1} of the facts dictates their proper assessment in C_{A2} . Let us say that C_{A1} *upward overrides* C_{A2} when a first-order truth (falsity) ascription 'It is true (false) relative to C_{A1} that p ' dictates the second-order ascription of truth (falsity) 'It is true (false) relative to C_{A1} that it is true relative to C_{A2} that p ', for any C_{A2} . C_{A1} *downward overrides* C_{A2} when a second-order truth (falsity) ascription 'It is true (false) relative to C_{A1} that it is true relative to C_{A2} that p ' dictates the first-order ascription of truth (falsity) 'It is true (false) relative to C_{A1} that p ', for any C_{A2} .¹⁸

In Mary's second-order falsity ascription to Bob's first-order truth ascription, her current context of assessment upward overrides Bob's context of assessment. Since Mary judges John's knowledge attribution false from her context of assessment, she also judges that Bob's truth ascription (in his context of assessment) to John's knowledge attribution is false. The epistemic standards in her context of assessment upward override those in Bob's context of assessment.

Upward and downward overriding allow relativists to make the right predictions for (2). It would seem that overriding is also compatible with other data relativists have been keen on.¹⁹ However, it commits them to second-order relativism.²⁰

5 Sceptical paradox and index-blindness

The error-theoretic strategy from speakers' inter-contextual judgments in sections 3 and 4 appealed to data involving the use of 'true relative to context C ' and other locutions allegedly expressing relativized truth claims. The replies available to the relativist correspondingly targeted the use of these expressions to avoid the charge of systematic error attribution. However, there is a simpler error-theoretic objection that makes no appeal to data involving relativized truth claims and is thus immune

¹⁸Moruzzi & Wright promote a similar idea of trumping for relativism about future contingents. Overriding contrasts with trumping in one crucial respect: 'Being trumped [...] involves that another perspective gets, not to *override* the mandates of one's own perspective, but to determine what they are.' (Moruzzi & Wright, 2009, 314 n.10)

¹⁹Overriding does not affect the predictions about intra-contextual and inter-contextual truth ascriptions. It also leaves the data from ordinary disquotational 'true' untouched, since it exclusively targets the explicitly relativized, technical truth predicate. Finally, retracting an earlier knowledge claim because it is false at the current context of assessment is compatible with the additional judgment that the knowledge claim also turns out to be false at the previous context of assessment.

²⁰Given first-order relativism and overriding, it follows that there is a first-order relativized truth ascription 'It is true (false) relative to C_{A1} ', which is true relative to a context of assessment C_{A2} and false relative to C_{A3} . A further question is whether first- and second-order relativism can coherently be combined with third-order absolutism, or whether higher-order relativism is required all the way up. It is beyond the scope of this paper to determine whether higher-order relativism of any sort would be a price worth paying for accommodating data from inter-contextual truth ascriptions.

to the above replies. As I argue in this section, the solution to sceptical paradoxes and epistemic closure puzzles that naturally falls out of the relativist semantics requires the implausible attribution of a kind of semantic blindness I call index-blindness. Although relativists have not explicitly addressed sceptical paradoxes, they cannot reject this solution without significant costs, as I show in section 7.

Consider the following sceptical argument:

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

As Schiffer (1996, 317) remarks, 'this argument presents a paradox because it tempts us to say three things that are mutually inconsistent: its first premise is true; its second premise is true; and its conclusion is false.' That is, the following three sentences are mutually inconsistent: 'I don't know that I'm not a BIV'; 'If I don't know that I'm not a BIV, then I don't know that I have hands'; and 'I know that I have hands.' Yet each of these sentences strikes us as intuitively true.

A 'fully satisfactory' solution, Schiffer continues, must accomplish two things: First, it must explain why (SA) in fact does not present a paradox. That is, it must show which one of the three sentences is false, and explain why the argument is in fact valid. And second, it must explain why (SA) *seemed* to present a paradox. That is, it must explain why the false sentence seemed true, and why we were tempted to believe that the premises are true but the conclusion false.

The relativist semantics has the resources to provide a fully satisfactory solution to sceptical paradox. Going for the first task, relativists may remind us that we assess knowledge attributions and denials from a context of assessment. Sceptical hypotheses, as introduced by the first premise of (SA), raise the epistemic standards in the context of assessment to a high level (HIGH), since one must be in quite a strong epistemic position to rule out that a sceptical hypothesis like the BIV scenario obtains. Thus, we assess 'I don't know that I'm not a BIV' as true at C_U and $C_{A(High)}$. Importantly, the epistemic standards are now raised for a context of assessment that encompasses the entire argument (SA). Hence, we must also assess the second premise and the conclusion of (SA) at C_U and $C_{A(High)}$. This allows the relativist to explain why (SA) in fact is valid, thus capturing the intuitive force of the sceptic's argument. To understand how, we need to have a look at the relativist's notion of validity.

The validity of an argument is commonly understood as necessary preservation of truth from premises to conclusion. Since truth, on a relativist semantics, is truth at a context of use C_U and context of assessment C_A , validity is necessary preservation of truth at C_U and C_A from premises to conclusion.²¹ Thus, (SA) is valid if, necessarily, whenever the premises are true at C_U and $C_{A(High)}$, its conclusion is also true at

²¹Cf. MacFarlane's definition of logical consequence (MacFarlane, 2011a, 167)

C_U and $C_{A(High)}$.²² Since both premises and its conclusion are true at C_U and $C_{A(High)}$, (SA) is valid and sound.

This gives relativists a neat explanation of the validity and soundness of (SA). It also allows them to show that (SA) does not present a paradox. A paradox is ‘a set of mutually inconsistent propositions each of which enjoys some plausibility when considered on its own’ (Schiffer, 1996, 324). Propositions, and sentences, are mutually inconsistent just in case they cannot be true together. Obviously, relativists will relativize truth here as well. Thus, a paradox is a set of sentences, or propositions, each of which enjoys some plausibility when considered on its own but which cannot be true together at any C_U and C_A . But there are no contexts C_U and C_A at which each of the following three sentences enjoys some plausibility on its own yet they cannot be true together: ‘I don’t know that I’m not a BIV’; ‘If I don’t know that I’m not a BIV, then I don’t know that I have hands’; ‘I know that I have hands.’ From any $C_{A(High)}$, the first two are true and plausible, but the third lacks plausibility – it is false at $C_{A(High)}$. From any $C_{A(Low)}$, the second and third are true and plausible, but the first is not.²³ Provided that the first premise of (SA) induces $C_{A(High)}$ for the entire argument, and so for the three sentences of the apparent paradox, ‘I know that I have hands’ turns out to be the apparently true, but in fact false, sentence – it is false at $C_{A(High)}$. This completes the first part of the relativist solution.

To deliver on the second part, relativists must explain why ‘I know that I have hands’ seemed true and why it appeared that while the premises of (SA) are true, the conclusion is false. Why our resistance to accept the sceptical conclusion ‘I don’t know that I have hands’? The only answer for relativists is this: We intuitively assess the conclusion from our everyday context in which low epistemic standards prevail, while we assess the sceptical hypothesis introduced by the premises from a context with extraordinarily high epistemic standards. In (SA), we reason from the premises’ truth at C_U and $C_{A(High)}$ to the conclusion’s falsity at C_U and $C_{A(Low)}$. Yet we are ignorant of our switching contexts of assessment in moving from premises to conclusion, and its effect on our truth and falsity judgments. In the context of assessment induced by the premises of the argument (HIGH), the conclusion is in fact true. But we believe it to be false because we unknowingly fall back to assessing it from Low. Likewise, the premises together with the unnegated conclusion strike us as paradoxical because we are unaware of the fact that we assess them from different contexts of assessment, and that this influences our judgments.²⁴

²²Note that to get necessary and sufficient conditions for the validity of (SA), it must be the case that, necessarily, for every C_U and C_A , if the premises are true at C_U and C_A , then the conclusion is true at C_U and C_A . It is easy to see that for any C_A with sufficiently high standards to make the premises of (SA) true, the conclusion will also be true.

²³For simplicity, I assume that all contexts $C_{A(High)}$ and all contexts $C_{A(Low)}$ together exhaust all possible contexts of assessment.

²⁴I have chosen a phenomenology of sceptical arguments according to which, as we reason from premises to conclusion, we go from corresponding truth judgments to a falsity judgment. Accordingly, the relativist solution diagnoses a switch of contexts of assessment in the move from premises to conclusion. Nothing hangs on this choice. One might prefer to say that in reasoning through (SA), we come to accept the sceptical conclusion and only then remember our usual acceptance of the unnegated conclusion in everyday contexts. These two judgments then strike

The second part of the relativist solution²⁵ involves the attribution of error to ordinary speakers – indeed to anyone who is tempted to think (SA) presents a paradox. We mistakenly think (SA) gives rise to paradox because we mistakenly switch contexts of assessment midway. Relativists must explain this error by appeal to a specific kind of semantic blindness. For relativists, knowledge attributions are assessment-index-sensitive. Ignorance of this sensitivity is what leads speaker to their mistaken judgments. They betray index-blindness:

INDEX-BLINDNESS

Speakers are blind to the fact that the truth value of contents expressed by sentences involving a particular expression ('know') can vary with a particular coordinate in the index (epistemic standards).

More precisely, we can say that relativists are committed to the thesis that speakers are epistemically assessment-index-blind. They are ignorant of the fact that the truth value of knowledge sentences can vary with the epistemic standards coordinate in the index, which is determined by the context of assessment. In assessing (SA), they are blind to the fact that the truth value of knowledge attributions is sensitive to the epistemic standards coordinate in the index, which is determined by their context of assessment. Speakers do not realise that when they assess the premises as true and the conclusion as false, they do so from different contexts, and that this affects their assessment.²⁶

The attribution of index-blindness to speakers spoils the relativist's self-proclaimed status as a no-blindness-theory and undermines the main advantage of relativism. But one need not share MacFarlane's general worry about blindness attributions to find index-blindness an unpalatable consequence of relativism. A more serious problem

is as inconsistent, creating the air of paradox. If we prefer this diagnosis, what needs explanation in terms of index-blindness is the mistaken feeling of inconsistency between one's truth judgment of the conclusion in the high-standards context of assessment of (SA) and one's falsity judgments of 'I don't know that I have hands' in everyday low-standards context of assessment.

²⁵Note that the relativist 'solution' to sceptical paradox, like the contextualist's, does not amount to a refutation of scepticism. On the contrary, the relativist account gives sceptical intuitions their due. This does not imply that we can never (in no context of assessment) truly attribute knowledge. Whether or not sceptical standards are reasonable ones to adopt is a question on which the relativist semantics is neutral.

²⁶Relativists need a *local* index-blindness thesis. That is, speakers are imputed with index-blindness only locally, namely regarding sceptical arguments, situations in which they are confronted with the three knowledge sentences from (SA), and epistemic closure puzzles, as I argue below. Local index-blindness must be distinguished from *global* index-blindness, according to which speakers are blind to the index-sensitivity of an expression (knowledge sentences) in any situation in which the expression is used or evaluated. The latter is not needed to explain sceptical paradox, and its attribution would undermine other data in favour of a relativist semantics. Note that the distinction between local and global blindness is orthogonal to the distinction between content-blindness and index-blindness. The extent to which a semantic theory needs a local content-blindness thesis is determined by the range of data that this thesis needs to explain.

with index-blindness is that it is particularly implausible in light of speakers' reliable competence in handling other kinds of index-sensitivity. Take as an example the implementation of temporalism according to which the contents expressed by sentences in contexts of use vary their truth value with a time coordinate in the index.²⁷ Consider the following little argument:

- (T) John is asleep at home.
If John is asleep at home, then he is not in his office.
-
- John is not in his office.

Clearly, the argument strikes us as valid. And we would easily detect a change in the value of the time coordinate in the index. Suppose Bob wants to convince Mary at midnight that John never makes it to his office before 11am. Bob gets Mary to agree to premise one and two. If he were then to say, 'You see, John is not in his office!', Mary would respond that John is not in his office at the time of their conversation (midnight), but that this did not show that he is not in his office mornings before 11am. Mary is acutely aware of the fact that she agreed to the premises' truth at midnight, and she is happy to concede the conclusion's truth at midnight. But she will immediately spot the oddness of Bob's switching the time coordinate in reasoning from the premises' being true at midnight to the conclusion's being true in the morning. Temporal (use-)index-sensitivity does not pose a challenge to ordinary speakers. But if more uncontroversial phenomena of index-sensitivity do not cause speaker error, why should we assume that sentences involving 'know' are index-sensitive despite speakers' blindness to this index-sensitivity?²⁸

6 Context confusion without index-blindness?

Relativists might reply that their solution does not require attribution of any form of semantic blindness. Instead, the appearance of paradox can be explained by speakers' confusion of contexts of assessment alone: Speakers are not index-blind, they are merely confused about the epistemic standards relevant to their truth value judgments. Given the nature of sceptical arguments, this confusion is neither surprising nor is its attribution a theoretical cost for relativists.

The exclusive appeal to confusion of contexts is tempting, but it falls short of explaining the paradoxical appearance of the sceptical argument. Let us be clear about

²⁷A similar case could be constructed for 'orthodox' semantics, on which sentences vary in truth value only with a worlds coordinate in the index.

²⁸The objection in this section applies to nonindexical contextualist solutions of the paradox as well. Nonindexical contextualism (Brogaard, 2008; Kompa, 2002; MacFarlane, 2009) shares with relativism the theses of invariant content and index-sensitivity (cf. theses 1 and 2 in section 2) but assumes that the epistemic standards coordinate in the index is determined by the context of *use*. But the kind of use-index-blindness required for nonindexical contextualist solutions is by no means more plausible than the one relativists are committed to.

the proposal. Speakers are supposed to be confused about the contexts of assessment they are occupying, yet to be fully competent with respect to the index-sensitivity of knowledge sentences. Call such speakers context-confused yet index-competent. Due to their context-confusion speakers do not notice the switch of contexts, which explains why they arrive at the puzzling combination of truth value judgments.²⁹

Now remember what the paradox consists in: We are inclined to accept the premises as true and reject the conclusion as false (affirm the unnegated conclusion), but we are also inclined to think that the conclusion's truth follows from the truth of the two premises (we are inclined to think that the two premises and the unnegated conclusion are inconsistent). The problem with the idea of index-competent but context-confused speakers is that it cannot explain both speakers' robust truth value judgments and their sense of inconsistency. Yet both are needed to explain why (SA) appears to present a paradox.

To see this, notice, on the one hand, that the appearance of paradox only arises from (SA) if speakers have robust truth value judgments in the first place. Speakers judge the premises true and the conclusion false. So index-competent speakers, whose truth value judgments are sensitive to the epistemic standards of their context of assessment, cannot be at a loss about which epistemic standards to use in evaluating (SA)'s premises and conclusion for truth and falsity. If they were, they would not arrive at robust judgments at all. Thus, speakers' truth judgments of the premises are competently informed by $C_{A(High)}$ and their falsity judgments of the conclusion by $C_{A(Low)}$. They switch contexts. On the other hand, to explain why index-competent speakers take the truth of the premises and the unnegated conclusion to be inconsistent, we must hold that they are judging from the same context of assessment, i.e. they do not switch contexts. For otherwise index-competent speakers would not get a sense of inconsistency. After all, the premises' truth at $C_{A(High)}$ is only inconsistent with the unnegated conclusion's truth at $C_{A(High)}$. It is not inconsistent with the unnegated conclusion's truth at $C_{A(Low)}$. So we must say that index-competent yet context-confused speakers take themselves to be at $C_{A(High)}$ and $C_{A(Low)}$ at the same time, or that they unknowingly oscillate between the two contexts. But then we lose an explanation of their unchanging, robust truth value judgments: index-competent speakers only have stable truth value judgments if they competently pick up on their context's epistemic standards, i.e. if they are not oscillating between or confused about their context.

In sum, we only get a full explanation of the appearance of paradox if we drop the assumption of index-competence and assume a certain degree of index-blindness.

²⁹For illustration, consider a *plausible* case of a context-confused but index-competent speaker. Suppose that on a rainy St. Andrew's Day, John falls into a coma, from which he awakens three weeks later on a sunny day without any memory of having been in a coma, nor any awareness of time having passed since St. Andrew's Day. Suppose John then incorrectly judges 'It's raining' to be true. Given a temporalist semantics, his judging so can be explained by the fact that John is mistaken about the time of his context of use. Unaware of time's passing, he assumes it is still St. Andrew's Day, when it *was* raining. The explanation of John's mistake in judging need not declare him blind to the time-parameter of the index (given a temporalist semantics). All John suffers from is mislocation in time. He is ignorant of his actual context, yet he is fully index-competent.

Speakers switch contexts of assessment and thus arrive at their truth value judgments, and they take these to be inconsistent because they are blind to the fact these judgments are informed by their contexts of assessment.

7 Epistemic closure puzzles

As a matter of fact, relativists have not addressed sceptical paradoxes in any detail.³⁰ So they might try to dodge the objection from sceptical paradox by rejecting any relativist solution. As concerns sceptical paradoxes, they might say, relativism has no advantage over traditional invariantist theories. Whatever works as a solution of the paradoxes for the latter will do for the former.

This rejection is not promising. Remember that relativists appeal to retraction data in support of their view. When John is presented with the possibility that thieves might have stolen his car, he will retract his earlier claim ‘I know that my car is parked in the driveway.’ Relativists explain this by pointing out that while in the earlier context *LOW* it was right to judge the knowledge claim as true, the standards have been raised to *HIGH*, so judged from John’s later context of assessment, the earlier claim is false. Now, consider this anti-sceptical argument corresponding to (SA):

(AS) I know that I have hands.
 If I know that I have hands, then I know that I’m not a BIV.

I know that I’m not a BIV.

The relativist story in line with the solution to sceptical paradox will say that, given the mention of a sceptical hypothesis, we will switch contexts of assessment from *LOW* to *HIGH* in moving from premises to conclusion. We will thus find the premises true at $C_{A(Low)}$ but the conclusion false at $C_{A(High)}$. But notice the similarity between the explanations of retraction data and of this anti-sceptical paradox. The crucial point is the switch of contexts of assessment. If relativists make so much of their explanation of retraction data, on what grounds could they refuse the similar explanation in the case of (anti-)sceptical paradox?

It cannot be the peculiarity of external-world scepticism that supports the relativist’s potential rejection of the relativist solution. The error-theoretic objection is not limited to sceptical paradoxes concerning BIVs, evil demons, and the Matrix. It generalizes to paradoxes from everyday conversation data about people, things, and activities – data that relativists are keen on explaining. Consider the following argument based on Vogel’s (1990) car theft case – data that MacFarlane readily appeals to

³⁰Perhaps with one exception: Mark Richard seems to advertise relativism as a plausible development, or fix of, epistemic contextualism and recommends the ‘insight into what is going on in skeptical arguments’ offered by contextualism as one of the view’s attractions (2008, 167). Presumably, then, Richard would not be tempted to reject the relativist treatment of sceptical paradoxes in section 5 above.

in support of relativism:³¹

- (C) I know that my car is parked in the driveway.
If I know that my car is parked in the driveway, then I know that it has not been stolen.

I know that my car has not been stolen.

Such ‘semi-sceptical’ cases (to borrow Vogel’s phrase) are equally puzzling, and similar formulations of them have been widely discussed in the literature on epistemic closure.³² But the relativist solution to external-world sceptical paradox is an equally good explanation of why we find these cases puzzling. So relativists would have to reject the relativist solution across the board, which undermines their contention to explain data concerning the use of ‘know’ in everyday situations. But everyday talk about cars, lotteries, and zebras is the very data with which relativists have made their case.

8 Conclusion

My purpose in this paper has been to show that relativism about knowledge sentences requires a semantic blindness thesis to explain the data from speakers’ use of knowledge sentences. I argued that this cannot be shown by appeal to speakers’ alleged use of a relativized truth predicate, or of any natural language stand-in for this predicate. Relativists have good reasons to reject the data, and even if they were to accept it, they could amend their semantics to accommodate the data. I suggested a simpler way to show the relativist’s need for an error theory. An adequate relativist solution to sceptical paradoxes and epistemic closure puzzles systematically commits relativists to the attribution of a particular kind of semantic blindness to speakers: epistemic assessment-index-blindness, i.e. blindness to the fact that the truth value of contents expressed by sentences involving the expression ‘know’ can vary with the epistemic standards coordinate in the index, which is determined by the context of assessment. This solution falls naturally out of the relativist semantics, and as I argued, its rejection

³¹Other cases include lottery-style cases and Dretske’s zebra case (see Dretske (1970); Hawthorne (2004); Vogel (1990)).

³²A simple version of single-premise epistemic closure principle is (EC) (cf. Hawthorne (2004, 31–50) for a discussion and refinement of the principle):

- (EC) Necessarily, if S knows p and S knows that p entails q , then S knows q .

Note that we can formulate the sceptical argument by explicitly using a contraposition instance of (EC) and the plausible assumption that one knows that one’s having hands entails that one is not a BIV: (1) I don’t know that I’m not a BIV. (2) Necessarily, if I don’t know that I’m not a BIV, then either I don’t know that I have hands or I don’t know that my having hands entails that I’m not a BIV. (3) I know that my having hands entails that I’m not a BIV. (4) Therefore, I don’t know that I have hands.

would undermine the relativist's explanation of a wide variety of data from speakers' use of knowledge sentences. But index-blindness is highly implausible in light of speakers' robust competence with other kinds of index-sensitivity. Being committed to index-blindness, relativists lose what MacFarlane advertised as their main advantage over invariantist and contextualist competitors: the ability to predict the empirical data while avoiding the 'double-edged sword' of positing speaker error.³³

References

- Akerman, J. & Greenough, P. (2010). Vagueness and non-indexical contextualism. In S. Sawyer (Ed.), *New Waves in Philosophy of Language*. New York: Palgrave Macmillan.
- Blome-Tillmann, M. (2008). The indexicality of 'knowledge'. *Philosophical Studies*, 138(1), 29–53.
- Brogaard, B. (2008). In defence of a perspectival semantics for 'know'. *Australasian Journal of Philosophy*, 86(3), 439–59.
- Cappelen, H. & Lepore, E. (2005). *Insensitive Semantics*. Malden, MA: Blackwell.
- Cohen, S. (1987). Knowledge, context, and social standards. *Synthese*, 73(1), 3–26.
- Cohen, S. (2004). Contextualism and unhappy-face solutions: Reply to Schiffer. *Philosophical Studies*, 119(1), 185–97.
- DeRose, K. (1995). Solving the skeptical problem. *The Philosophical Review*, 104(1), 1–52.
- DeRose, K. (2006). "Bamboozled by our own words": Semantic blindness and some arguments against contextualism. *Philosophy and Phenomenological Research*, 73(2), 316–38.
- Dretske, F. I. (1970). Epistemic operators. *The Journal of Philosophy*, 67(24), 1007–1023.

³³Earlier versions of this paper were presented in Edinburgh, Miami, St Andrews, at the Australasian Association of Philosophy Conference 2010, and at the conference 'Contexts, Perspectives and Relative Truth' in Bonn. I am greatly indebted to the audiences at all these occasions. Thanks to Derek Ball, Jochen Briesen, Jessica Brown, Herman Cappelen, Stewart Cohen, Torfinn Huvenes, Jonathan Ichikawa, Martin Montminy, Dilip Ninan, Jonathan Schaffer, Daniele Sgaravatti, Brian Weatherson, and all the members of Arché's Contextualism & Relativism Project for their helpful comments and discussion. I am particularly grateful to Patrick Greenough and Jonathan Schaffer for their continuing support and their invaluable feedback on various drafts. In preparing this work, I have benefitted from support by the AHRC, the University of St Andrews, and the DAAD.

- Egan, A. (2007). Epistemic modals, relativism, and assertion. *Philosophical Studies*, 133, 1–22.
- Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth* (pp. 131–67). Oxford: Clarendon Press.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Kaplan, D. (1989). Demonstratives. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563). New York, Oxford: Oxford University Press.
- Kölbel, M. (2009). The evidence for relativism. *Synthese*, 166(2), 375–95.
- Kompa, N. (2002). The context sensitivity of knowledge ascriptions. *Grazer Philosophische Studien*, 64, 79–96.
- Lewis, D. (1980). Index, context, and content. In S. Kanger & S. Öhman (Eds.), *Philosophy and Grammar*. Dordrecht: Reidel.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74(4), 549–67.
- MacFarlane, J. (2005a). The assessment-sensitivity of knowledge attributions. In T. S. Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology 1* (pp. 197–233). Oxford: Oxford University Press.
- MacFarlane, J. (2005b). Making sense of relative truth. *Proceedings of the Aristotelian Society*, 105, 321–39.
- MacFarlane, J. (2007). Truth and subjectivity. Unpublished ms.
- MacFarlane, J. (2009). Nonindexical contextualism. *Synthese*, 166, 231–50.
- MacFarlane, J. (2010). Assessment sensitivity. Relative truth and its applications. Unpublished Book Ms.
- MacFarlane, J. (2011a). Epistemic modals are assessment-sensitive. In B. Weatherson & A. Egan (Eds.), *Epistemic Modality* (pp. 144–78). Oxford: Oxford University Press.
- MacFarlane, J. (2011b). Relativism and knowledge attributions. In S. Bernecker & D. Pritchard (Eds.), *Routledge Companion to Epistemology* (pp. 536–44). London: Routledge.
- Montminy, M. (2009). Contextualism, relativism and ordinary speakers' judgments. *Philosophical Studies*, 143(3), 341–56.
- Moruzzi, S. & Wright, C. (2009). Trumping assessments and the Aristotelian future. *Synthese*, 166(2), 309–31.

- Richard, M. (2004). Contextualism and relativism. *Philosophical Studies*, 119(1), 215–42.
- Richard, M. (2008). *When Truth Gives Out*. Oxford: Oxford University Press.
- Schaffer, J. (2004). From contextualism to contrastivism. *Philosophical Studies*, 119(1), 73–103.
- Schaffer, J. & Szabó, Z. G. (forthcoming). Epistemic comparativism: A contextualist semantics for knowledge ascriptions. *Philosophical Studies*.
- Schiffer, S. (1996). Contextualist solutions to scepticism. *Proceedings of the Aristotelian Society*, 96, 317–33.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.
- Stephenson, T. (2007). Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30(4), 487–525.
- Vogel, J. (1990). Are there counterexamples to the closure principle? In M. Roth & G. Ross (Eds.), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht: Kluwer.
- Williamson, T. (2005). Contextualism, subject-sensitive invariantism, and knowledge of knowledge. *The Philosophical Quarterly*, 55(219), 213–35.