To appear in *The Philosophical Quarterly*

Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge

Timothy Williamson

Abstract

§I schematises the evidence for an understanding of 'know' and other terms of epistemic appraisal that embodies contextualism or subject-sensitive invariantism, and distinguishes between those two approaches. §II argues that although the cases for contextualism and sensitive invariantism rely on a principle of charity in the interpretation of epistemic claims, neither approach satisfies charity fully, since both attribute metalinguistic errors to speakers. §III provides an equally charitable anti-sceptical insensitive invariantist explanation of much of the same evidence as the result of psychological bias caused by salience effects. §IV suggests that the explanation appears to have implausible consequences about practical reasoning, but also that applications of contextualism or sensitive invariantism to the problem of scepticism have such consequences. §V argues that the inevitable difference between appropriateness and knowledge of appropriateness in practical reasoning, closely related to the difference between knowledge and knowledge of knowledge, explains the apparent implausibility.

A much-discussed idea in current epistemology is that some important features of the use of epistemic terms such as 'know' are best explained on the hypothesis of a shift in the standards for their correct application. Many versions of this idea are grouped together as 'contextualism'. Recently, a subtly different claim of standard-shifting has been distinguished from contextualism under the title 'subject-sensitive invariantism' (henceforth, simply 'sensitive invariantism'). Both contextualists and sensitive invariantists typically hold that mundane changes in how much is practically at stake or in what is conversationally relevant can raise or lower epistemic standards, and that the same underlying mechanisms also help to explain the pull of sceptical arguments, because sceptics shift the discourse to extraordinarily high standards by which they speak truly. This paper aims to demonstrate some explanatory resources of insensitive invariantism, the view on which the epistemic terms at issue undergo no shift in the standard for their correct application. If so, the standard motivations for contextualism and sensitive invariantism fail.

I

We may begin by setting out the evidence for shifting standards as an abstract schema. We fix on the word 'know', which is central to the debate; similar arguments apply to other terms of epistemic appraisal, such as 'justified', and to words in languages beyond English. We consider examples with three crucial features.

First, in some context C, it is reasonable and natural for a well-informed speaker to assert the term 'know' of some thinking subject S with respect to some true proposition *p*, time *t* and possible world *w*.[1] Here *t* and *w* are asserted to be the time and world of the knowing; they need not be the time and world of the asserting, nor need they be asserted to be part of the content of the knowledge. To assert 'know' of S with respect to *p*, *t* and *w*, the speaker in C utters a sentence K[S,*p,t,w*]. No special assumption is made here about the structure of K[S,*p,t,w*], except that it includes the word 'know'. It might have the explicit form 'S knows at *t* in *w* that P'. More usually, the speaker allows dimensions of the context of utterance to fix some of the parameters. For example, in a first-person present-tense case, S is the speaker of C, *t* is the time of C and *w* is the world of C, so K[S,*p,t,w*] may just have the form 'I know that P'. Let us provisionally assume that the speaker is right; 'know' is indeed true in C of S with respect to *p, t* and *w*. Equivalently:

(1) K[S,*p*,*t*,*w*] is true in C.

The second feature of the example is that in some other context C*, it is reasonable and natural for a well-informed speaker to deny the term 'know' of some thinking subject S* with respect to some true proposition *p**, time *t** and possible world *w**. We make no assumption at this point as to whether S*, *p**, *t** and *w** are distinct from S, *p*, *t* and *w* respectively. To deny 'know' of S* with respect to *p**, *t** and *w**, the speaker in C* utters a sentence ~K*[S*,*p**,*t**,*w**]. No special assumptions are made about the structure of ~K*[S*,*p**,*t**,*w**], except that it includes the word 'know'; it can differ from that of the negation of K[S,*p*,*t*,*w*]. ~K*[S*,*p**,*t**,*w**] might have the explicit form 'S* does not know at *t** in *w** that P*', or the speaker may allow dimensions of the context of utterance to fix some of the parameters. In the first-person present-tense case, ~K*[S*,*p**,*t**,*w**] may just have the form 'I do not know that P*'. Let us provisionally assume that the speaker is right; 'know' is indeed false in C* of S* with respect to *p**, *t** and *w**. Equivalently:

(1*) ~K*[S*,*p**,*t**,*w**] is true in C*.

The third feature of the case is that, by traditional standards, S* is epistemically as well-positioned at *t** in *w** with respect to *p** as S is at *t* in *w* with respect to *p*. For example, by those standards the probability of *p* on the evidence that S has at *t* in *w* is no higher than the probability of *p** on the evidence that S* has at *t** in *w**, the grounds on which S at *t* in *w* believes *p* are no better than those on S* at *t** in *w** believes *p**, S is no more confident at *t* in *w* of *p* than S* is at *t** in *w** of *p**, and so on. Perhaps matters of life and death for S* turn on *p**, while only trivial matters for S turn on *p*, or perhaps S* is obsessing in an epistemology seminar about sceptical scenarios while S is getting on with everyday life without a thought for scepticism, but according to traditional epistemology such differences make it no easier for S to know *p* than for S* to know *p**, given (as by hypothesis it is) that S* believes *p** just as S believes *p*. Thus, according to traditional epistemology:

(2) If S knows *p* at *t* in *w*, then S* knows *p** at *t** in *w**.

Are (1), (1*) and (2) mutually consistent?

Traditional epistemology ignores any indexicality in the word 'know'. It treats 'know' as having the

same reference, and making the same contribution to what proposition is expressed by sentences in which it occurs, irrespective of the context in which the utterance is made. On this view, the reference to C in (1) and to C* in (1*) is redundant, for in those contexts 'know' has the same reference as it has in the context in which we are now theorizing. In that case, truth behaves disquotationally with respect to 'know':

(3) K[S,*p*,*t*,*w*] is true in C iff S knows *p* at *t* in *w*.

(3*) ~K*[S*,*p**,*t**,*w**] is true in C* iff S* does not know *p** at *t** in *w**.

Given (3), (3*) and elementary logic, we can move from (1) and (1*) to the following, respectively:

(4) S knows *p* at *t* in *w*.

(4*) S* does not know *p** at *t** in *w**.

But (4) and (4*) form an inconsistent trio with the conditional (2): (4) is its antecedent and (4*) is the negation of its consequent.

The contextualist response is to deny at least one of (3) and (3*), on the grounds that 'know' is an indexical term. Of course, since it occurs outside quotation marks in (3) and (3*), contextualists will hold that those sentences too express different propositions as uttered in different contexts; which of them expresses a falsehood depends on the context in which the theorist is operating. Nevertheless, contextualists will insist, the reference of 'know' differs between C and C*, and therefore must differ either between the present context and C or between the present context and C* (or both). Given the indexicality of 'I', it is obviously false for me to assert the following biconditional: 'I am Madonna' is true in a context in which Madonna is the speaker iff I am Madonna. Similarly, (3) fails if C differs relevantly from the present context and (3*) fails if C* differs relevantly from the present context. In particular, if 'know' refers in C to the relation of K'ing, and in C* to the relation of K*'ing, then K[S,*p*,*t*,*w*] is true in C iff S K's *p* at *t* in *w*, while ~K*[S*,*p**,*t**,*w**] is true in C* iff S* does not K* *p** at *t** in *w**. But if K*'ing is sufficient but unnecessary for K'ing ('epistemic standards are higher in C* than in C'), then it is compatible with (2) and its grounds that S K's *p* at *t* in *w* and S* K's *p** at *t** in *w**,

while S does not K* *p* at *t* in *w* and S* does not K* *p** at *t** in *w**. Consequently, (1) and (1*) are jointly compatible with (2) and its grounds.

Contextualists give different accounts of the underlying mechanisms which, they say, generate the contextual shifts in the reference of 'know'. Naturally, those mechanisms should engage with non-epistemic differences between the contexts C and C*: for example, in how much is practically at stake, or in what error possibilities are conversationally salient. This paper abstracts from the details of specific contextualist accounts: it concerns the general structure of the argument.[2]

For contextualists, what propositions the antecedent and consequent of (2) express also varies with context, since 'know' occurs in them outside quotation marks: but (2) may still express a true proposition in any context, because whenever 'know' expresses a relation of K**'ing in a context C**, if S* is as well-positioned epistemically at *t** in *w** with respect to *p** as S is at *t* in *w* with respect to *p*, then S K**'s *p* at *t* in *w* only if S* K**'s *p** at *t** in *w**. Roughly: 'know' always refers to a purely *epistemic* standard, although which epistemic standard it refers to varies contextually. Thus contextualism reconciles (1), (1*) and (2).

Contextualists use such cases to argue for contextualism. For (1), (1*) and (2) are all highly plausible. Both the speaker's assertion endorsed by (1) and the speaker's denial endorsed by (1*) seem utterly natural, integral to the ordinary use of 'know', while (2) is far from distinctive of contextualism: traditional epistemology treats it as obvious. The conjunction of (1), (1*) and (2) with (3) and (3*) is inconsistent (given uncontentious background assumptions). Thus (1), (1*) and (2) together entail the disjunction of the negations of (3) and (3*). But the negation of (3) entails contextualism (the indexicality of 'know'); so does the negation of (3*). Thus (1), (1*) and (2) together entail contextualism.

Only recently has contextualism been clearly distinguished from *sensitive invariantism*, an anti-contextualist approach that appeals to a different kind of shift in epistemic standards.[3] Like contextualists, sensitive invariantists accept (1) and (1*); they vindicate the claims at issue made by speakers in C and C*. Unlike contextualists, sensitive invariantists treat 'know' as non-indexical, *invariant* in reference across contexts; so they accept both (3) and (3*). Since the conjunction of (1), (1*), (2), (3) and (3*) is inconsistent (given uncontentious background assumptions), sensitive invariantists deny (2); they assert both (4) and (4*). For them, what epistemic standard the subject must meet in order to know is *sensitive* to non-epistemic features of the subject's circumstances. For example, suppose that *p** makes more practical difference to S* at *t** in *w** than *p* makes to S at *t* in *w*, or more error possibilities are socially salient to S* at *t** in *w** than to S at *t* in *w*;

then, under sensitive invariantism, S* may have to meet a higher epistemic standard at *t** in *w** in order to know *p** than S has to meet at *t* in *w* in order to know *p*. Consequently, even though S* is as well-positioned epistemically at *t** in *w** with respect to *p** as S is at *t* in *w* with respect to *p*, S may know *p* at *t* in *w* while S* does not know *p** at *t** in *w**. Whereas contextualism exploits differences in the situation of the *speaker* who applies the word 'know', between the contexts C and C*, sensitive invariantism exploits differences in the situation of the *subject* to whom the word 'know' is applied, between S at *t* in *w* and S* at *t** in *w**.[4] Of course, everyone agrees that the situation of the subject is relevant to the correct application of 'know', for that includes circumstances of the usual epistemic kind, such as the validity of S's reasoning or the reliability of S's perceptual faculties. What distinguishes sensitive invariantism from more traditional conceptions of knowledge is the role that it assigns to circumstances of the subject not of the usual epistemic kind.

'Sensitive invariantism' is a less precise term than 'contextualism'; it is defined in vague epistemic terms such as 'epistemically well-positioned'. Unlike contextualism, sensitive invariantism does not assign 'know' to a well-recognized general semantic category (like that of indexicals). Even the notion of an epistemic standard is obscure. According to sensitive invariantism, S* at *t** in *w** meets the same epistemic standards with respect to *p** as S at *t* in *w* meets with respect to *p*, yet S at *t* in *w* knows *p* while S* at *t** in *w** does not know *p**. Thus knowing is not itself classified as an epistemic standard. Why not? Any sensitive invariantist claim to have a general concept of an epistemic standard that is independent of the concept of knowledge should be treated with caution, for concepts usually linked to the concept of an epistemic standard, such as those of evidence, justification and epistemic probability, themselves arguably depend on the concept of knowledge (Williamson 2000).

A particular unclarity in the notion of an epistemic standard concerns the case in which *p** follows from *p* but not *vice versa*. For example, let *p* be the proposition that the animal in the cage is a zebra, *p** the proposition that the animal in the cage is not a painted mule, and S*, *t** and *w** be just S, *t* and *w* respectively. Suppose that S believes at *t* that the animal in the cage is not a painted mule on the basis of competent deduction from the premise that the animal in the cage is a zebra. Consider a theory on which S knows at *t* that the animal in the cage is a zebra but does not know at *t* that the animal in the cage is not a painted mule, because (to a first approximation) if the animal in the cage had not been a zebra, S would not still have believed at *t* that the animal in the cage was a zebra, whereas if the animal in the cage had been a painted mule, S would still have believed at *t* that the animal in the cage was not a painted mule (Dretske 1970, Nozick 1980). Is the theory applying

different epistemic standards to determine whether S knows the two propositions? On any sensible probability distribution, the probability of *p** is at least as high as the probability of *p*, which might suggest that a higher epistemic standard is being applied to *p** than to *p*. On the other hand, satisfaction of the tracking counterfactual (that if the proposition had been false the subject would not have believed it) might itself be regarded as a fixed epistemic standard by which *p* does better than *p**. The question is evidently vague. In practice, recent sensitive invariantists typically deny such synchronic failures of the closure of knowledge under competent deduction, but that does not resolve the taxonomic issue. It would be too restrictive simply to exclude consideration of cases in which *p** is not *p*. Sensitive invariantists tend to allow that, at a time *t* before deducing that it is not a painted mule, S may know that the animal in the cage is a zebra, while at a later time *t** after making the deduction (and as a result of doing so), S still believes but fails to know that the animal in the cage is a zebra. Arguably, if the beliefs at issue are present-tensed, then in effect the belief at *t* is the belief that the animal in the cage at *t* is a zebra, which differs in content from the belief at *t** that the animal in the cage at *t** is a zebra.

Fortunately for sensitive invariantists, the difficulties in providing a general definition of 'sensitive invariantism' do not refute the strategy of denying the conditional (2) in a particular case. We can still use the expression 'sensitive invariantism' at least as a family resemblance term for a vaguely specified cluster of recognizably similar approaches.

Sensitive invariantism can to some extent mimic the successes of contextualism by transferring what the contextualist says about differences in the situation of the speaker to differences in the situation of the subject. Indeed, no transfer is needed when the situation of the speaker *is* the situation of the subject, as it is in first-person present-tense 'know'-ascriptions. When one says something of the form 'I know that P', the speaker is the subject, and the time and world of the context are those with respect to which the ascription is made. Such ascriptions are of course central to much traditional epistemological debate; they cannot separate contextualism from sensitive invariantism.

Nevertheless, the parallelism between contextualism and sensitive invariantism is imperfect. For let S*, *p**, *t** and *w** be S, *p*, *t* and *w* respectively. In this special case, (2) is a tautology; sensitive invariantists cannot deny (2). Here (1) and (1*) alone jointly entail contextualism, via the disjunction of the negations of (3) and (3*), independently of considerations about the subject's epistemic position. Indeed, this case provides the clearest and simplest argument for contextualism: speakers in different contexts make verbally inconsistent epistemic assessments of the same subject with respect to the same time, the same world and the same

proposition (DeRose 2004). Perhaps more is practically at stake in C* than in C, for example in whether the speaker should trust S's testimony *p*; perhaps S's fallibility has more conversational relevance in C* than in C. The difference in the contexts does all the work; there is no difference in the subject's situation to which sensitive invariantists might appeal. Here they cannot mimic the contextualist strategy. Instead, they must reject at least one of (1) and (1*), and argue that the speaker in C or C* is not speaking truly. However, this paper aims not to adjudicate the dispute between contextualism and sensitive invariantism but to question the motivation for both views.

II

The common basis of contextualism and sensitive invariantism comprises pairs of claims such as (1) and (1*), which in effect endorse the dispositions of speakers in various contexts to assert or deny 'know' of various cases. Presumably, the endorsement rests on a methodological principle of charity, by which, very roughly, we should prefer to interpret speakers as speaking and thinking truly rather than falsely (*ceteris paribus*). Shifting standards seems to give us more flexibility to assign 'know' a charitable reference. We saw at the end of section I that sensitive invariantists' anti-contextualist commitments limit their charity, by making (1) and (1*) mutually inconsistent in some cases. The present section suggests that the limitations on charity for both contextualists and sensitive invariantists are far more extensive than that: given reasonable assumptions, any case of the kind that motivates claims of shifting standards is also liable to manifest cross-contextual patterns of dispositions to make mutually irreconcilable claims. If so, any reasonable view will involve the attribution to speakers of some systematic errors. Of course, that does not refute the principle of charity, nor does it prove that contextualism or sensitive invariantism is not the most charitable account. But it does show that an insensitive invariantist account that denies (1) or (1*) should not be dismissed as uncharitable on those grounds alone.[5]

Consider a familiar dialogue at the zoo:

John: I know that this is a zebra.

Mary: How do you know that it isn't a mule cleverly painted to look like a zebra?

John: Hmm, for all I know it *is* a painted mule. So I was wrong. I don't know that it is a zebra after

all.

The sensitive invariantist tries to reconcile John's first and last remarks by postulating that Mary's question caused John to cease to know that it is a zebra, although that move would not work if John's last sentence were the past-tense 'I didn't know that it was a zebra after all'. Whether John uses the past or present-tense, a contextualist can reconcile his first and last remarks by postulating that Mary's question causes the word 'know' to shift its reference. But neither move makes sense of John's admission 'So I was wrong', for both involve explaining how John was not wrong in his first speech. More subtly, in context his final words 'after all' imply an admission of error. A bomb disposal expert might report 'The bomb is armed' and then, after disarming it, 'The bomb is not armed', but to say 'The bomb is not armed after all' in the latter circumstances would be to invite the interpretation that one was admitting an error in the original claim 'The bomb is armed'. If John claims that his earlier claim was false, not all his claims are true.

Contextualists and sensitive invariantists might object that the inclusion of the words 'I was wrong' and 'after all' in the dialogue was tendentious. But in fact those words seem entirely natural there. In adjusting their epistemic claims to previously unnoticed sceptical scenarios as they are brought to their attention, speakers often feel caught out, corrected. Similarly, when the epistemology seminar is over and epistemic standards drop, the sceptical claims made in the seminar are liable to seem ridiculous and false, not pedantic but true. In general, speakers in the high standards context C* tend to regard speakers in the low standards context C as not taking the error possibilities seriously enough, while speakers in C tend to regard speakers in C* as taking the error possibilities too seriously. Imagine Lo in the local bar, notorious for its low standards, saying 'John knows that it is a zebra' and Hi in the zoology seminar, famous for its high standards, saying 'John does not know that it is a zebra'. Hi looks down on Lo, whose remark she regards as an ignorant underestimate of the problems of animal identification, disrespectful of the required specialist expertise; Lo resents Hi, whose remark he regards as a snobbish piece of professional empire-building, disrespectful of lay experience. Of course, were Hi and Lo contextualists themselves, they might think differently: but it is an occupational hazard for those defending a philosophical theory to project it onto ordinary speakers of the language. Ordinary speakers who have not consciously signed up to contextualism often have attitudes to what others say that contextualism makes inappropriate.

Schematically, in effect, the speaker in C* denies (1) and the speaker in C denies (1*). Let us pretend

that the speaker in C* utters the contradictory of (1), and the speaker in C utters the contradictory of (1*). Realistically, of course, they will say something much less formal, such as 'He's wrong', but the pretence is harmless and clarifies the structure of the argument. On pain of inconsistency, the theorist who asserts (1) and (1*), in charity to the first-order utterances of the speakers in C and C* respectively, cannot also assert that the contradictories of (1) and (1*) are true, in charity to the second-order utterances of the speakers in C* and C respectively. Endorsing all the claims speakers are disposed to make in C and C* lands us in a double contradiction, independently of (3) and (3*), denied by contextualists, and of (2), denied by sensitive invariantists.

One response for contextualists is to go contextualist about 'true' as well as 'know'. They endorse speakers' claims quotationally:

(5) '~K*[S*,*p**,*t**,*w**] is not true in C*' is true in C.

(5*) 'K[S,*p*,*t*,*w*] is not true in C*'* is true in C*.

In order to get an inconsistency from the combination of (5) and (5*) with (1) and (1*), one needs disquotational principles analogous to (3) and (3*) (all in the theorist's context):

(6) '~K*[S*,*p**,*t**,*w**] is not true in C*' is true in C iff ~K*[S*,*p**,*t**,*w**] is not true in C*.

(6*) 'K[S,*p*,*t*,*w*] is not true in C*'* is true in C* iff K[S,*p*,*t*,*w*] is not true in C.

But the contextualist about 'true' rejects (6) and (6*) on the grounds that the word 'true' differs in reference between C or C* and the context in which the theorist advances (5) and (5*). Without something like (6) or (6*), the contextualist can consistently combine (5) and (5*) with (1) and (1*).

One trouble with this generalized contextualist move is that it is utterly indiscriminate. Faced with any apparent disagreement, no matter how blatant, between speakers in different contexts, one could effect such a 'reconciliation'. However angrily each party denies that what the other says is 'true', the generalized contextualist could insist that they are all speaking truly, given the variable reference of their words (including

'true' and 'false') across the contexts in which they are being used. That looks suspiciously like undergraduate relativism ('whatever a person sincerely believes is true (for them)'). Charity was never intended to do so much. We may suspect that the purported 'reconciliation' was achieved at the price of trivializing the relevant uses of 'true', and that no reasonable principle of charity should be satisfied by an interpretation on which everything that speakers say and think is trivially true. While these sketchy remarks are no substitute for a proper discussion of the generalized contextualist position, they will do for present purposes because most contextualists about 'know' would deny that their contextualism about 'know' commits them to such contextualism about 'true'. They intend a far more limited doctrine. For example, they do not regard their dispute with anti-contextualists as merely apparent.

Contextualists about 'know' and sensitive invariantists might respond differently to (5) and (5*). They might argue that (1) and (1*) take priority over (5) and (5*), on the grounds that charity should be applied to non-metalinguistic utterances before it is applied to metalinguistic utterances. If (1) and (1*) have more weight than (5) and (5*), then the original argument for contextualism or sensitive invariantism still has force. Is this marginalization of the metalinguistic *ad hoc*? We expect speakers to be more reliable over simple claims than over complex ones. However, some metalinguistic utterances play a key role in the interpretation even of non-metalinguistic expressions. For suppose that in a context C0 speakers apply a non-metalinguistic expression E to an object *x* in straightforward ignorance of some highly relevant facts about *x*; thus the application of E to *x* in C0 may have been incorrect (false). The interpreter will hope to discover such errors by observing speakers correct themselves, withdraw their earlier application of E in C0 to *x*, when they learn the relevant facts, or by observing better-informed speakers correct worse-informed speakers, criticize their application of E in C0 to *x*. Otherwise, the charitable interpreter would wrongly count the application of E in C0 to *x* as correct (true). To avoid such misinterpretations, one must give weight to metalinguistic utterances. It is insufficient merely to observe speakers in a later context C1 denying E of *x*, for that is compatible with the hypothesis that E is an indexical term, true of *x* in C0 but false of *x* in C1 (a possibility on which contextualism relies).[6] Similarly, the interpreter needs to use metalinguistic utterances to identify and filter out cases of linguistic incompetence. Indeed, metalinguistic responses to utterances in which an expression occurs constitute an important test of its indexicality. Thus the methodology of interpretation tells against the marginalization of metalinguistic utterances such as the denial in C of (1*) and the denial in C* of (1).

When S*, *t**, *w** and *p** are S, *t, w* and *p* respectively (the case most relevant to distinguishing

contextualism from sensitive invariantism), there is a salient explanation why the speaker in C denies (1*). For suppose that the speaker believe, truly or falsely, that 'know' is non-indexical. Consider the following argument as uttered in C, where the sentence names 'K[S,*p,t,w*]' and '~K[S,*p,t,w*]' are to be replaced by the sentence K[S,*p,t,w*] itself and its contradictory respectively (but '~K*[S*,*p**,*t**,*w**]' is left as a sentence name):

~K*[S*,*p**,*t**,*w**] is true in C* iff ~K[S,*p,t,w*].

K[S,*p,t,w*].

Therefore: ~K*[S*,*p**,*t**,*w**] is not true in C*.

The argument is valid. The first premise is true if 'know' is non-indexical (the differences between ~K*[S*,*p**,*t**,*w**] and ~K[S,*p,t,w*] take care of any indexicality in constituents of the former sentence distinct from 'know'). The speaker in C accepts the second premise. Hence the speaker in C who believes that 'know' is non-indexical is predicted to accept the conclusion, and so to deny (1*).

Similarly, still given that S*, *t**, *w** and *p** are S, *t, w* and *p* respectively, there is a salient explanation why the speaker in C* denies (1). For suppose that the speaker believe, truly or falsely, that 'know' is non-indexical. Consider the following argument as uttered in C*, where the sentence names '~K*[S*,*p**,*t**,*w**]' and 'K*[S*,*p**,*t**,*w**]' are to be replaced by the sentence ~K*[S*,*p**,*t**,*w**] itself and its contradictory respectively (but 'K[S,*p,t,w*]' is left as a sentence name):

K[S,*p,t,w*] is true in C iff K*[S*,*p**,*t**,*w**].

~K*[S*,*p**,*t**,*w**].

Therefore: K[S,*p,t,w*] is not true in C.

The argument is valid. The first premise is true if 'know' is non-indexical (the differences between K[S,*p,t,w*] and K*[S*,*p**,*t**,*w**] take care of any indexicality in constituents of the former sentence distinct from 'know'). The speaker in C* accepts the second premise. Hence the speaker in C* who believes that 'know' is non-indexical is predicted to accept the conclusion, and so to deny (1).

Thus if speakers treat 'know' as non-indexical, in the very situations in which only the hypothesis of shifting standards is supposed to allow a charitable treatment of speech dispositions, speech dispositions do not

admit of fully charitable treatment. Other phenomena do indeed indicate that speakers treat 'know' non-indexically, for example in reported speech (Hawthorne 2003). If Lo in the low standards context says 'I know that P', Hi in the high standards context may report Lo by saying 'Lo says that he knows that P'. Given a platitudinous schema about truth, 'In saying that Q, one speaks truly iff Q', Hi can conclude 'Lo speaks truly iff he knows that P', which the contextualist will reject as uttered in the high standards context. Similarly, if Hi says 'I do not know that P*', Lo may report Hi by saying 'Hi says that she does not know that P*'. Given the platitudinous schema, Lo can conclude 'Hi speaks truly iff she does not know that P*', which the contextualist will reject as uttered in the low standards context. This version of the problem involves the complexities of indirect speech reports of propositional attitude ascriptions: we have just seen that the problem can be raised even without those complexities. While everyday indirect speech reports sometimes ride roughshod even over the indexicality of words that are generally agreed to be indexical, such as 'tall', speakers are much less prone with the latter terms than with 'know' to treat what are supposed to be contextual differences as disagreements.

Given that speakers treat 'know' as non-indexical in various respects, it does not follow that they are right to do so.[7] The position is rather that all theorists will be forced to postulate and explain systematic errors in our use of epistemic terms at some point or other. To make an adequately reasoned choice between the competing theories, one will need to compare their explanations of the postulated errors. The rest of this paper merely sketches the kind of explanation available to the insensitive invariantist.

III

Some insensitive invariantists are sceptics, for whom 'know' invariably refers to a maximal epistemic standard that we cannot meet. They attempt to explain the illusion of knowledge in various ways. This paper, however, focusses on anti-sceptical insensitive invariantism: 'know' invariably refers to an epistemic standard that we can and do meet quite easily; everyday ascriptions of knowledge are often true. Apparent cases of standard-shifting typically involve such everyday ascriptions in the low standards context. Let us therefore assume that S knows *p* at *t* in *w*: (4) is true. By insensitivity, (2) is also true, so S* knows *p** at *t** in *w**: (4*) is false. By semantic invariance, (3) and (3*) are true, so (1) is true and (1*) false. The assertion of knowledge in the low standards context C is true; the denial of knowledge in the high standards context C* is false. What the anti-sceptical

insensitive invariantist must explain is the illusion of ignorance. But is that so hard?

Whether someone knows something depends on a multitude of considerations, which rarely all point in the same direction. For example, the reliability of the process by which a belief was acquired and retained makes a difference to whether it constitutes knowledge: notoriously, a particular event of acquisition and retention instantiates many types of process of different degrees of specificity, which may consequently differ in degree of reliability too (Conee and Feldman 1998). Whether a given case of true belief is a case of knowledge depends on both the probability and the similarity of cases of false belief; those notions of probability and similarity are themselves vague. Analogous questions arise for other considerations on which knowing depends: for example, whether the subject's confidence is sufficient for belief. Evidently, we cannot judge whether someone knows something by feeding answers to all such questions into a formula. We have no choice but to judge on an impressionistic basis, perhaps using convenient but highly fallible rules of thumb, trying to remain receptive to anything relevant. In particular, we have no rule that tells us the exact relative weights of the considerations that support knowing and the considerations that support not knowing.

In these circumstances, it would be a miracle if we did *not* give more weight to considerations when they are psychologically salient to us than when they are not. Just that happens when we have to weigh complex conflicting considerations in reaching a difficult moral or practical decision: we may oscillate between two options, as we switch our attention between the considerations that favour one and the considerations that favour the other; it does not follow that what really is the morally or practically 'best' thing to do oscillates likewise.[8] But the high standards case is constructed in just such a way as to focus the ascriber's attention on considerations that tell against the ascription of knowledge, more specifically, on possibilities of error. They may be psychologically salient because the practical costs of error are high for the subject or the ascriber, or simply because they have been evoked in vivid and convincing detail. No wonder the ascriber gives them increased weight. One effect of fictional violence on television is to make viewers overestimate their chances of being the victims of violent crime: they suffer an illusion of danger. Might not an illusion of epistemic danger result from exposure to lurid stories about brains in vats, evil demons, painted mules, or gamblers who bet the farm?[9]

We cannot filter out the psychological bias effect by constructing cases in which more is at stake for ascribers than they realize. For if we judge the sentence K*[S*,*p**,*t**,*w**] false as uttered in C*, because the falsity of *p** would be disastrous for the speaker in C* even though she does not realize that it would be

disastrous, our judgements as theorists about K*[S*,*p**,*t**,*w**] is still coloured by *our* awareness of the danger to the speaker.

Arguments for shifting standards typically fail to control for psychological bias effects. The inconsistencies in our use of 'know' in the pairs of cases to which such arguments appeal (see section II) give us good reason to take the possibility of such effects seriously. Surely they do occur. Nevertheless, we shall see that, by themselves, they do not provide a fully satisfying treatment of the cases.

IV

There is an intuitively attractive link between knowledge and practical reasoning. If it is not true for me to say 'I know that my door is locked', something is wrong with practical reasoning in which I rely on the premise that my door is locked, for in some sense I am not entitled to that premise, however much it seems to me that I am. If I have good evidence that 'I know that my door is locked' is true, although it is in fact false for reasons beyond my ken, then I have good evidence that I am entitled to the premise that my door is locked, and therefore a good *excuse* for relying on it, but a good excuse for doing something is not an entitlement to do it. I may have good evidence that I am entitled to take books out of the library, and therefore a good excuse for taking them out, even though I am not in fact entitled to do so, because the rules have recently been changed. Conversely, if it is true for me to say 'I know that my door is locked', so far nothing is wrong with practical reasoning in which I rely on the premise that my door is locked, for I am entitled to that premise. In other words, 'I know that my door is locked' is true in my context iff that my door is locked is an *appropriate premise* for my practical reasoning. More generally:

(KPR) A first-person present-tense ascription of 'know' with respect to a proposition is true in a context iff that proposition is an appropriate premise for practical reasoning in that context.[10]

Although the notion of an appropriate premise in practical reasoning is imprecise, we seem to have enough grip on the difference between sensible and silly practical reasoning to use KPR as an effective constraint on 'knowledge' ascriptions. In particular, we can use KPR to reinforce the original argument for

shifting standards in some crucial cases. Consider a case as in section I with these special features: *p* makes little practical difference to the speaker in context C (given their purposes); *p** makes an enormous practical difference to the speaker in context C* (given their purposes); S, *t* and *w* and S*, *t** and *w** are the speaker, time and world of C and C* respectively; K[S,*p,t,w*] expresses a first-person present-tense assertion in C of 'knowledge' of *p*; ~K*[S*,*p*,t*,w**] expresses a first-person present-tense denial in C* of 'knowledge' of *p**. Such cases are not hard to construct. We therefore have these applications of KPR:

(7) K[S,*p,t,w*] is true in C iff *p* is an appropriate premise for practical reasoning in C.

(7*) ~K*[S*,*p*,t*,w**] is true in C* iff *p** is not an appropriate premise for practical reasoning in C*.[11]

Since *p** makes much more practical difference in C* than *p* makes in C, even though the speaker in C* is epistemically as well-positioned with respect to *p** as the speaker in C is with respect to *p*, both these claims are plausible in suitably constructed examples:

(8) *p* is an appropriate premise for practical reasoning in C.

(8*) *p** is not an appropriate premise for practical reasoning in C*.

Pretheoretically: it is too risky for the speaker in C* to use *p** as a premise in practical reasoning and not too risky for the speaker in C so to use *p*. But (7) and (8) jointly entail (1), while (7*) and (8*) jointly entail (1*). Thus, using KPR, we can revive the original argument for shifting standards, while replacing the original reliance on bare linguistic intuitions about the applicability of 'know' for (1) and (1*) by appeal to a general epistemic principle about practical reasoning and our ability to discriminate in particular cases between silly and sensible decision-making. The insensitive invariantist must reject at least one of KPR, (8) and (8*).

The insensitive invariantist cannot plausibly deny KPR without further explanation; its appeal is too great. Moreover, without KPR the concept of knowledge would lose some of its significance: one reason why it matters whether you know something is that, if you do, you are entitled to use it in ways in which you would not otherwise be so entitled. Once KPR is granted, it is hard to see how there could not be true pairs (8) and (8*), for

surely what it is appropriate to use as a premise for practical reasoning at least partly depends on how much is at stake, not just on purely epistemic matters of the traditional kind.

Although appealing, KPR is not self-evidently correct. If it fails, so does the envisaged strengthening of the argument for shifting standards, and the insensitive invariantist's task is correspondingly easier. More generally, the more the practical is separated from the epistemic, the easier it is to reply to arguments from practical differences to shifting semantic standards for epistemic terms. A full evaluation of KPR will not be undertaken here. Rather, the aim is to assess the prospects for insensitive invariantism granted KPR.

Considerations about practical reasoning do not favour all the applications of contextualism and sensitive invariantism that have been made; KPR is a two-edged sword. Most contextualists and sensitive invariantists hope to explain the pull of scepticism by arguing that sceptical discourse creates contexts in which sceptics speak truly. Thus, in an epistemology seminar with an obsessive focus on radical sceptical scenarios, one cannot truly say 'I know I have hands'; one cannot even truly say 'I know it is more probable than not that I have hands', for the sceptic argues that my evidence is simply the appearance that I have hands, which, being equally consistent with the hypothesis that I have hands and the hypothesis that I lack them, according to the sceptic gives no support to one hypothesis over the other. Imagine that, in the middle of the epistemology seminar, a rich friend offers you a bet on which you gain £5 if you have hands and lose £50 otherwise. If reference in this context suits the sceptic, then by KPR you have no appropriate premise to use for sensible practical reasoning to the conclusion that you should accept the bet, for such a premise would imply that it is more probable than not that you have hands, since otherwise it would not rationalize accepting a bet on which the potential loss is much larger than the potential gain. Nevertheless, despite all the talk of brains in vats, it is sensible to accept the bet, not to walk away from £5. Of course, scepticism makes it unclear how the bet could be enforced. Indeed, the sceptic will ask how you know that you have been offered a bet, and if so by whom − perhaps it was by an evil demon rather than your rich friend. Taken together, such doubts might constitute a reason for refusing the bet. But since it *is* sensible to accept the bet, that would be just further evidence against the truth of scepticism even in the seminar. Might proposing the bet lower epistemic standards? That suggestion looks ominously *ad hoc*. If the sceptical possibilities are as serious as mundane possibilities in the seminar, why should merely proposing a bet exclude the former but not the latter? If the sceptical possibilities are not as serious as mundane possibilities in the seminar, did the seminar effect a genuine shift in truth-conditions? The appeal to sensible practical reasoning undermines the shifting-standards explanation of the pull of scepticism.

Nevertheless, the insensitive invariantist still has to confront the argument from sensible practical reasoning to shifting standards in cases of contrasting stakes.

V

Let us temporarily bracket KPR and consider appropriateness in practical reasoning more abstractly.

Suppose that an agent relies on the premise *q* in otherwise impeccable practical reasoning, and (in her context) *q* is indeed appropriate, but she is in no position to know that *q* is appropriate. Then our assessment of her reasoning should be ambivalent: it is good that she used an appropriate premise, bad that she was in no position to know that it was appropriate. No simple verdict captures the complexity of the position.

It is tempting to object that such a position cannot arise, because it is the nature of appropriateness as a premise for practical reasoning to be epistemically accessible to the agent. For she must reason from what is accessible to her and cannot be condemned for not taking account of anything else. On the objector's view, whenever *q* is appropriate, the agent is in a position to know that *q* is appropriate: in the terminology of Williamson (2000), that *q* is an appropriate premise for practical reasoning is a *luminous condition*. But it is argued there that only trivial conditions are luminous, for instance those that obtain in all cases or in none. That result will be assumed henceforth, although to save space the argument will not be properly rehearsed here. The strategy is to construct a sorites series between a case in which the condition clearly obtains and one in which it clearly fails to obtain, and then to argue that such a series cannot exist for a luminous condition. Luminosity must fail close to the boundary between cases where the condition obtains and cases where it does not, just on the obtaining side. Now for most propositions *q*, the condition that *q* is an appropriate premise for practical reasoning is highly non-trivial. Take the condition that an appropriate premise is that my door is locked; it obtains in some cases and not in others; a continuum runs between the former and the latter of exactly the kind that the anti-luminosity argument requires. Thus, in some cases, the proposition that my door is locked is an appropriate premise for my practical reasoning even though I am in no position to know that it is appropriate. Think of *q*'s appropriateness as the agent's possession of authority to use *q* as a premise in practical reasoning. It is one thing to have the authority to do something, another to know that one has that authority, even when both having the authority and knowing that one has it matter.

How harshly should we judge practical reasoning in which the agent relies on an appropriate premise without being in a position to know that it is appropriate? A natural answer is: it depends on how much is at stake. If not much, then it seems unreasonably pedantic to condemn the reasoning. But if matters of life and death are at stake, the charge that the agent was not in a position to know that the premise was appropriate becomes more serious.

Now recall KPR as the epistemic standard of appropriateness. The insensitive invariantist can state it in simplified form:

(KPR+) One knows *q* iff *q* is an appropriate premise for one's practical reasoning.

Thus when *q* is appropriate but one is in no position to know that *q* is appropriate, one in effect knows *q* without being in a position to know that one knows *q*. The argument for the possibility of such cases is a special case of the anti-luminosity argument; the details are presented in (2000).

It would be pointless to respond by insisting that the proper condition for *q* to be an appropriate premise for one's practical reasoning is not that one knows *q* but that one knows that one knows *q*. For, by another application of the anti-luminosity argument, one can know that one knows *q* without being in a position to know that one knows that one knows *q*, so the modified version of KPR+ would be open to an 'objection' exactly analogous to the one that prompted the modification. For any natural number *n*, one can in principle have *n* iterations of knowledge of a proposition *q* without being in a position to know that one has *n* iterations of knowledge of *q*.[12] Although there is no purely logical objection to having infinitely many iterations of knowledge of a proposition *q* (that is, to having *n* iterations for every natural number *n*), it is doubtful that such cases arise in practice.

The insensitive invariantist can now start to explain why (8) and (8*) look jointly plausible, even though the agent in C* is epistemically as well-positioned with respect to *p** as the agent in C is with respect to *p*. For suppose that the agent in C knows *p* and the agent in C* knows *p**, but the agent in C is in no position to know that she knows *p* and the agent in C* is in no position to know that she knows *p**. Since stakes are higher in C* than in C, we as theorists may view the failure of second-order knowledge in C* more sternly than its failure in C, and therefore regard *p* as appropriate in C but *p** as inappropriate in C*. The agent herself may take the same view, if she is aware of the difference in stakes. Similar explanations are possible on the assumption

that the agent in C has *n* iterations of knowledge of *p* and the agent in C* has *n* iterations of knowledge of *p**, while the agent in C is in no position to know that she has *n* iterations of knowledge of *p* and the agent in C* is in no position to know that has *n* iterations of knowledge of *p**, for fixed *n*. How many iterations are relevant depends on how much is at stake.

The insensitive invariantist could try to build variation in the required number of iterations of knowledge into appropriateness itself, with a corresponding revision of KPR and KPR+: in some cases *q* would be appropriate iff one knew *q*, in others iff one knew that one knew *q*, and so on, depending on the stakes. If such a move is acceptable, it provides the insensitive invariantist with a systematic response to arguments from practical differences to shifting semantic standards for epistemic terms. However, the move does not guarantee the epistemic accessibility of appropriateness, for the anti-luminosity argument is quite general. Moreover, it unperspicuously mixes together considerations at different levels. In knowing that one knows *q*, one knows a truth about one's own epistemic state; in knowing *q* itself, one may know a truth simply about the external world. One might therefore prefer the attractively simple KPR+. Even under that assumption, we shall see that the insensitive invariantist has a systematic response to practical arguments for shifting standards.

How specifically does the general reasoning of the anti-luminosity argument apply in the cases to which contextualists and sensitive invariantists appeal? They argue that, after a casual glance at the timetable, 'I know that the plane stops in Chicago' is true as uttered by Lo, false as uttered by Hi. In setting up such examples, they emphasize the lack of further checks and the possibility of errors, while still making it clear that a denial of knowledge by the insensitive invariantist would require extensive revision of the everyday practice of knowledge ascription in a sceptical direction. Their examples are often envisaged as ones in which the subject counts as knowing by ordinary standards, but not by very much. For the anti-sceptical insensitive invariantist, these are cases close to the boundary between knowledge and ignorance, just on the knowledge side. But the anti-luminosity argument predicts that one will know without being in a position to know that one knows in exactly such cases. In such cases, both Lo and Hi know that the plane stops in Chicago ; neither Lo nor Hi is in a position to know that they know that the plane stops in Chicago.[13] Since stakes are higher for Hi than for Lo, the lack of second-order knowledge is more serious for Hi than for Lo. That the plane stops in Chicago is an appropriate premise for practical reasoning for both of them (given KPR+). However, Hi has far more reason than Lo has to check on such practical reasoning: to engage in second-order practical reasoning about whether to trust the first-order practical reasoning. Since Hi is in no position to know that the first-order premise that the

plane stops in Chicago is appropriate, the second-order premise that the first-order premise is appropriate is, although true, inappropriate (given KPR+). Thus second-order reasoning is in no position to give a clean bill of health to first-order reasoning based on the premise that the plane stops in Chicago. Although that applies to both Hi and Lo, Hi needs the bill of health more.

We can reformulate the point by using the idea that knowing *q* is the condition for having warrant to assert *q* (Williamson 2000). Both Hi and Lo have warrant to assert 'The plane stops in Chicago'; neither has warrant to assert 'I know that the plane stops in Chicago' or 'The other guy knows that the plane stops in Chicago' (each knows what checking the other has done). Correspondingly, neither has warrant to assert 'I have warrant to assert that the plane stops in Chicago' or 'The other guy has warrant to assert that the plane stops in Chicago'. But, given the temptation to assert that the plane stops in Chicago, the cautious question 'Have I warrant to make that assertion?', which neither has warrant to answer, is far more urgent for Hi.

Similar problems arise when more, but not infinitely more, iterations of knowledge are available. If stakes are high enough, prudent human agents will engage in third-order reasoning about whether to trust their second-order reasoning about whether to trust their first-order reasoning, and so on. Consider this dialogue (with oneself or another):

Q1: Is *q* the case?

A: Yes.

Q2: Did you have warrant for your answer to Q1?

A: Yes.

Q3: Did you have warrant for your answer to Q2?

A: I don't know.[14]

At any point in such an interrogation, anything less than a positive answer seems to destabilize the previous positive answer, and therefore all the earlier positive answers in a domino effect. For what use is an assertion if the speaker is unwilling to stand over it at the next level up? Yet each further question in effect demands a further iteration of knowledge. If the account in (2000) is at least approximately on the right lines, in virtually any case a human agent would sooner or later run out of iterations of knowledge, even if all the required higher-order beliefs were in place. This extension of the point to higher iterations of knowledge enables the insensitive

invariantist to handle cases that contextualists may offer in which by ordinary standards subjects not only know but know that they know, being far from cases in which they falsely believe that they know; the Chicago flight example could be reconstrued as such a case. For although they know that they know, they lack some higher iteration of knowledge.

Fortunately, in practice we often have enough iterations of knowledge to withstand the interrogations to which it is reasonable to respond, perhaps while uneasily aware that we could not withstand more extensive interrogations. For creatures who can know without being in a position to know that they know, or who can know that they know without being in a position to know that they know that they know, the capacity in principle to withstand all such interrogations is simply not a good test of their knowledge: as often, the rules of dialectic are a poor guide to epistemology. High-level failure does not legitimately unravel low-level success.

More has to be explained. Hi does not merely lack warrant to assert 'Lo knows that the plane stops in Chicago'. Hi seems positively to possess warrant to assert 'Lo does not know that the plane stops in Chicago' (although Lo lacks warrant to assert 'Hi does not know that the plane stops in Chicago'). But higher-order ignorance does not warrant denying first-order knowledge. For the insensitive invariantist, to have warrant to assert 'Lo does not know that the plane stops in Chicago' is to know that Lo does not know that the plane stops in Chicago; but nobody knows that, because Lo does know that the plane stops in Chicago. This is where the considerations of section III are needed.

When we as theorists contemplate Hi's position, we are struck by the disastrous consequences of believing falsely (in counterfactual circumstances) that the plane stops in Chicago; the same applies when Hi contemplates Hi's position, if Hi knows how much is at stake. That makes salient to us (perhaps including Hi) the weaknesses of the epistemic position that Hi and Lo share with respect to the proposition that the plane stops in Chicago. Consequently, when we consider from Hi's point of view whether Lo knows that the plane stops in Chicago, we give more weight to considerations that favour a negative answer. It therefore appears to us that Hi can truly assert 'Lo does not know that the plane stops in Chicago'. But the appearance is deceptive, the effect of psychological bias. Neither Lo nor Hi really has warrant to deny that either of them knows that the plane stops in Chicago. The bias infects our assessment of truth in Hi's context even if Hi herself is unaware of the high stakes.[15]

Insensitive invariantists can therefore use the bias effects to explain the systematic errors which they (like virtually all theorists) are forced to discern in our practice of knowledge attribution. Those errors are not

necessary for prudent decision-making. Hi can be cautious by not believing outright that Lo knows that the plane stops in Chicago; to believe outright that Lo does not know that the plane stops in Chicago is overkill. Insensitive invariantists can use the considerations about iterations of knowledge to explain the more functional features of apparent standard-shifting, which do contribute to habits of good decision-making. The two sorts of explanation complement but also reinforce each other: bias in our beliefs about which first-order beliefs constitute knowledge makes those second-order beliefs less reliable, thereby making it harder for them to constitute second-order knowledge.

The resultant anti-sceptical insensitive invariantist picture is complex: but the complexities are ones that we should have expected anyway for good independent reasons. Contextualism and sensitive invariantism seem to involve at least as many complexities; it is unclear how far they can explain those complexities on the basis of independently plausible considerations. Given the explanatory resources available to anti-sceptical insensitive invariantism, the case for contextualism or sensitive invariantism looks weak.[16]

Notes

1 The assumption that the objects of knowledge are propositions is convenient, but inessential to the argument.

2 For recent defences of contextualism in this sense see Cohen (1999), DeRose (1992, 2004) and Lewis (1996). Features of Cohen's and DeRose's examples are used at several points in the text.

3 Hawthorne (2003) tentatively defends sensitive invariantism. Jason Stanley has developed another version of the view in unpublished writings. It is unclear whether to classify some earlier work by other authors as contextualist or as sensitive invariantist.

4 The distinction corresponds to that between context of utterance and circumstance of evaluation, in the terminology of Kaplan (1989).

5 Williamson (2004b) discusses the role of principles of charity in philosophical methodology, arguing (compatibly with the claims of this paper) for a principle that maximizes knowledge rather than true belief (and true assertion). Should one apply contextualism or sensitive invariantism in determining what it means to 'maximize knowledge', with consequent risk of circularity? Fortunately, the present application of charity is too coarse-grained for that complication to matter.

6 That metalinguistic terms may themselves be indexical does not imply that it is unnecessary to take account of their use.

7 Schiffer (1996) complains that ordinary speakers are unaware of indexicality in 'know'. Some contextualists reply that the same goes for 'flat'. Hawthorne (2003, pp. 109-11) and Williamson (200?) explain how ignorance of indexicality makes problems for the preservation of information in memory.

8 Williamson (200?) develops the analogy with moral and practical problems in detail. It is mentioned by Feldman (2001, pp. 72-3), to whom Cohen (2001) replies.

9 See Kahneman and Tversky (1984) for related frame effects. The relevance of such work in psychology to sceptical arguments was pointed out in Vogel (1990); Spicer (2004) has further discussion.

10 See Hawthorne (2003) and Williamson (2000, 200?) for related discussion.

11 The derivation of (7*) from KPR assumes that ~K*[S*,*p**,*t**,*w**] is true in C* iff K*[S*,*p**,*t**,*w**] is not true in C*. In suitably constructed examples this assumption is uncontentious.

12 The failure of higher-order analogues of the 'KK' principle is one of many differences between the anti-luminosity objection to it and the simple-minded idea that externalist accounts of knowledge have us know without knowing that we know; the latter idea is often based on an illicitly internalist understanding of 'know' when the sentence in its scope includes epistemic vocabulary. Note also that the anti-luminosity objection does not invoke limits on the ability to conceive or believe complex contents.

13 As the case is sometimes envisaged, it is not obvious that Hi even believes that the plane stops in Chicago, still less that Hi believes that they 'know' that it does. If Hi does not believe something, that already explains why Hi does not know it, without any appeal to shifting standards. Since what matters most is whether Hi is in a position to know, assume for the sake of argument that Hi has the relevant beliefs (perhaps Hi is reckless, or does not realize how much is at stake).

14 'I don't know' here is tantamount to the assertion that one does not know whether one knows that one knows *q*. To have warrant for that assertion, one must know that one does not know whether one knows that one knows *q*. But one may know that one knows *q*, yet be in no position either to know that one knows that one knows *q* or to know that one does not know whether one knows that one knows *q*

(exercise: illustrate this with a model in possible worlds semantics for epistemic logic). In such cases, one is not even in a position to admit ignorance. Silence is golden.

15 The same bias may affect Lo in judging whether Hi knows. The most straightforward low standards situations involve a failure even to consider high standards situations.

16 I thank participants at the Stirling conference and in Oxford classes, especially Jessica Brown, Stewart Cohen, Keith DeRose, Maria Lasonen, Duncan Prichard and Nico Silins, and the referees for helpful comments on earlier versions of this material.

Bibliography

Cohen, S. 1999. 'Contextualism, skepticism, and the structure of reasons'. *Philosophical Perspectives*, 13, pp. 57-89.

Cohen, S. 2001. 'Contextualism defended: Comments on Richard Feldman's "Skeptical problems, contextualist solutions"'. *Philosophical Studies*, 103, pp. 61-85.

Conee, E., and Feldman, R. 1998. 'The generality problem for reliabilism'. *Philosophical Studies*, 89, pp. 1-29.

DeRose, K. 1992. 'Contextualism and knowledge attributions'. *Philosophy and Phenomenological Research*, 52, pp. 913-29.

DeRose, K. 2004. 'The problem with subject-sensitive invariantism'. *Philosophy and Phenomenological Research*, 68, pp. 346-50.

Dretske, F. 1970. 'Epistemic operators'. *Journal of Philosophy*, 67, pp. 1007-23.

Feldman, R. 2001. 'Skeptical problems, contextualist solutions'. *Philosophical Studies*, 103, pp. 87-98.

Hawthorne, J. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press.

Kahneman, D. and Tversky, A. 1984. 'Choices, values, and frames'. *American Psychologist*, 39, pp. 341-50.

Kaplan, D. 1989. 'Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals'. In J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan* (New York: Oxford University Press), pp. 481-563.

Lewis, D. 1996. 'Elusive knowledge'. *Australasian Journal of Philosophy*, 74, pp. 549-67.

Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.

Schiffer, S. 1996. 'Contextualist solutions to scepticism'. *Proceedings of the Aristotelian Society*, 96, pp. 317-33.

Spicer, F. 2004. 'Epistemic intuitions and epistemic contextualism'. Typescript.

Vogel, J. 1990. 'Are there counterexamples to the closure principle?'. In M. Roth and G. Ross (eds.), *Doubting* (Dordrecht: Kluwer), pp. 13-27.

Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. 2004. 'Philosophical "intuitions" and scepticism about judgement'. *Dialectica*, 58: 109-53.

Williamson, T. 200?. 'Knowledge, context and the agent's point of view'. In G. Preyer and G. Peter, eds., *Contextualism in Philosophy*, forthcoming.