

3.2 Knowledge attributions and relevant epistemic standards

Dan Zeman

Abstract: The paper is concerned with the semantics of knowledge attributions (K-claims, for short) and proposes a position holding that K-claims are context-sensitive that differs from extant views on the market. First I lay down the data a semantic theory for K-claims needs to explain. Next I present and assess three views purporting to give the semantics for K-claims: contextualism, subject-sensitive invariantism and relativism. All three views are found wanting with respect to their accounting for the data. I then propose a hybrid view according to which the relevant epistemic standards for making/evaluating K-claims are neither those at the context of the subject (subject-sensitive invariantism), nor those at the context of the assessor (relativism), but it is itself an open matter. However, given that we need a principled way of deciding which epistemic standards are the relevant ones, I provide a principle according to which the relevant standards are those that are the highest between those at the context of the subject and those at the context of the assessor/attribution. In the end I consider some objections to the view and offer some answers.

1. Invariantism versus context-sensitivity

It is a fairly widespread view in epistemology today that knowledge attributions (K-claims, for short¹) are context-sensitive. Epistemic terms such as “know” introduce variability in the sentences in which they appear in such a way that different utterances of the same sentence in different contexts get to have different truth values. This variability in truth value is due to factors that pertain, in ways that I will explore shortly, to context. On an intuitive level, such variability is brought to the fore by examples such as Keith DeRose’s (1992) famous Bank Cases, in which the issue is whether someone, driving by the bank on a Friday afternoon in order to deposit her paycheck, finding the lines very long and thus considering postponing the deposit for the next day, knows or doesn’t know that the bank will be open on Saturday. While in the first case there is not much of a difference between the deposit being made the next day or in one of the following days, because there is nothing at stake for the person making it, in the second

case, since there is a lot at stake for the person making the deposit, it is really important that the deposit is made on Saturday. It is thus very important that in the second case the person knows that the bank is open on Saturday. Under the assumption that, as a matter of fact, the bank is indeed open on Saturday, the example is supposed to elicit the intuition that “what is at stake” for the person in the two cases makes the difference between the person knowing that the bank will be open tomorrow and not knowing it. Since “what is at stake” is a contextual matter, we have here a clear example of context influencing the truth value of K-claims.

One important issue that arises in connection to the variability in truth value of K-claims across contexts concerns the nature of the factors responsible for the variability in question. These factors have been thought of in many ways: as possibilities that need to be ruled out, as possibility that could be safely ignored, as a specific amount of evidence that needs to be possessed, as the practical interests or “what is at stake” for the relevant individual (the attributor, the subject or the assessor – according to different views that I am going to explore below). In DeRose’s example above, the factor responsible for the variability in truth value was “what is at stake”. However, I don’t want to commit myself to the claim that stakes are the only contextual factor to which variability of K-claims could be traced down. Thus, when talking more abstractly I will use the term “epistemic standards” to denote whatever one takes these factors to be; I will follow DeRose (and later Jason Stanley) in speaking about stakes as the way in which epistemic standards governing our use of K-claims are influenced by context only as a mere convenience, hoping that both the data to be presented and what I have to say about the data will hold even on other ways of thinking about the factors responsible for the variability of K-claims.²

Any view holding that K-claims are context-sensitive, in the above sense, contrasts with *invariantism*. Invariantism is the view that there is no variation in the truth value of K-claims across contexts. Correlatively, invariantists claim that epistemic terms such as “know” have constant semantic values across contexts.³ However, invariantists have to face the appearance of such a variation, which is vividly brought to the fore by common examples employed by contextualists, like the Bank Cases presented above. Therefore, invariantists need to explain away the appearance of variability. One common strategy used by invariantists was to retort to what has come to be known as “warranted assertibility maneuvers”. According to such maneuvers, what varies with context is not the truth value of sentences, but the assertibility conditions for those sentences: K-claims are either true or

false in any context, and the fact that we are warranted in asserting them doesn't have any influence on their truth value. However, the "warranted assertibility maneuvers" have been seriously discredited by pointing out, as Keith DeRose (2002) does, that one could employ such maneuvers in order to illegitimately evade serious objections and that, in most cases, the way in which the strategy is supposed to work is ad-hoc.

There might be other ways to hold an invariantist position, without appealing to "warranted assertibility maneuvers" (Bach (2005) seems to be an example). Be that as it may, I won't be concerned with invariantism's strategies in this paper. Instead, my aim is to lay down the possible positions which commonly hold that K-claims are context-sensitive and then assess them with respect to the data they purport to explain. The data I will use are the cases presented by Jason Stanley in the Introduction to his book *Knowledge and Practical Interests*. I will present these cases, in quite a bit of detail, in section II. Then in section III I will present the competing views and some criticisms, at the same time evaluating them with respect to the data put forward in section II. The fact that, as it will turn out, each view has problems squaring with the data might give one reason to renounce the idea that K-claims are context-sensitive in any way. In the last section, however, instead of handing over the game to the invariantist, I will put forward an alternative that combines the virtues of all the views examined, but arguably does not inherit their vices.

2. The data: "the Stanley cases"

In the Introduction to his book *Knowledge and Practical Interests*, Stanley presents a battery of cases against which he tests the views he considers later in the book. I will proceed in exactly the same way here. I think that accounting for those cases ("the Stanley cases", as I will call them) constitutes a fair test for any view about knowledge attributions. Any semantic theory that purports to account for how we make knowledge attributions, I claim, should match with the intuitions we have regarding the truth values of K-claims in the specific cases in which those are made. Furthermore, I will claim that the intuitions Stanley presents us as having in the cases that will be put forward are, indeed, "the right ones". Both these two claims are highly controversial. Regarding the first, one might point out that intuitions are too shaky and confused to form the basis for claims about the semantics of natural language expressions. That is certainly true, but one thing should

not pass unnoticed: in the debate about the semantic of knowledge attributions each participant starts with presenting some specific cases that are meant to elicit in us some intuitions. Then it is argued that a specific view, the one the participant in the debate defends, is to be preferred over the competing views precisely because it does a better job in capturing those intuitions. Discarding intuitions in general would be thus like going to war against your most precious ally. However, one might not be prepared to accept the intuitions I claim we have in *all* the cases. That is, philosophers would deny that they have certain intuitions, in specific cases – usually, in exactly those cases that turn out to be problematic for their own views. Here I have nothing else to say except that I hope that the specific intuitions I claim we have in the specific cases I will present will coincide with those of the reader. Besides, I think it is a good general methodological rule that if one starts trusting intuitions, one shouldn't give them up when one's preferred theory enters in conflict with them. I will thus take any view that respects the intuitions as I present them as being more adequate than one that disregards them or renders them in any way inappropriate.

Let us now move to the description of the cases. The first two cases involve a person attributing knowledge to herself in different contexts. To put a bit more flesh on the bones, I will present the cases as they can be found in Stanley's book and then explain my terminology in connection with them. The first case is

Low Stakes (L). Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. It is not important that they do so, as they have no impending bills. But as they drive past the bank, they notice that the lines inside are long, as they often are on Friday afternoons. Realizing that it isn't very important that their paychecks are deposited right away, Hannah says, "I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit our paychecks tomorrow morning."

So in this case the attributor/subject is Hannah, who is in a context in which epistemic standards are low.⁴ She makes a positive K-claim about herself ("I know the bank will be open tomorrow"), which we intuitively think is *true* in the given case. The second case is

High Stakes (H). Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks.

Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah notes that she was at the bank two weeks before on a Saturday morning, and it was open. But, as Sarah points out, banks do change their hours. Hannah says, “I guess you’re right. I don’t know that the bank will be open tomorrow.”

In this case the attributor/subject is again Hannah, who is now in a context in which epistemic standards are high. She makes a negative K-claim about herself (“I don’t know that the bank will be open tomorrow”), which we intuitively think is, again, *true* in the case given.

Now, the two cases in which the attributor coincides with the subject are cases that already raise problems for the invariantist, making her owe us an explanation for the coincidence in truth value of the two K-claims. Whatever strategy she will adopt, the invariantist will eventually end up saying that our intuitions are misleading in (at least) one of the cases. (The same result can be obtained by distinguishing the attributor from the subject, but make them share the epistemic standards that prevail in a given context.) Since I prefer views that do justice to all intuitions, I take this to be an unacceptable result. However, the fact that the attributor and the subject are the same person (or that they share the epistemic standards) also serves to blur the difference between distinctively different views that commonly hold that K-claims are context-sensitive. For, as we will see in the next section, all these views account equally well for the two cases mentioned above. Therefore, in order to better assess the views, we should consider more complicated cases that, instead of obscuring the problems those views might have, bring them to the fore. It is thus a good idea to supplement the data with more complex cases, such as those in which the attributor and the subject of a K-claim are different people, situated in contexts in which the epistemic standards are different. The first such more complex case is

Low Attributor – High Subject Stakes (LA-HS). Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Two weeks earlier, on a Saturday, Hannah went to the bank, where Jill saw her. Sarah points out to Hannah that banks do change their hours. Hannah utters, “That’s a good point. I guess I don’t really know that the bank will be open on Saturday.” Coincidentally, Jill is thinking of going to the bank on Saturday, just for fun, to see if she meets Hannah there. Nothing is at stake for Jill,

and she knows nothing of Hannah's situation. Wondering whether Hannah will be there, Jill utters to a friend, "Well, Hannah was at the bank two weeks ago on a Saturday. So she knows the bank will be open on Saturday."

Here, we have an attributor, Jill, who is in a context in which epistemic standards are low, and a subject, Hannah, who is in a context in which epistemic standards are high. The attributor, Jill, makes a positive K-claim about the subject, Hannah ("she knows the bank will be open on Saturday"), which we think is *false* in the case given.

The second more complex case that needs to be accounted for is

High Attributor – Low Subject Stakes (HA-LS). Hannah and her wife Sarah are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. Since they have an impending bill coming due, and very little in their account, it is very important that they deposit their paychecks by Saturday. Hannah calls up Bill on her cell phone, and asks Bill whether the bank will be open on Saturday. Bill replies by telling Hannah, "Well, I was there two weeks ago on a Saturday, and it was open." After reporting the discussion to Sarah, Hannah concludes that, since banks do occasionally change their hours, "Bill doesn't really know that the bank will be open on Saturday."

In this final case, we have an attributor, Hannah, who is in a context in which epistemic standards are high, and a subject, Bill, who is in a context in which epistemic standards are low. The attributor, Hannah, makes a negative K-claim about the subject, Bill ("Bill doesn't really know that the bank will be open on Saturday"), which we think is *true* in the case given.

The following table summarizes the intuitions we have in the cases presented:

Case	K-claim	Truth-value
L	<i>Positive</i>	T
H	<i>Negative</i>	T
LA-HS	<i>Positive</i>	F
HA-LS	<i>Negative</i>	T

With these results in front of us, I can now proceed to present and assess the competing views.

3. The views and their problems

There is more than one way in which K-claims can be context-sensitive. In this section I will briefly describe three views incorporating the thesis that the truth value of K-claims varies across contexts. After presenting each view, I will mention some common objections and then confront it with the data presented in the preceding section. How well each view comes out from this confrontation will be the criterion on the basis of which we should choose the best view.

3.1. Contextualism

According to epistemic contextualism, terms like “know” denote different relations in different contexts of utterance. This has the consequence that, as DeRose puts it in his pioneering article, “the truth conditions of [K-claims] vary in certain ways according to the context in which the sentences are uttered” (DeRose, 1992: 914). This variation in semantic value across contexts is in turn responsible for the variation in truth value (that is, for K-claims being context-sensitive in my sense of the term). Contextualists differ in conceiving the “certain ways” DeRose speaks about, but their main claim is quite straightforward.

Thinking of “know” as denoting different relations in different contexts of utterance is not the only way to hold that K-claims have different semantic values in different contexts. Bach (2005) distinguishes between two varieties of epistemic contextualism: “the indexed version” and “the relativized version”. The difference between these varieties lies in the way the epistemic standard plays a role in establishing the semantic content of K-claims. Thus, according to the first version, K-claims are elliptical for sentences in which epistemic standards (*e*) are indexed to the verb “know” (like in “S knows_{*e*} that *p*”); according to the second version, K-claims are elliptical for sentences in which “know” is relativized to epistemic standards (like in “S knows relative to *e* that *p*”). However, it is arguable that both these versions will have to face some serious objections. Let me mention some of these in what follows.

Contextualism in general faces some notorious problems which transfer also to epistemic contextualism. The first is that, as it stands, contextualism cannot account for the sense of disagreement we detect in exchanges like the following: A utters the sentence “Avocado is tasty”, and B replies with

“No, it is not”. Disagreement is a thorny issue, and I’m not going to address it here⁵, but one might legitimately grant that in exchanges like these we at least have the intuition that A and B disagree.⁶ Now, the same seems to happen in the epistemic case: if, say, the dogmatic utters “Moore knows he has hands” and the skeptic replies “No, Moore does not know he has hands”, there is a strong intuition that the skeptic and the dogmatic disagree.⁷ But if contextualism were right, there won’t be any disagreement in such a case: since the skeptic and the dogmatic are in different contexts, in which different epistemic standards prevail, the semantic values of the two sentences will not be contradictory. So, contextualism cannot account for disagreement. Of course, the contextualist is free to deny that such cases really *are* cases of disagreement, but this comes with a cost: to maintain that ordinary speakers don’t know the workings of their own language, that they are semantically “blind”. And semantic blindness has been found by many, including Schiffer (1996) and Hawthorne (2004), to be too hard a pill to swallow.⁸

Second, there is the problem of reporting K-claims. As Mark Richard (2004) has asked, what exactly is one saying in reporting a debate between two people to the effect that one subject knows something? According to contextualism, since the semantic value of K-claims depends on the context in which they are uttered, it is the context of the reporter that establishes the semantic value of the reported K-claim. Thus, when one is reporting an exchange between a dogmatist and a skeptic pertaining to the issue whether the dogmatic knows she has hands, the K-claim will have the semantic value it has in the reporter’s context, which might well be different from that it has in both the dogmatic and the skeptic’s context. Yet it is quite unreasonable to suppose that what was initially said could differ radically from what it is said by the reporter. Contextualism predicts a totally unintuitive result in cases like these.

Finally, as John MacFarlane (2005a) has forcefully argued, epistemic contextualism has another unpleasant consequence: namely, admitting that we were wrong and therefore retracting an earlier K-claim when we come to know better becomes meaningless. To see this clearer, imagine the following situation: we are both in a cab that carries us from the city to the airport. When the cab passes by the beautiful medieval city center, I utter “It is very nice around here”. Twenty minutes later, when we are approaching the airport, you say “It is not very nice around here”; your claim doesn’t make me retract my earlier remark. If “know” would be context-sensitive in the way “here” is, the same would happen with K-claims; but in

their case it is different. When I move from a context in which epistemic standards are low to one in which they are high, I will admit that I was wrong and therefore retract my earlier K-claim that I knew. Retraction seems to be a common practice among (rational) agents; yet, contextualism renders it totally meaningless.⁹

I don't take these objections to lead to a definitive rejection of contextualism, but they do indeed pile up a certain amount of evidence against the view.¹⁰ There is, however, a closely related but still significantly different view that evades some of the objections presented. *Non-indexical contextualism*¹¹ (NIC), as MacFarlane (2009) calls it, is a view according to which K-claims are context-sensitive, but their semantic value is constant across contexts of utterance. NIC is a development of the Kaplanian framework in which the crucial distinction is that between contexts of utterance and circumstances of evaluation.¹² A context of utterance is the particular situation in which an utterance of a sentence is produced, whereas circumstances of evaluation are "both actual and counterfactual situations with respect to which is appropriate to ask for the extensions of a given well-formed expression" (Kaplan, 1989: 502). Kaplan's claim was that circumstances of evaluation comprise, besides possible worlds, also times and (maybe) locations. Following this model, a number of non-indexical contextualist views have been proposed for a number of discourses. According to the work done by Kölbel (2004) and Lasersohn (2005) in the evaluative sphere, in order to evaluate sentences containing evaluative terms, we need to introduce new parameters into the circumstances of evaluation ("perspectives" in Kölbel's view, a "judge" parameter in Lasersohn's), because the traditional parameters (world and time) are not enough to yield a definite truth value for those sentences. The newly introduced parameters get activated, as in the case of contextualism, by features of the context in which an evaluative sentence is uttered; the difference is that this time, instead of providing elements that will be part of the semantic value of the sentences, these features provide parameters in the circumstance of evaluation that is operative in the context. In the same vein, a non-indexical contextualist about knowledge attributions will hold that circumstances of evaluation will comprise, besides possible worlds, also epistemic standards (and maybe other unorthodox parameters as well, although she is not committed to that). K-claims will be evaluated with respect to such enriched circumstances, and epistemic features of the context in which a K-claim is uttered will provide a parameter in the circumstance of evaluation (an epis-

temic standard), and not an element in the sentence's semantic value in that context.¹³

Although, from a semantic point of view, there are significant differences between these two views, they both face the same problem with respect to the Stanley cases. Let me illustrate with considering in detail how contextualism fares with respect to them. Note first that contextualism easily explains the coincidence in truth value of the K-claims made in the **L** and **H** cases. Since the positive and the negative K-claim are uttered in different contexts, they will have different semantic values; therefore, there is no problem with both being true. Moreover, the **HA-LS** case is also accounted for: since the relevant epistemic standards are those at the context in which the K-claim is uttered, and since in the attributor's context the epistemic standards are high, this will make her negative K-claim true. And this is also what our intuitions tell. However, the problem appears in the **LA-HS** case. Since the relevant epistemic standards are those at the context in which the K-claim is uttered, and since in the attributor's context the epistemic standards are low, this will make her positive K-claim true. But this is *not* what our intuitions tell. Therefore, contextualism has problems squaring with the data.

It is easy to see that NIC will follow the same pattern as contextualism in all the four cases. The only difference, as already explained above, stems from the role different factors play in arriving at the truth values of sentences in contexts. Thus, the coincidence in truth value of the two K-claims made in the **L** and **H** case, respectively (one positive, the other negative) is explained under the current view not by a difference in semantic content, but by the fact that the two sentences get evaluated with respect to different circumstances: one in which epistemic standards are low (case **L**), the other in which epistemic standards are high (case **H**). Thus, although one K-claim is the negation of the other, there is no problem with both being true, since they get evaluated with respect to different circumstances. The same considerations are true in the other two cases: what the context of the attributor contributes is not an element in the semantic value of the sentence uttered, but an element in the circumstance. However, this does not prevent NIC having the same problem as contextualism: the troublesome case is **LA-HS**. I thus conclude that both views fall short of accounting for the Stanley cases.

3.2. Subject-sensitive invariantism

One might hold a view about the context-sensitivity of K-claims which, instead of focusing on the attributor of a K-claim, focuses on its subject. Such a subject-centered view is the view known under the label *subject-sensitive invariantism* (SSI).¹⁴ SSI's main claim is that whether a subject knows something crucially depends on the subject's epistemic standards, with these standards being in turn dependent on non-standard features of the subject itself and its context. Those non-standard features of the subject have to do with the subject's interests, with "what is at stake" for the subject – the variation of those factors across contexts being responsible for the variation in the subject's epistemic standards. Although SSI's main insight is into the nature of knowledge, linking it with practical interests and action, the view has also direct consequences for semantics. Thus, one of the semantic claims supported by SSI is that epistemic terms' semantic value does not change with the context of their utterance. Consequently, the K-claims in which they appear also have constant semantic values across contexts of utterance. In this, SSI agrees with invariantism and NIC, and contrast with contextualism. However, from this fact we shouldn't conclude that K-claims also have constant truth values across contexts. Given that the subject's epistemic standards vary from context to context, whether a subject knows something will be a context-dependent matter as well. The variation in epistemic standards clearly affects the truth value of K-claims, allowing that the same K-claim could be true in one context and false in another.

The fact that, according to SSI, K-claims have constant semantic values across contexts might provide the view with better answers to the problems encountered by contextualism. But that is not to say that SSI has no problems of its own. Schaffer (2006), for example, argues that an important drawback of the view is that it doesn't square very well with extremely plausible views about those aspects of knowledge that pertain to its social role: activities such as inquiry, expertise, testimony, conversational norms and the fact that we take knowledge to have a high value. Blaauw (2008) argues that SSI has severe problems with embracing a widely accepted principle concerning transmission of knowledge through memory. Finally, MacFarlane (2005b), although not explicitly rejecting SSI, accuses the view of "knowledge laundering" and of clashing with a fairly entrenched view about testimony.

Let us see, however, how SSI accounts for the Stanley cases. First, let us note that SSI has no problem with the **L** and **H** cases: since the positive and the negative K-claim are made in different contexts, and since the relevant standards are those at the context of the subject (either identical with the attributor, or having the same epistemic standards as the attributor), each K-claim comes out true. Further, SSI nicely accounts for the **LA-HS** case: since the relevant epistemic standards are those at the context of the subject, and since in the subject's context the epistemic standards are high, this will make the attributor's positive K-claim false. And this is also what our intuitions tell. However, as MacFarlane (2005a) points out and DeRose (2004) explicitly argues, SSI fares badly with respect to the **HA-LS** case. Since the relevant epistemic standards are those at the context of the subject, and since in the subject's context the epistemic standards are low, this will make the attributor's negative K-claim false. But this is *not* what our intuitions tell. Therefore, SSI fails the test too.¹⁵

3.3. Relativism

A third view which purports to give a semantic account of K-claims is the one recently proposed by MacFarlane (2005a)¹⁶: relativism. MacFarlane holds that K-claims are "assessment-sensitive", a new phenomenon that he claims to have unveiled for a number of discourses. As well as SSI and NIC, relativism agrees with invariantism that epistemic terms and the sentences containing them have constant semantic values across contexts. But, in opposition to invariantism, relativism claims (together with SSI and NIC) that the truth values of K-claims in different contexts are not constant. However, the reasons why this is so are entirely different from those invoked by NIC and SSI. Let me spend a while presenting in detail the relativist framework.

Recall the Kaplanian distinction, which played an important role in NIC, between context of utterance and circumstance of evaluation. In Kaplan's framework, the truth of sentences is relativized to both these factors: a sentence is true relative to a context *and* to a circumstance of evaluation. However, this is consistent with utterances having truth values absolutely: given a context, and given that any sentence uttered in that context has to be evaluated at the circumstance operative in that context, that utterance of the sentence will be true or false absolutely. Thus, not only that a context of utterance already comprises all that is necessary for the evaluation of a

sentence in that context, but it also necessarily *determines* the circumstance with respect to which the sentence has to be evaluated.¹⁷ Now, in MacFarlane's view, sentences are also doubly relativized, but this time to contexts of utterance and *contexts of assessment*. A context of assessment is simply a context from which a sentence, uttered at a (possibly different) context, is evaluated for its truth. The context of assessment supplies the circumstance with respect to which the sentence is evaluated. The context of utterance can coincide with the context of assessment, but in the case in which it doesn't, there is nothing in the context of utterance that determines the circumstances with respect to which the sentence has to be evaluated. The circumstance operative in the context of assessment trumps, so to speak, the circumstance operative in the context of utterance. This has the consequence that utterance-truth is *not* absolute, in contrast with NIC. The same utterance could thus get different truth values, depending on the context from which it is assessed, with the circumstances of that context now playing the crucial role in arriving at the truth value of the sentence.¹⁸ In order to get relativism about knowledge attributions, all we have to add to this machinery is the claim that circumstances comprise, besides possible worlds, epistemic standards (and maybe other unorthodox parameters if one is relativist about other domains as well).

Now, it is arguable that relativism does not encounter the problems that the other views did. But the data that motivated the move towards relativism could and have been disputed. However, I will not pursue this issue here; I think that, even granting the solidity of the data MacFarlane relies on, his view is still not trouble-free. To see this, let us again consider how the view squares with the Stanley cases. Now, in order to be able to compare relativism with the other views in how well it squares with the data, I have to make the following simplification: I will consider that, in the scenarios considered, the assessor is the attributor herself. In other words, I will consider only the case in which the context of assessment and the context of utterance coincide. This simplification might be thought of as belittling relativism. But it is not: first, note that the case in which the context of assessment coincides with the context of utterance is certainly one that the relativist has to, and as we have seen, does indeed allow. This, of course, is not to deny that there is a distinction between them. Second, my point can be made even if we introduce more complicated cases featuring three distinct characters: the attributor, the subject and the assessor, each being in different context in which different epistemic standards prevail. However,

this would only complicate matters and make the comparison between the competing views more cumbersome. Thus, my simplification is harmless.

Now, let us see how the view comes out when confronted with the data. Let us first note that relativism handles easily the **L** and **H** cases: since the positive and the negative K-claim are made (and assessed) in different contexts, different epistemic standards are part of their respective contexts of assessment/utterance; thus, each K-claim is true relative to the epistemic standard that is operative in the context of the assessor/attributor. Further, unlike SSI, relativism handles nicely the **HA-LS** case: since the relevant epistemic standards are those at the context of the assessor/attributor, and since in the assessor/attributor's context the epistemic standards are high, this will make the assessor/attributor's negative K-claim true. And this is in line with what our intuitions tell. In contrast, consider the **LA-HS** case: since, again, the relevant epistemic standards are those at the context of the assessor/attributor, and since in the assessor/attributor's context epistemic standards are low, this will make the assessor/attributor's positive K-claim true. But this is *not* what our intuitions tell. Therefore, relativism has troubles with accounting for the Stanley cases.¹⁹

Before proceeding to the next section, let me first address an objection that I have been presented with on several occasions. One might oppose the above diagnosis on the following grounds: One of the core claims of relativism is that there is no absolute utterance-truth. This claim, I take it, could be put in more mundane words as the claim that there is no neutral perspective, something like a God's-eye view from which our utterances get an established truth-value once and for all. But – and this is the answer to the diagnosis – the Stanley cases, as I have presented them, presuppose exactly such a neutral view. That is, we have the intuitions (I claim) we have just because we situate ourselves (or we are, as readers, asked to situate ourselves) in a neutral, all-encompassing position; in fact, we are never in a position to judge from outside, being instead confined to judge only from within our own particular perspective. Since presenting the data in the way I presented them implies reliance on the God's-eye view, this simply begs the question against relativism.

Despite its initial plausibility, the objection is misguided. It is certainly true that in presenting the cases as I did I gave the reader substantial information about the contexts in which the characters of the cases were situated. Whether this amounts to tacitly relying on a God's-eye view or not I'm not entirely sure, but the point is that I cannot see how else I could even describe such cases without giving the reader such kind of information. If

one wants to elicit some intuitions, the more information is given about a case, the likelier to have a clearer intuition about it. Moreover, this kind of information about the context of hypothetical attributors, subjects or assessors of K-claims has to be given by *any* of the participants in the debate – including the relativist. The result that relativism has troubles with the Stanley cases, I therefore conclude, is not endangered by the objection just considered.

4. Going invariantist?

The problems for the three views discussed above may motivate one to shy away from the idea that K-claims are context-sensitive in any of the ways I described above. The following table synthesizes how well the views coped with the data:

	L	H	LA-HS	HA-LS
Contextualism/NIC	OK	OK	Not OK	OK
SSI	OK	OK	OK	Not OK
Relativism	OK	OK	Not OK	OK

I haven't discussed invariantism's problems in detail, but the results in the table above don't seem to situate the context-sensitivity views in a much better position than invariantism. However, before giving up the idea that K-claims are context-sensitive, let me try to spell out a way to modify one of the views that holds on to this idea so that it would get the right results in the Stanley cases.

In reply to some objections to contextualism, DeRose has argued that the view is not committed to always considering the attributor's context as "calling the shots". The contextualist is free to say that sometimes the epistemic standards at the context of the *subject* are those which settle the semantic value of "know". Thus, "[t]here's nothing in contextualism to prevent a speaker's context from selecting epistemic standards appropriate to the subject's context, even when the subject being discussed is no party to the speaker's conversations" (DeRose, 2004: 348). This move does indeed save contextualism from predicting the wrong results in the problematic Stanley case (namely, the **LA-HS** case). However, the move does not save contextualism from the other objections I mentioned in section III.

It is not my intention to claim that contextualism will not eventually be able to provide acceptable solutions to the problems mentioned. But instead of working out a suitable version of contextualism, what I want to do is to see whether one of the other views presented – namely, relativism – cannot be modified such that to be able to account for all the intuitions in the Stanley cases. To that end, I will start from DeRose’s insight above. Thus, the proposal I want to put forward will consist in incorporating this insight into a view that is not contextualist but still holds that K-claims are context-sensitive. What I want to claim is that the flexibility that contextualism gains by allowing the subject’s standards to settle the semantic value of “know” in the attributor’s context is also available to relativism (although not by the same means, of course). In this connection, it is instructive to have a look on the debate surrounding epistemic modals. Here, too, we find basically the same contenders, with relativism being one of them (due to the work of MacFarlane (forthcoming) and others). However, as Dietz (2007) has noted, relativism about epistemic modals has problems with what he calls “ignorant assessor cases” – cases structurally similar to with my **LA-HS** case. Now, interestingly, in the case of epistemic modals MacFarlane has acknowledged the problem and has agreed that what is needed is more flexibility. The solution, he says, must be one in which

[t]he semantics must track both what is known by the asserter and what is known by the assessor, and then amalgamate these two bodies of knowledge into a single body of known facts with respect to which the epistemic modal is to be evaluated. (MacFarlane, forthcoming: 53)

So, my claim is that we need to allow for such flexibility within a relativist view about knowledge attributions. How exactly we should design the details of such a flexible relativism is a tricky issue. What we want is a view holding that K-claims have constant semantic values across contexts, but allowing their truth-value to vary; moreover, sentence-truth will be doubly relativized, both to contexts of utterance and to contexts of assessment. Epistemic standards, as in the case of NIC, will be part of the circumstances of evaluation. Now, since the context of assessment provides the circumstance with respect to which K-claims will be evaluated, the key point in order to gain more flexibility will be to allow contexts of assessment to be trumped by other contexts, in the sense that it will be the circumstances at those other contexts that will be relevant for evaluating a K-claim. Paraphrasing DeRose’s claim, there is nothing in relativism to prevent an assessor’s context from selecting epistemic standards appropriate to other contexts. In the simplified version of relativism I was operating so far

(in which the context of assessment and the context of utterance coincide), this flexible approach will allow for the epistemic standards of the subject to be the relevant ones in evaluating a certain K-claim *in a given context of assessment*. Thus, the view remains relativist, but with the twist of being able to allow for the desired flexibility.

Now, this flexibility must not be confounded with arbitrariness. It would be too easy if allowing for flexibility within a relativist framework (or, for that matter, in any framework) would be enough. I take it that in order to have a solid and coherent view, something must be said about how this flexibility should be implemented. This surely is a quite demanding task, and I'm not taking myself to be able to provide a solution to such an important problem. However, in the remainder of the paper I will try to flesh out a way to implement flexibility in a principled way.

Since the modification of the relativist view to the point of allowing more flexibility was motivated by the failure to account for all the Stanley cases, it might be a good idea to follow closely the intuitions we had in those cases. Reflecting on the case that created the problem (the **LA-HS** case), one immediate feature of it stands out: the problem came from the fact that the epistemic standards in the subject's context were high. On the other hand, the case that relativism handled successfully (besides the simple **L** and **H** cases – namely, the **HA-LS** case) was the one in which the epistemic standards in the assessor/attribution's context were high. This suggests the following principle (aimed as guiding an assessor in evaluating K-claims):

Highest Standards Principle (HSP): In assessing a K-claim, the relevant epistemic standards are those at the context in which the epistemic standards are the highest.

What is important to note right away is that “highest” here is intended to mean “highest between the context of the assessor and the context of the subject”.²⁰ This prevents HSP from collapsing into a skeptical invariantist view, in which “highest” would always mean the skeptic's standards. That is, “highest” is comparative and contextual, not absolute – even if “contextual” now means taking into account more than one context. Assessment-sensitivity allows K-claims to get different truth values according to the epistemic standards operative in the context of the assessor; since the assessor's context changes, so do the truth values of K-claims. According to HSP, the same holds: K-claims still change their truth value as a function of

the context of assessment; the only difference now is that the assessor is allowed to retort to the standards at the context of the subject when she assesses a given K-claim.

It is easy to see that simply by following HSP the Stanley cases are satisfactorily handled. Since the assessor/attributor is the same as the subject, the **L** and **H** cases don't pose any problem. In the **HA-LS** case, since the highest epistemic standards are those at the context of the assessor/attributor, they will be the ones selected as relevant by the HSP. This will make the assessor/attributor's negative K-claim true. And this also what our intuitions tell. In the **LA-HS** case, since the highest standards are those at the context of the subject, they will be selected as the relevant ones by the HSP. This will make the assessor/attributor's positive K-claim false. And, here too, this is also what our intuitions tell.

Let us now survey some objections to the account just given. One problem that immediately springs to mind is related to cases in which the assessor and the subject are in different contexts. The view implies that, whatever her context is, the assessor will be able to decide in which of the two contexts (her own or the subject's) the epistemic standards are the highest. That is, it is presupposed that in any situation the assessor could have a glance into the subject's context, compare the epistemic standards found there with those in her own context, and then decide which epistemic standards are the highest. Yet, nothing guarantees that this is possible. In fact, for various reasons, there are plenty of cases in which this is actually impossible. The frequency of such cases might undermine the view, rendering it if not incoherent, then practically useless.

I have to admit that this is a problem for the view, but I don't think the situation is hopeless. One first try to fix this problem would be to note that, as in general with interpreting other human beings, one's own situation might function as a guide for judging others' situation. In order for one to apply the principle, sometimes all what is needed is being aware of one's own epistemic situation. And this kind of awareness, although not always guaranteed either²¹, is nevertheless more common than that of the situation of others. Thus, if the assessor is in a context in which epistemic standards are low, she should automatically privilege the subject's standards; conversely, if the assessor is in a context in which epistemic standards are high, she should automatically privilege her own standards. Although this strategy might work in some situations, it certainly has its limitations: for it would only work on the premises that the assessor is at the extreme limits of a series of standards – that her standards are either the lowest or the

highest in this series. And, even if the idea of a strict ordering of epistemic standards seems hard to defend, there is at least some plausibility in the idea that there are more than two contexts with different epistemic standards in which assessors could find themselves. So this solution cannot entirely assuage the worry.

However, I think that even accepting that assessors will mistakenly evaluate K-claims for reasons having to do with the opacity of the subject's context is not disruptive for the view. First, let me note that there being situations in which the assessor has no access to the subject's context is a kind of limitation that applies to us as human beings, and which manifests itself in any area in which we behave as interpreters. "For all the data about knowledge attributions show, it could well be that we often attribute knowledge to people who don't have it and often resist attributing it to people who do have it" (Bach, 2005: 86). Bach himself is an invariantist, but I don't see any reason why someone embracing context-sensitivity couldn't avail herself of such fallibility. Second, the fact that assessors would make mistakes in assessing K-claims, for reasons having to do with the opacity of the subject's context, should not count against the theory. Consider the view we have about demonstratives like "this". The semantic value of an utterance of "this" is supposed to be the object that is intended, or is pointed to, by the speaker. Now, imagine the following situation: you overhear a dialogue taking place in a room in which you have no access, in which the expression "this" occurs. Not being able to see what has been pointed to, or not having a clue about what the speaker intended to refer to, you are not able to fully understand what has been said. Is this a reason to modify our theory about demonstratives? No. Whatever weaknesses that theory might have, it certainly doesn't come from the failure, in some cases, to identify the referent of demonstratives. The same, it seems to me, applies to the case of knowledge attributions: the impossibility to glance in the subject's context in some situations should not count against the theory.

Another complaint might be voiced at this point, in connection with the **LA-HS** case. Suppose assessors follow HSP in assessing K-claims; also, suppose there is no barrier for the assessor to glance into the subject's context and retrieve the relevant information. Everything goes well, and the assessment of a K-claim about the subject conforms to the intuitions we have in such a case. But, the objections goes, it is precisely because the assessor has followed the principle that she has actually managed to change her own context so that the case won't count as a **LA-HS** case anymore! This is a fair complaint, but I think it is not as unpalatable as it might seem

at a first glance. For, I take it, the destruction of the **LA-HS** case is a desirable consequence of the fact that, at least in some situations, it is rational to inquire into the subject's context in order to gather information about her standards before assessing a given K-claim. Situations in which, for example, it is very important for the assessor to know whether the subject knows something would presumably be of this kind. Of course, the (rational) assessor might not succeed in retrieving the required information, but the point remains. Thus, although the view cannot totally prevent assessors from making mistakes in assessing K-claims, it renders them rational, as in fact they are.²²

Summing up: I've proposed a way to modify relativism such as to avoid the problems posed by one of the Stanley cases. The view proposed is committed to the following claims: epistemic terms and K-claims have constant semantic values across contexts of utterance; the truth-value of K-claims varies across contexts of assessment; the relevant epistemic standards for assessing a K-claim are not necessarily those of the assessor (the idea of flexibility). The HSP fleshes out the view, with the result of portraying a knowledge-attributor that is fallible but rational. Moreover –and this is supposed to be the view's strong point– unlike the other views scrutinized, the view proposed leaves the intuitions with which we started intact. If indeed coherent, it provides a serious alternative to those views, without having to succumb to invariantism.

Notes

1. More precisely, I will use the term "K-claim" to denote any *sentence* involving the attribution or denial of knowledge to a subject. I will also speak both of *positive* K-claims, when knowledge is attributed to a subject (sentences of the form "X knows that p") and of *negative* K-claims, when knowledge is denied to a subject (sentences of the form "X does not know that p" – in both cases X signifying the subject and p a proposition).
2. Schaffer (2005) argues that "what shifts" (that is, what is different from context to context) must be what he calls "epistemic alternatives" – basically, relevant possibilities that need to be ruled out by a subject in order to count as a knower. Schaffer rejects the view that it is "epistemic standards" that shift, but his use is different from mine. As I said, I will use "epistemic standards" as a blanket term, and use stakes as a particular way of thinking about the factors responsible for the variability.
3. I'm using the term "semantic value" to not prejudge the issue whether what is expressed by sentences in contexts are propositions, in the traditional sense

(contents that are true or false relative to possible worlds), or “relativized propositions” (contents that are true or false relative to possible worlds and other unorthodox parameters, such as time, location, various standards, etc.). Although I will not directly address this issue here, some arguments for or against the positions I will discuss are arguments for or against one of these ways of conceiving semantic content. I thus use “semantic value” as a blanket term for whatever one might think is expressed by sentences in contexts.

4. I’m not presupposing, in what follows, that there is a strict ordering of epistemic standards: all I need is a clear contrast between two sets of epistemic standards: “low” and “high”. Think of “low” and “high”, if you wish, as describing ordinary contexts and skeptical contexts, respectively.
5. I am going to assume, though, that a necessary condition for two people to disagree is for them to endorse sentences having contradictory semantic values.
6. For talk about the intuition of disagreement in matters of taste, see Lasersohn (2005).
7. The same point is made by Richard (2004) and Kompa (2005). That the dogmatist and the skeptic don’t seem to talk past each other has been enough reason for Bach (2005) to conclude that contextualism doesn’t really provide the much-trumpeted solution to the skeptical puzzle.
8. See, though, (DeRose 2006) for an attempt to assuage this worry.
9. One might reply to the last objection that the analogy with indexicals is not a fortunate one. Contextualists have been keen on taking as models for epistemic terms other context-sensitive expressions than indexicals. Thus, one common contextualist claim was that “know” is similar to gradable adjectives (for example, (Cohen 1999)). But as Stanley (2004) and Partee (2004) have forcefully argued, this analogy is worse than the one with indexicals, because epistemic terms fail two main tests for gradability: they don’t allow modifiers like “very” or “really” and they don’t allow for comparatives. (See though (Ludlow 2005) for the claim that “know”, being a verb, cannot pass these tests which are designed for adjectives. Ludlow finds Stanley’s arguments correct, but misplaced.) But even if these arguments eventually fail, the analogy between gradable adjectives and epistemic terms is of no help: people don’t seem to retract sentences containing gradable adjectives.
10. There are other objections to contextualism that I didn’t consider, such as those in Williamson (2005) to the effect that contextualism cannot account for the preservation of information through memory and testimony, that it cannot give a coherent picture of our practical reasoning, or that it violates the factivity of knowledge.
11. NIC has been present on the market under different names. Thus, the view is identical with what Kölbel (2004) has called “genuine relativism” (in contrast with “indexical relativism”). Lasersohn (2005) refers to it simply as “relativism”. However, MacFarlane (2009) has claimed that NIC is not relativism in

its true sense: the difference between the views lies in the different kind of relativization each view takes to be essential. The difference will become clear in my discussion of relativism below.

12. Compare also with (Lewis 1998) distinction between context and index. There are significant differences between Lewis' view and Kaplan's, but they are irrelevant for my purposes here.
13. Kompa (2002, 2005) seems to defend NIC for knowledge attributions.
14. There are two prominent versions of SSI in the literature: Stanley's (2005) Interest-Relative Invariantism and Hawthorne's (2004) Sensitive Moderate Invariantism. For the purposes of this paper, I will count these two views as equivalent.
15. As far as Interest-Relative Invariantism is concerned, Stanley is well aware of this difficulty for his view. This is why he argues at length (2005, chapter 5) that, after all, the intuition we have in the **HA-LS** case is not the one we should have. Hawthorne is also aware of the problem cases like **HA-LS** pose for his Sensitive Moderate Invariantism. His strategy is to explain away the intuition we have in such cases, by employing the idea of projection: people tend to project their own epistemic standards to the subjects of the K-claims they are evaluating. Thus, both authors are thus willing to give up the task of accounting for all the intuitions. As I said before, my position is that a view that accounts for all the intuitions is to be preferred over a view that explains (some of) them away.
16. See also (Richard 2004). It is not clear, however, to what extent Richard's view overlaps with MacFarlane's.
17. This is not to say that we don't need other circumstances to evaluate a given sentence in a context. Complex sentences sometimes comprise other sentences whose truth value must be evaluated at other circumstances than those of the context of their utterance. But such complex sentences also comprise expressions whose unique role is precisely that of shifting the circumstance: operators. The existence of operators in language has been used by Kaplan (1989) as a premise in an argument in favor of introducing unorthodox parameters in the circumstances, such as times and (perhaps) locations.

However, this general argument, known as "the operator argument", has been questioned. First, one might doubt, as Stanley (2005) does in his criticism of both relativism and NIC, that there is not enough evidence that there are expressions in English that could play the role of operators shifting epistemic standards. But as Ludlow (2005) has pointed out, there seem to be plenty of natural language expressions of the kind Stanley looks for: expressions like "for X", "according to Y's standards" or "by the standards of science" are just some examples. However, there is a deeper worry lurking here, and a number of authors have been prone to forcefully point it out. The worry, as Stanley (2005), King (2003) and, more recently, Cappelen and Hawthorne (2009) claim, is whether the expressions that the relativist takes as circumstance-

shifting operators are rightfully interpreted as such. King (2003), following the lead in contemporary linguistics, argues that time, for example, should not be part of the circumstances with respect to which we evaluate tensed sentences, because tense and temporal expressions such as “yesterday” are better interpreted as quantifiers operating on times. A similar point could be made with respect to the introduction of other unorthodox parameters in the circumstances (for all the authors mentioned, the only acceptable parameter is the world). Cappelen and Hawthorne (2009) dedicate an entire chapter to the operator argument, and argue against it on a number of scores. In the face of such challenges, the relativist (as well as the non-indexical contextualist) has two options. The first is to try to show, as Recanati (2007) does, that the interpretation of tense and temporal expressions as quantifying over times is not forced on us, and thus that the operator argument could be employed to argue for the introduction of time in the circumstances. Similar considerations might then be used to argue that the operator argument works for other unorthodox parameters as well (in our case, epistemic standards). The second option is to point out, as MacFarlane does, that the introduction of unorthodox parameters in the circumstances is not dependent of the existence of operators in the language and is thus independently motivated. One such independent motivation is connected to the kind of contents one needs in a broader theory of communication. Another motivation could come from the fact that postulating such contents simply gives us a better explanation of the data to be accounted for. I take myself to follow this last strategy here. So, although this discussion is by far not conclusive, and much more needs to be said to substantiate these claims, I will take it that the introduction of unorthodox parameters in the circumstances such as epistemic standards is not a theoretically unacceptable move.

18. Although non-absoluteness of utterance truth is the main feature of relativism, what I will say below does not trade on this feature. The reason is that in the way I constructed the cases, there will never be an assessment of the same utterance in two different contexts of assessment; instead, in each case a different sentence will be used. However, I could have constructed the cases in such a way to put this feature of the view at work; the results I claim we get would have been exactly the same. See below.
19. How do things look in the more complex cases alluded to above, in which the context of assessment is different from the context of utterance? Let us imagine such a scenario. Let attributor AT make a K-claim about subject S, and let them be in different contexts in which different epistemic standards prevail. Now, let assessor AS, situated in a context in which epistemic standards are different from those in AT and S’s contexts (or at least different from one of those), evaluate AT’s utterance. Imagine now that AS is in a context in which epistemic standards are low, whereas S is in a context in which epistemic standards are high. It is my intuition that any assessment by AS of a positive

K-claim about S as true is incorrect in such a scenario. Moreover, this seems to me to be the case regardless of the prevailing standards at AT's context. But even if this last claim is debatable, imagine that AT is also in a context in which epistemic standards are low. Here the intuition I have certainly goes in the direction I mentioned. So, there seems to be at least one case (call it **LAS-LAT-HS**), quite similar to **LA-HS** above, that creates problems for the relativist.

20. What about the more complicated case in which the attributor, the subject and the assessor are all in different contexts? HSP instructs the assessor to go for the highest standard. The difference between this case and the simpler one is that now the attributor's context might be the one in which the epistemic standards are the highest, and not that of the subject. Does that have any influence on the view? I don't think so. For it seems to me that a positive K-claim assessed by an assessor in a context in which epistemic standards are low, K-claim made by an attributor in a context in which epistemic standards are high, even if those in the subject's context are low (case **LAS-HAT-LS**), is still false. But intuitions are shakier in such complicated cases.
21. Stanley (2005) includes in his original battery of cases one in which the subject is unaware of the epistemic standards in play in her own context (the case he dubs "Ignorant High Stakes"). I haven't considered that case in presenting the data because my case against the three views presented could be made without it.
22. Compare this idea with what Wright (2007) has to say in reply to Dietz's (2008) objection to relativism in the case of epistemic modals. Wright's point is highly significant in that the case he is responding to is one structurally similar to my **LA-HS** case.

References

- Bach, K.
2005 The Emperor's New 'Knows', in Preyer and Peter (2005), pp.51-90;
- Blaauw, M.
2008 Subject Sensitive Invariantism: *In Memoriam. The Philosophical Quarterly*, 58: 318-325;
- Cappelen, H. and Hawthorne, J.
2009 *Relativism and Monadic Truth*, Oxford University Press;
- Cohen, S.
1999 Contextualism, Skepticism, and the Structure of Reasons. *Philosophical Perspectives*, 13: 57-89;
- DeRose, K.
1992 Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52: 913-929;

- 2002 Assertion, Knowledge, and Context. *The Philosophical Review*, 111: 167-203;
- 2004 The Problem with Subject-Sensitive Invariantism. *Philosophy and Phenomenological Research*, 68: 346-350;
- 2006 'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments against Contextualism. *Philosophy and Phenomenological Research*, 73: 316-338;
- Dietz, R.
2008 Epistemic Modals and Correct Disagreement, in M. Garcia-Carpintero and M. Kölbel (Eds.), *Relative Truth*, Oxford University Press: 239-262;
- Hawthorne, J.
2004 *Knowledge and Lotteries*, Clarendon Press;
- Kaplan, D.
1989 Demonstratives, in J. Almog, J. Perry and H. Wettstein (Eds.), *Themes from Kaplan*, Oxford University Press: 481-563;
- King, J.
2003 Tense, Modality and Semantic Values. *Philosophical Perspectives*, 17: 195-245;
- Kompa, N.
2002 The Context Sensitivity of Knowledge Ascriptions. *Grazer Philosophische Studien*, 64: 1-18;
- 2005 The Semantics of Knowledge Attributions. *Acta Analytica*, 20: 16-28;
- Kölbel, M.
2004 Indexical Relativism versus Genuine Relativism. *International Journal of Philosophical Studies*, 12: 297-313;
- Lasersohn, P.
2005 Context Dependence, Disagreement, and Predicates of Personal Taste. *Linguistics and Philosophy*, 28: 643-686;
- Lewis, David
1998 Index, Context, and Content, in D. Lewis, *Papers in Philosophical Logic*, Cambridge University Press: 21-44;
- Ludlow, P.
2005 Contextualism and the New Linguistic Turn in Epistemology, in Preyer and Peter (2005): 11-50;
- MacFarlane, J.
2005a The Assessment-Sensitivity of Knowledge Attributions, in T. Szabo Gendler and J. Hawthorne (Eds.), *Oxford Studies in Epistemology 1*, Oxford University Press: 197-233;
- 2005b Knowledge Laundering: Testimony and Sensitive Invariantism. *Analysis*, 65: 132-138;
- 2009 Non-Indexical Contextualism. *Synthese*, 166: 231-250;

- forthcoming Epistemic Modals are Assessment-Sensitive”, retrievable at <http://johnmacfarlane.net/epistmod.pdf>;
- Partee, B.
2004 Comments on Jason Stanley’s “On the Linguistic Basis for Contextualism”. *Philosophical Studies*, 119: 147-159;
- Preyer, G. and Peter, G. (Eds.)
2005 *Contextualism in Philosophy. Knowledge, Meaning, and Truth*, Clarendon Press;
- Recanati, F.
2007 *Perspectival Thought. A Plea for Moderate Relativism*, Oxford University Press;
- Richard, M.
2004 Contextualism and Relativism. *Philosophical Studies*, 119: 215-242;
- Schaffer, J.
2005 What Shifts? Thresholds, Standards, or Alternatives, in Preyer and Peter (2005): 115-130;
- Schaffer, J.
2006 The Irrelevance of the Subject: Against Subject-Sensitive Invariantism. *Philosophical Studies*, 127: 87-107;
- Schiffer, S.
1996 Contextualist Solutions to Skepticism. *Proceedings of the Aristotelian Society*, 96: 317-333;
- Stanley, J.
2004 On the Linguistic Basis for Contextualism. *Philosophical Studies*, 119: 119-146;
2005 *Knowledge and Practical Interests*, Clarendon Press;
- Williamson, T.
2005 Knowledge, Context, and the Agent’s Point of View, in Preyer and Peter (2005): 91-114.
- Wright, C.
2007 New Age Relativism and Epistemic Possibility: The Question of Evidence. *Philosophical Issues*, 17: 262-283.