

Knowledge, Context and the Agent's Point of View

Timothy Williamson

1. Contextualism and Relativism

Contextualism is relativism tamed.

Relativism about truth is usually motivated by the idea of no-fault disagreement. Imagine two parties: one (she) says 'P'; the other (he) says 'Not P'.¹ Apparently, if P then 'P' is true and 'Not P' false, so she is right and he is wrong; if not P then 'P' is false and 'Not P' true, so he is right and she is wrong. In both cases, there is an asymmetry between the two parties. Since P or not P (by the law of excluded middle), there is indeed an asymmetry between them, one way or the other. Yet the two parties may strike a neutral observer as on a par, equally intelligent, informed, perceptive and alert. Relativists about truth strive to dissolve the unpleasant asymmetry: "'P" is true for her; "Not P" is true for him'. Trouble starts when we ask what the relativists mean by 'for' in the construction 'true for X'. If to call something true 'for' X is just to say that X believes that it is true, then the attempted dissolution amounts to this: 'She believes that "P" is true; he believes that "Not P" is true'. But that is to add no more than that both parties believe that they are right; it does nothing to undermine the argument for an asymmetry between them. Relativists had better mean something else by 'true for X'. When asked to explain what else they mean, wild relativists bluster incoherently.

Contextualists, by contrast, have a clear answer. A sentence is true for X if and only if it is true as uttered by X, true relative to a context in which X is the speaker. Such relativism is tame because the relativity to context in the truth-value of a sentence allows for absoluteness in the truth-value of what the sentence is used to say in a given context. When she says 'P', she speaks truly: not just truly for her, but absolutely truly. When he says 'Not P', he too speaks truly: not just truly for him, but absolutely truly. The argument for asymmetry

assumes that, when she says 'P', she speaks truly if and only if P, and when he says 'Not P', he speaks truly if and only if not P. But that assumption is only as good as the assumption that, when she says 'P', she says that P, and when he says 'Not P', he says that not P. Contextualism denies that assumption in the cases at issue. If I report a speaker who utters the sentence 'P' as having said that P, in effect I assume, contrary to contextualism, that she said in her context what I would have said in my context by uttering the same words.

One great strength of contextualism is that it is uncontroversially correct about some cases. If she says 'I am a woman' while he says 'I am not a woman', it would not normally occur to us even for a moment to think of them as thereby disagreeing, although verbally his sentence is the negation of hers. We automatically apply the rule that 'I' as used in a given context refers to the speaker of that context: when she uses 'I', it refers to her; when he uses 'I', it refers to him. The disquotational clause that I am the referent of 'I' as used by him or her has no allure. Thus when she says 'I am a woman', she says that she is a woman; consequently, she speaks truly if and only if she is a woman; she does not say (falsely) that I am a woman. When he says 'I am not a woman', he says that he is not a woman; consequently, he speaks truly if and only if he is not a woman; he does not say that I am not a woman. In reporting their speech correctly, we make the possibility manifest that both parties spoke truly. It should be manifest even to the parties themselves. Although they can and will disquote on 'I' in reporting their own speech, they cannot and will not do so in reporting the speech of the other.

When the context-relativity is less clearly rule-governed, we may still adjust to it equally smoothly. In the sentence 'It is one of them', the pronouns can refer to almost anything, but usually no confusion results. Sometimes we make allowances for less obvious context-relativity. If she says 'Jan is tall' in a conversation about jockeys while he says 'Jan is not tall' (speaking of the same person at the same time) in a conversation about basketball players, we may allow that both parties spoke truly. Disquotation in reporting their speech is initially somewhat more tempting than in the previous cases. I might report her as having said that Jan is tall and him as having said that Jan is not tall. On reflection, however, I had better not stick to both reports in a single context. For if she said that Jan is tall, then she spoke truly if and only if Jan is tall, and if he said that Jan is not tall, then he spoke truly if and only if Jan is not tall. Therefore, if she said that Jan is tall and he said that Jan is not tall, they did not both speak truly (assume that I am not myself shifting contexts in mid-sentence). Our difficulty in reporting them is that we lack other words to express in our context what the word 'tall' expresses in contexts with other standards, in which 'tall' does not express what it expresses in our context. Nevertheless, we can describe what seems to be going on clearly enough in schematic terms. In her context, the word 'tall'

expresses the property of being tall(hers) (tall for a jockey); in saying 'Jan is tall', she says that Jan is tall(hers). In his context, 'tall' expresses the property of being tall(his) (tall for a basketball player); in saying 'Jan is not tall', he says that Jan is not tall(his). Being tall(hers) and being tall(his) are different properties, for although whatever is tall(his) is tall(hers), some things are tall(hers) without being tall(his). Thus it may be that both parties are speaking truly.

Contextualists can elucidate their point by talking of propositions. In that framework, if p is the proposition that P, uttering a sentence in a context in which it expresses p with the appropriate declarative force is necessary and sufficient for saying that P. The contextualist claim about a given sentence is that it expresses different propositions as uttered in some different contexts. As uttered by her, the sentence 'I am a woman' expresses the true proposition that she is a woman, so its negation expresses the false proposition that she is not a woman; as uttered by him, 'I am a woman' expresses the false proposition that he is a woman, so its negation expresses the true proposition that he is not a woman. The difference between the two contexts in the reference of 'I' makes a difference between them in what proposition is expressed. In all this, the linguistic meaning of the sentence and its constituent words is held fixed: both speakers are using the word 'I' with its standard English meaning; one rule determines the contextual variation in reference. A dictionary of English does not need separate entries for 'I' as used by different speakers. Understanding the word 'I' requires mastering this contextual variation, in order not to misinterpret other speakers. Contextualism does not concern the trivial point that physically indistinguishable words in different languages may have different meanings (the Italian word 'burro' means *butter*; the Spanish word 'burro' means *donkey*).

If the sentence s as uttered in a context c expresses the proposition that P, then s is true in c if and only if P, and false in c if and only if not P. As uttered in the present context, the sentence 'P' expresses the proposition that P; thus 'P' is true in the present context if and only if P, and false if and only if not P. As uttered in some other context, 'P' expresses some other proposition, that Q; thus 'P' is true in that other context if and only if Q. Since a sentence may be true in one context and false in another, truth and falsity for sentences are relative to a context. At least some apparent disagreements dissolve if the parties speak in different contexts. But truth and falsity for propositions are as absolute as you like. If 'P' expresses the proposition that P(hers) as uttered in her context and the distinct proposition that P(his) as uttered in his context, then P(hers) may be absolutely true while P(his) is absolutely false, independently of context, because P(hers) does not entail P(his). When he says 'Not P', he expresses the absolutely true proposition that not P(his). Although the sentence that he

utters is the negation of the sentence that she utters, the proposition that he expresses is not the negation of the proposition that she expresses. They are talking past each other. In this tame way, contextualists restore the two parties to parity.

Normally, a difference between the propositions that a particular sentence expresses in different contexts can be traced to a difference between the contributions that a particular constituent word such as 'I' or 'tall' makes in those contexts to what proposition the sentence expresses. Sometimes, as with 'I', it is a blatant difference in reference. In other cases, the difference must be inferred from a difference in the truth-value of the whole sentence, which involves eliminating other subsentential constituents as the source of the contextual variation.

The notion of a difference in the truth-value of a sentence must be applied with care. The sentence 'Descartes died in 1650' expresses a truth as uttered in the actual world; it expresses a falsehood as uttered in a counterfactual world in which Descartes died in 1651. But such a difference in truth-value does not imply any contextual variation in the relevant sense in 'Descartes died in 1650'. As uttered in the counterfactual world, the sentence expresses the very same proposition that it expresses as uttered in the actual world, the proposition that Descartes died in 1650. If Descartes had died in 1651, someone who said 'Descartes died in 1650' (meaning by the words what they actually mean in English) would still have said that Descartes died in 1650. The difference is that the proposition that Descartes died in 1650 is actually true but would have been false if Descartes had died in 1651, because Descartes died in 1650 in the actual world but in 1651 in the counterfactual world. The truth-value of the proposition is not an essential property. In the terminology of David Kaplan (1989), we have varied not only the context of utterance (in particular, the world in which a proposition is relevantly expressed) but also the circumstance of evaluation (in particular, the world in which the proposition is relevantly true or false). For variation in the truth-value of a sentence as uttered in different contexts to constitute a genuinely contextual effect, the sentence must be evaluated with respect to a fixed circumstance. Unless the circumstances agree on whether P, the sentence 'P' may express the proposition that P in both cases and differ in truth-value in a non-contextualist way.

The criteria just sketched for isolating contextualist phenomena are not always straightforward to apply, even when we are clear about the truth-values of the sentences at issue as uttered in the relevant contexts. The temporal dimension is particularly tricky. On the majority view, defended by Frege, propositions cannot change in truth-value; since the sentence 'It is Monday' expresses a truth as uttered on Monday and a falsehood

as uttered on Tuesday, it expresses different propositions as uttered on different days. More generally, any significant use of tenses creates a genuinely contextualist effect. By contrast, on a minority view, defended by Prior, propositions can change in truth-value; the sentence 'It is Monday' expresses a single proposition, that it is Monday, which is true on Monday and false on Tuesday. That would no more be a contextualist effect than is the contingency in truth-value of 'Descartes died in 1650'. Fortunately, we need not decide the proper treatment of tenses for the purposes of this essay, although we must be circumspect in applying the framework to diachronic cases.

2. The Contextualist Strategy

According to contextualists, many philosophical problems result from subtly concealed forms of context-relativity. The mark of paradox is that apparently obvious claims are apparently jointly inconsistent. On the contextualist diagnosis, the inconsistency is really only between the sentences that express those claims, which we accept in different contexts: each sentence expresses a true proposition as uttered in the context in which we accept it, but in no single context do they all express true propositions, or do we accept them all; what varies is not the truth-value of any given proposition (the circumstance of evaluation is fixed), but rather which proposition a given sentence expresses. The contextualist resolution of apparent conflicts has become one of the most fashionable strategies in contemporary philosophy, not least through its systematic application and advocacy by David Lewis.²

The classic example is contextualism about the word 'know'. Many epistemologists find each claim in triplets such as the following highly plausible:

- (1) Mary knows that she had her purse yesterday morning.
- (2) If Mary knows that she had her purse yesterday, then Mary knows that the universe was not created an hour ago (with misleading apparent traces of millions of years of past history ...).
- (3) Mary does not know that the universe was not created an hour ago (with misleading apparent traces of

millions of years of past history ...).

For (1) seems to report a typical case of ordinary common sense knowledge; (2) seems to register Mary's capacity to extend her knowledge by deductive reasoning, for we may assume that she believes the conclusion that the earth was not created a week ago on the basis of competent deduction from the premise that I had my purse yesterday (and perhaps on many other equally good bases too); (3) seems to record her inability to rule out possibilities to which her methods of belief formation are wholly insensitive. Yet the sentences (1)-(3) constitute an inconsistent triad of the form { 'P', 'If P then Q', 'Not Q' }. To resolve the difficulty, contextualists distinguish between an ordinary context, in which we use the word 'know' in speaking of everyday concerns, according to comparatively low standards for its correct application, and an extraordinary epistemological context, in which we use the word 'know' to obsess about sceptical scenarios, according to much higher standards for its correct application.³

In the ordinary context, 'know' expresses a relation of know(low)ing. One can know(low) that P even though one has no way of persuading a sceptic that one is not in an apt sceptical scenario in which one falsely believes that P, for such sceptical scenarios are in some sense irrelevant to the ordinary context. Mary know(low)s that she had her purse yesterday morning. Moreover, deduction is a way of extending know(low)ledge; if one believes the conclusion that Q on the basis of competent deduction from the premise that P (and perhaps some other premises too), and one know(low)s that P (and those other premises too, if any), then one know(low)s that Q. Thus if Mary know(low)s that she had her purse yesterday morning then she know(low)s that the universe was not created an hour ago. Therefore, she know(low)s that the universe was not created an hour ago.

In the epistemological context, 'know' expresses a relation of know(high)ing. One cannot know(high) that P if one has no way of persuading a sceptic that one is not in an apt sceptical scenario in which one falsely believes that P, for such sceptical scenarios are in some sense relevant to the epistemological context. Mary does not know(high) that the universe was not created an hour ago. Equally, she does not know(high) that she had her purse yesterday morning. Therefore, on a truth-functional reading of the conditional, if Mary know(high)s that she had her purse yesterday morning then she know(high)s that the universe was not created an hour ago. More generally, deduction is a way of extending know(high)ledge, given that we are not here concerned with scepticism about the validity of the deductive reasoning itself. If one believes the conclusion that Q on the basis

of competent deduction from the premise that P (and perhaps some other premises too), and one know(high)s that P (and those other premises too, if any), then one know(high)s that Q. Thus, again, if Mary know(high)s that she had her purse yesterday morning then she know(high)s that the universe was not created an hour ago.

In the ordinary context, sentences (1) and (2) express truths while (3) expresses a falsehood. In the epistemological context, (2) and (3) express truths while (1) expresses a falsehood. Since (1)-(3) are jointly inconsistent, in no context do they all express truths. When we initially consider (1), we tend to remain in an ordinary context, for no sceptical scenario has been mentioned. But (2) and (3) advert to a sceptical scenario in which the universe was created an hour ago with misleading apparent traces of millions of years of past history. Thus considering (2) and (3) tends to put us into the epistemological context, and merely reconsidering (1) is not enough to get us out again: we may simply become more doubtful of (1). It takes less to summon up an evil demon than to exorcise him.

Different contextualists postulate different context-shifting mechanisms, although most of them make some play with the idea of change in the contextual relevance of various possibilities of error. This paper prescind from the details of particular contextualist accounts in order to investigate more general theoretical issues about epistemological contextualism as such.

It is important to check whether specifically contextualist ideas are crucial to the resolution of the paradox. Contextualism about 'know' does not follow merely from this possibility: when Mary is in a police station at noon, reporting the theft of her purse, it is true to say 'Mary knows that she had her purse yesterday morning', but when she is in an epistemology seminar that afternoon, discussing scepticism, it is false to say 'Mary does not know that she had her purse yesterday morning'. For we can make sense of that story by assuming that 'know' stands for a single temporary relation which a subject may have to a truth at one time and not at another, as in learning and forgetting, although here what makes the difference is her conversational situation (if she continues to believe even in the seminar that she had her purse yesterday morning).⁴ Even if that is not the best treatment of the present tense in 'knows', the upshot would be at most a quite general form of contextualism about tenses, not anything specific to 'know': that is not epistemological contextualism. Similarly, one cannot argue for contextualism about 'tall' merely by pointing out that the sentence 'Michael is tall' may express a truth when Michael is an adult and a falsehood when he is a little boy. For genuine contextualism about 'know', we must keep the time and world of evaluation fixed. But contextualists argue that we can do so. Even if Mary herself is in an ordinary context at the relevant time, worrying about the loss of her

purse rather than sceptical scenarios, contextualists will still ascribe truth to both ‘Mary knows that she had her purse yesterday morning’ as uttered in the ordinary context of a conversation about Mary’s loss and ‘Mary does not know that she had her purse yesterday morning’ as uttered simultaneously in the epistemological context of a seminar about scepticism. It is not a question of her knowing something at one time and not at another, or in one world and not in another. Saving the truth of the two simultaneous and apparently contradictory assertions is a specifically contextualist move.

Sceptical arguments have a pull on our assent that cries out for explanation. Their power is felt immediately by a high proportion of those who are willing to listen to them carefully. Contextualists explain this power by conceding a semantics for ‘know’ on which sceptics speak truly, for example in uttering (3) and the negation of (1). But this is not total capitulation to scepticism, for contextualists also explain the pull that common sense has on our assent, for on their semantics we also speak truly when we make ‘knowledge’-ascriptions in ordinary contexts, for example in uttering (1).⁵ Sceptics speak truly only by creating contexts with extraordinarily high standards for the correct application of ‘know’, in which their favourite sentences express truths.

Contextualist explanations of the power of scepticism need not be *ad hoc*, for they may appeal to independently testable generalizations about all contexts of utterance, not just those at issue in sceptical problems. For example, such generalizations may also predict a rise in standards for the correct application of epistemic terms when more is at stake practically. ‘The man in the information booth knows that no train is coming this afternoon’ may be true as uttered by the disappointed trainspotter on the platform, false as uttered simultaneously by the sleepy sunbather on the line. The predictions may also cover many terms of epistemic appraisal beyond ‘know’, such as ‘justified’. These predictions may seem to be verified.

3. Contextualism, Vagueness, and the Transmission of Information

One feature of many philosophically significant terms that apparently predisposes them to contextualist treatment is their *vagueness*. We understand them not by learning precise definitions but by extrapolating from examples which leave their application to ranges of borderline cases unclear. In many contexts, speakers find it convenient to resolve some of this vagueness in one way or another, according to their practical purposes.

Naturally, they will sometimes find it convenient to resolve the vagueness in opposite ways in different contexts. One local stipulation about the extension of 'red' makes it include x ; elsewhere, another local stipulation about the extension of 'red' makes it exclude x . Vague terms appear not to cut nature at the joints, not to pick up hidden but sharp and uniquely natural divisions into kinds that might stabilize their reference: on the epistemicist view, a vague term has hidden sharp boundaries as used in a given context, but that does not stop it from having different hidden sharp boundaries as used in another context. Of course, context-relativity is not the very same phenomenon as vagueness: that 'I' refers to John as uttered by John and to Mary as uttered by Mary is an example of context-relativity without being an example of vagueness. Nevertheless, one might think that the vagueness of a term makes contextual variation in its reference practically irresistible (even though it also makes the variation hard to measure). In this light, it looks as though it would be an amazing coincidence if a vague term did have exactly the same reference in all contexts.⁶

Epistemic terms such as 'know' and 'justified' are surely at least somewhat vague; it seems obvious that they are not perfectly precise. Between the clear cases of their application and the clear cases of their non-application, they seem to have a range of borderline cases, just as one might expect. If that gives us independent reason to predict their context-relativity, then surely it is a natural strategy to put that context-relativity to work in epistemologically significant ways, as contextualists try to do.

It is beginning to look as though context-relativity must be virtually ubiquitous in natural languages, just as vagueness is. However, it comes at a cost, as we can see by considering the preservation of information in memory and its transmission by testimony.⁷ Let the sentence 'P' as uttered in different contexts express different propositions, some true, some false. Suppose that in a context c you acquire an item of information, expressed by 'P' as uttered in c . You could store the sentence 'P' in your memory. But when the time comes to retrieve the information, for guiding action or passing on to someone else, you will be in a new context c^* ; even if 'P' expresses a truth as uttered in c , it need not express a truth as uttered in c^* . If in c^* you remember that you acquired 'P' in c , and what c was like, you might be able to construct another sentence 'P*' that as uttered in c^* expresses the very truth that 'P' expresses as uttered in c . But that procedure imposes a heavy burden on memory. Most people cannot remember where they got much of the information on which they rely. Alternatively, when you first acquire the information, in c , you could seek another sentence 'Q' that expresses independently of context the same proposition that 'P' expresses in c ; you could then store the sentence 'Q' in your memory. But if such an eternal sentence 'Q' were readily available for the occasion sentence 'P' on each

occasion, then the context-relativity of 'P' would be, if not idle, at least underemployed.

One obvious problem is that you might not always know the value of some contextual parameters that affect the reference of terms in 'P'. If you do not know where you are, how can you replace 'here' by a context-insensitive designation of the place? If you do not know what time it is, how can you replace 'now' by a context-insensitive designation of the time? You could use 'there' and 'then' as memory demonstratives, but that requires preserving some memory of the incident. You could existentially quantify out by using 'somewhere' and 'sometime', but the resort to such generality involves significant loss of information. Obviously, these problems do not show that we would be better off without 'here' and 'now'; they even illustrate the utility of the words. But they also illustrate their limitations when information is to be preserved and transmitted.

The reference of paradigm indexicals such as 'here' and 'now' varies with easily isolated parameters of context, such as place and time, whose values we at least often know enough to designate non-indexically. I know that in my current context 'here' corefers with 'Room 2, Staircase 10, Old Buildings, New College, Oxford OX1 3BN, UK' and 'now' corefers with '11.23 a.m., 9 December, 2003' (give or take a little specificity). But contextualism makes the reference of epistemic terms such as 'know' and 'justified' vary with a more elusive parameter, epistemic standards. We have no received scale for conveniently specifying epistemic standards in terms that contextualists would regard as context-insensitive. Phrases such as 'high standards' and 'low standards' are themselves context-sensitive, and in any case far too vague and unspecific for identifying location on a continuum of standards. In answer to the questions 'How high?' and 'How low?', one finds oneself giving a detailed account of who said what. One could try a description such as 'The epistemic standards in force in Room 2, Staircase 10, Old Buildings, New College, Oxford OX1 3BN, UK at 11.23 a.m., 9 December, 2003', but such a description would clearly be of little use to those who did not know what was going on at that place at that time. If I simply store the sentence 'At 11.23 a.m., 9 December, 2003, I did not know whether I had my wallet the previous morning', it may express a falsehood the next time I produce it, owing to a change of epistemic standards. If I store the sentence 'At 11.23 a.m., 9 December, 2003, I did not have the relation then expressed in my context by the word "know" to whether I had my wallet the previous morning', it will do me little good the next time I produce it if I cannot remember enough about the past context to have much idea of what relation the word 'know' then expressed. Evidently, the more extensively 'know' varies in reference with context, the worse the problems of using it to preserve and transmit information will be. Yet

'know' does not seem to be designed like 'here' and 'now' primarily for immediate consumption; we need to preserve and transmit information about who has what kind of knowledge, or who knew what when.

Epistemological contextualism has been criticized for postulating contextual variation of which ordinary speakers are unaware (Schiffer 1996). Normal competence with paradigm indexicals and demonstratives such as 'I', 'you', 'here', 'there', 'now', 'then', 'he', 'she', 'it', 'this', 'that' and 'these' makes their context-sensitivity manifest to one; no philosophical argument is needed to persuade one of its existence. The hypothesis of hidden context-sensitivity is therefore taken to be implausible: if our word 'know' is context-sensitive, why is that not obvious to us? Whatever one thinks of that argument, the use of 'know' in preserving and transmitting information is even more problematic if it is context-sensitive in ways of which ordinary speakers are unaware. For will they not be liable to store information in the shape of sentences involving the word 'know' and produce those sentences again in new contexts, in which they no longer express truths? The same point applies to other philosophically contentious contextualist claims. For their very contentiousness suggests that the context-sensitivity is not manifest to ordinary speakers, who are therefore liable to run afoul of it in attempting to preserve and transmit information.

The practical difficulties in using context-sensitive terms of epistemic appraisal do not constitute a decisive argument against contextualism (Hofweber (1999) and Rysiew (2001) argue for similar conclusions). To some extent as speakers we could work round them in various *ad hoc* ways, especially if the contextual variation is not too extensive. Even if it is hard for us to become consciously aware of the contextual variation, we might somehow unconsciously adjust to it in preserving and transmitting information: much of native speakers' competence with a natural language is not open to their view. Nevertheless, the practical difficulties show that context-sensitivity carries significant disadvantages; if one were designing a language, one would not want too much of it. But the previous consideration of vagueness suggested that context-sensitivity might be unavoidable in virtually all terms of interest.

The aim of the second half of this essay is to develop a model that makes terms such as 'know' and 'justified' context-invariant without requiring them to be precise. The next section explores a related issue in practical affairs. Section 5 applies its moral to contextualism in epistemology.

4. Contextualism and Practical Reason

Imagine that Clare faces a difficult practical decision: whether to resign her good job tomorrow for reasons of principle. She thinks ‘If it would be wrong for me not to resign, I will resign; if it would not be wrong for me not to resign, I will not resign. But would it be wrong for me not to resign?’⁸ The relevant considerations are complex. Some tell in favour of the conclusion that it would be wrong for her not to resign; others tell against it. She is unsure what relative weights to assign the conflicting factors. When the considerations in favour of resigning are more salient to her, she is disposed to judge that it would be wrong for her not to resign; when the considerations against resigning are more salient to her, she is disposed to judge that it would not be wrong for her not to resign. As the moment for decision approaches, she oscillates, agonizing, between the opposed views.

Suddenly, a contextualist appears and says to Clare:

Do not worry. You are mistaken in supposing that there is a disagreement between what you think when you think ‘It would be wrong for me not to resign’ and what you think when you think ‘It would not be wrong for me not to resign’. Both thoughts are true.⁹ The sentence ‘It would be wrong for me not to resign’ expresses different propositions as uttered in different contexts. Sometimes, the considerations in favour of resigning are more salient to you; that creates a context in which ‘It would be wrong for me not to resign’ expresses a truth, to which you are disposed to assent. At other times, the considerations against resigning are more salient to you; that creates a context in which ‘It would not be wrong for me not to resign’ expresses a truth, to which you are also disposed to assent. You have no need to reject one of the two thoughts.

This contextualist resolution of Clare’s problem is liable to strike us as glib and shallow. When the moment for action comes, Clare must either resign or not resign. Given her intention expressed in the words ‘I will resign if and only if it would be wrong not to do so’, by resigning she goes only with the thought ‘It would be wrong for me not to resign’, and by not resigning she goes only with the thought ‘It would not be wrong for me not to resign’. She cannot have it both ways.

Of course, context-sensitive elements occur in the sentence ‘It would be wrong for me not to resign’:

for example, the first person pronoun and the tense of the finite verb. But no variation in their reference explains Clare's oscillation. Throughout, 'me' refers to Clare, the time reference is to the time (which we may assume to be fixed) when she must resign if she is to do so at all, and so on.

Equally unappealing would be a contextualist resolution of the apparent disagreement between two commentators, one saying 'It would be wrong for Clare not to resign' while the other says simultaneously 'It would not be wrong for Clare not to resign'.

Can we go any deeper in explaining why the contextualist resolution misses the point of Clare's problem? In cases of decision-making, one context is distinguished above all others: that of the agent at the moment of action. The primary question is whether the sentence 'It would be wrong for me not to resign' expresses a truth as uttered in the context in which the speaker is Clare and the time is that for resigning if she is to do so at all.¹⁰ Call that context the *agent's context*, and the proposition which the relevant first person present tense sentence (such as 'It would be wrong for me not to resign') expresses in it the *agent's proposition*. If at some point in her agonizing Clare uses the sentence to express a proposition other than the agent's proposition, which might fail to match the agent's proposition in truth-value, she is no longer concentrating on the relevant practical problem. Similarly, an external commentator who uses the sentence 'It would be wrong for Clare not to resign' to express a proposition other than the agent's proposition is no longer concentrating on the relevant practical problem for Clare.¹¹ But if the sentence expresses the agent's proposition in all the contexts at issue, then it expresses the *same* proposition in all those contexts, and contextualism fails for this case.

Where contextualism applies, the speaker's context is *autonomous* with respect to the agent's context in determining what propositions the relevant sentences express: the speaker's proposition need not be the agent's proposition. We have seen reason to suspect that, in practical matters, the speaker's context lacks such autonomy. The lack of autonomy would not be epistemic: for all that it implies, the speaker (a disinterested observer, or the agent herself in a cooler hour, or ...) may be in a better position than the agent at the time of action to know what is right or wrong. That Clare is the one who stands to lose her job does not necessarily help her to know whether it would be wrong for her not to resign. The speaker's context lacks autonomy only in setting the content of the practical question. In some sense 'Would it be wrong?' is above all a question *for* the agent, in the context of agency (which is not to deny that the agent may fail to raise it); it is a question for others, or for the agent in other contexts, *because* it is a question for the agent in the context of agency. Of course, others can ask the question for some instrumental reason, for example in making predictions about the

agent's behaviour in other situations. But the true answer to a question is not always the same as the most useful answer for some immediate practical purpose. To think of the true answer to the moral question about a fixed action as oscillating with the practical purposes of the questioner seems to lose sight of the meaning of the question. Plausibly, therefore, shifting from the context of agency to other contexts does not shift the content of the question.

If the contextualist misdescribes the practical case, and Clare's utterances of 'It would be wrong for me not to resign' and 'It would not be wrong for me not to resign' express genuinely incompatible contents, then for part of her deliberations Clare is disposed to misjudge whether it would be wrong for her not to resign. The contextualist might challenge us to explain why Clare would commit such an error. But an explanation is not far to seek. For, as the case was originally described, there were complex, conflicting considerations for and against resigning, with no obviously correct standard for weighting them against each other. One would expect Clare to assign more weight to considerations when they are psychologically salient to her than when they are not. As she goes through the considerations in her mind, one after another, naturally they fluctuate in salience to her. In the circumstances, it would be astonishing if her judgement did *not* fluctuate too. No special semantic hypothesis is needed to explain such a familiar psychological phenomenon.

The argument has evidently done nothing to show that, unexpectedly, 'wrong' is precise. The point is rather that, however vague the term, its meaning does not permit variation in content across contexts. In that respect we might compare 'wrong' with a proper name. The nature of a proper name is to have a constant character in Kaplan's sense: in each context it has the same content (it refers to the same object). In this sense, proper names which differ in reference are different names, even if they look and sound the same ('John Smith' and 'John Smith'). But it does not follow that proper names are perfectly precise.

The argument does not imply that the orthographic type 'wrong' is always used in English with the same sense. It may well have distinct moral and nonmoral readings in the sentence 'It is wrong to answer "Thirteen" to the question "What is the sum of seven and five?"', with a corresponding difference in truth-value. But that is ambiguity, as in 'bank', not indexicality. It is a given sense of an expression that is indexical; one must be sensitive to the indexicality in order to understand the expression with that sense. An ambiguous orthographic type has several senses, understanding it with any one of which does not imply understanding it with any other (it may be best to regard distinct senses as marking out distinct but orthographically coincident expressions). Contextualists allege indexicality rather than ambiguity: for example, a contextualist

understanding of 'know' requires sensitivity to the contextual variation in its reference. On the account just sketched, the orthographic type 'wrong' has at least one non-indexical sense, for which contextualism fails.

The argument is suggestive rather than irresistible. Nevertheless, let us see what happens when we try to develop a similar account for the word 'know'.

5. Practical Epistemic Evaluation

According to contextualism about 'know', a sentence of the form 'S knows at t that P' may be true as uttered in a context c yet false as uttered in another context c^* , even though in both contexts 'S' refers to the same subject, S, ' t ' refers to the same time, t , and 'P' expresses the same proposition, that P: what differs is the referent of 'know'.¹² In this case the agent's context, the context being spoken about, is the context of S as agent at time t .¹³ Then at least one of c and c^* differs from the agent's context in the referent of 'know', since otherwise they would not differ from each other in that respect. Similarly, the present-tensed 'I know that P' in the agent's context differs in truth-value from 'S knows at t that P' in one of c and c^* , otherwise c and c^* would not differ from each other in the truth-value of the latter sentence. Given that 'S knows at t that P' expresses different propositions in c and c^* , at least one of those propositions differs from the agent's proposition, the proposition that 'I know that P' expresses in the agent's context. Contrapositively, if the speaker's context lacks autonomy in ascriptions involving 'know', in the sense that 'know' must have the reference in the speaker's context that it has in the agent's context, then contextualism fails for 'know'.

Even in the epistemic case, the lack of autonomy of the speaker's context would be non-epistemic in the same sense as in the practical case: for all that it implies, the speaker (an external observer, or S herself at a time other than t , or ...) may be in a better position than the agent S at t to know whether S knows at t that P. S may think that she knows that the animal in the cage is a zebra, while the zoo-keeper knows that S does not know that it is a zebra, because he knows that it is a cleverly disguised mule. The speaker's context would lack autonomy only in setting the content of the epistemic question. In some sense 'Does S know at t that P?' is above all a question *for* the agent S, at time t (which is not to deny that S may fail to ask the question); it is a question for others, or for S at other times, *because* it is a question for S at t . Of course, others can ask the question for some instrumental reason, for example in deciding whether to rely on S's testimony. But the true

answer to a question is not always the same as the most useful answer for some immediate practical purpose. To think of the true answer to the epistemological question as oscillating with the practical purposes of the questioner risks losing sight of the meaning of the question. Shifting from the context of agency to other contexts may not shift the content of the question.

The hypothesis that the speaker's context lacks autonomy in 'know'-ascriptions already explains how 'know' could work in a non-contextualist way without being precise. No amazing coincidence in reference is required: just a primary role for the agent's context that cannot be trumped by the speaker's context.

If there is any reason for the speaker's context to lack autonomy in 'know'-ascriptions, then epistemological contextualism may well fail non-miraculously. But is there any reason for the speaker's context to lack autonomy in 'know'-ascriptions? Traditional epistemology has concentrated on the first-personal present-tensed question 'What do I know?', as if the agent's context had primacy (Descartes is the prime example). One might expect that primacy to deprive the speaker's context of autonomy. But the tradition is far from obviously warranted in treating the first person present tense question as the key to epistemology. For example, Edward Craig (1990) has argued that the primary point of the concept of knowledge is in distinguishing between reliable and unreliable informants. On his view, the third person question 'What do they know?' takes precedence over the first person 'What do I know?', so one might expect the speaker's context to be autonomous in relation to the agent's context. More generally, many epistemological externalists emphasize the significance for knowing of reliable causal or counterfactual connections between belief states and states of the external environment, irrespective of the subject's access to those connections; such externalism is naturally regarded as a third person approach to epistemology.¹⁴ Speakers, in their own contexts for their own purposes, classify others as 'knowing' or as 'not knowing'; why should they defer to those others for the content of their classification?

Such defences of the autonomy of the speaker's context in 'know'-ascriptions are not decisive. After all, a politician might classify the actions available to his more scrupulous opponents as 'right' or 'wrong', merely in order to predict their behaviour; that would not make his context as speaker autonomous in relation to their contexts as agents.¹⁵ However different his interests are from theirs, a sense seems to remain in which whether their potential actions are to be classified as 'right' or 'wrong' is a question for him because it is a question for them (even if they fail to raise it).

The epistemic case may be closer to the practical case than might initially be expected. For if one

knows that P, then one can hardly be wrong to believe that P; conversely, given that one does not know that P, it arguably is wrong to believe that P. 'Believe' here is used in the sense of outright belief; assigning a high subjective probability (short of 1) to the proposition that P does not suffice for believing that P. For example, in believing on statistical grounds that my ticket has a chance of 0.9999 of not winning the lottery, and consequently assigning a subjective probability of 0.9999 to the proposition that my ticket will not win, I do not have the outright belief that my ticket will not win, for that would involve outright belief in the crucial premise of an argument for not buying the ticket in the first place (however large the prize). Let us probe the hypothesis that what it is wrong to believe goes with what one fails to know.

Suppose that the animal in the cage is a cleverly disguised mule, and John is in no position to distinguish his state from one of knowing by sight that it is a zebra. Since it is not a zebra, he does not know that it is a zebra. Is he wrong to believe that it is a zebra? The answer 'No' might be defended on the grounds that John is justified on his evidence in believing that it is a zebra. We may grant straight away that on John's evidence it is highly probable that the animal is a zebra. That is at least some *excuse* for believing it to be a zebra, although by itself it is consistent with knowing that he does not know that it is a zebra (consider the lottery case). A better excuse is that he is not in a position to know that he does not know that it is a zebra. Indeed, on John's evidence, it is highly probable that he knows that it is a zebra. But why regard any of that as more than an excuse? Excuses are offered in mitigation of an offence; they do not make the offending act not wrong. If it had not been wrong, no excuse would have been needed. In particular, if John had known that it was a zebra, he would have needed no excuse for believing that it was a zebra.

To vary the example, suppose that the animal in the cage really is a zebra, although to save money most of the other animals in the zoo have (unbeknownst to John) been replaced by cleverly disguised farm animals; he is again in no position to distinguish his state from one of knowing by sight that it is a zebra. On the usual view, he still does not know that it is a zebra.¹⁶ He has the same excuses as before for believing it to be a zebra. On his evidence, it is highly probable that he knows that the animal is a zebra. Given that he does not in fact know that it is a zebra, he still seems to need some excuse for believing that it is a zebra, in which case it is wrong for him to believe that it is a zebra. That contrasts with knowing that it is a zebra, which is a full justification for believing that it is a zebra, not a mere excuse.¹⁷

Whether one believes that P naturally has many further implications for what one does with the proposition that P, in particular, for whether one uses it as a premise in practical reasoning. Indirectly, therefore,

it has widespread ramifications for what one does.¹⁸

Imagine a typical scenario for contextualism about 'know'. On the basis of memory, Mary believes truly that she had her purse yesterday morning; background conditions are normal. In an ordinary context, Speaker(low) says 'Mary knows that she had her purse yesterday morning'. Simultaneously, in a seminar on scepticism, Speaker(high) says, of the same person, 'Mary does not know that she had her purse yesterday morning'. The contextualist insists that both speakers speak truly. Suppose that both speakers endorse and apply the argument of the preceding paragraphs, as uttered in their respective contexts. Thus Speaker(low) says 'It is not wrong for Mary to believe that she had her purse yesterday'. Speaker(high) says 'It is wrong for Mary to believe that she had her purse yesterday'. Both speak truly only on a contextualist account of 'wrong'. The relevant sense of 'wrong' is doubtless not specifically moral. Nevertheless, this contextualist resolution of the apparent disagreement still seems glib and superficial. Since epistemic standards vary so wildly between the ordinary context and the epistemological context, they vary markedly between at least one of those contexts and the agent's context. Mary's context. Yet the primary question seems to be 'Is it wrong for me to believe that I had my purse yesterday?' as uttered by Mary. The epistemic standards relevant to answering that question are those operative in Mary's context. Hence at least one of Speaker(low) and Speaker(high) is judging Mary by inappropriate epistemic standards.

A contextualist about 'know' but not about 'wrong' might accept that diagnosis, and adjust the original line of thought according to those commitments, as follows. If 'I know that P' is true in the context of agent S at time *t*, then it is not wrong for S at *t* to believe that P; if 'I know that P' is false in the context of S at *t*, then it is wrong for S at *t* to believe that P.¹⁹ Thus epistemic standards in the agent's context, not the speaker's, determine what it is wrong for the agent to believe. Suppose, for example, that Mary, like Speaker(low), is in an ordinary context (she might even be Speaker(low)). Thus epistemic standards in Speaker(high)'s context are much higher than in Mary's. Speaker(high) should say 'Although Mary does not know that P, it is not wrong for her to believe that P, because she can truly say "I know that P"'.²⁰

From Speaker(high)'s perspective, the context of the epistemology seminar, is Mary's outright belief that P justified or just excused? Her belief is true, but she is ignoring various sceptical possibilities of error. Since they do not obtain, she gets away with it. Pragmatically, her attitude may be the most sensible or only one to take. Taking account in real time of many heterogeneous distant possibilities is computationally infeasible for a human brain. But that sounds like a very good *excuse* for cutting corners. If Mary knew that P, she would not

be cutting corners in believing outright that P; no pragmatic excuse would be needed, for she would be justified in a stronger sense. Since Mary does not know that P, it is strictly speaking wrong for her to believe outright that P, because some excuse is needed. But then the attempt to be contextualist about 'know' but not about 'wrong' collapses.

The argument so far has relied on a contested connection between 'know' and justification for belief. However, even without assuming such a connection, we can develop a variant of the anti-contextualist argument. For if contextualism about 'know' is well-motivated, so is contextualism about 'justified'. The contextual shifts in epistemic standards that appear to involve the reference of 'know' appear equally to involve the reference of 'justified'. For example, confronted with a sceptical scenario in which the universe was created an hour ago (with misleading apparent traces of millions of years of past history ...), Mary is liable to think 'Oh dear, I was not really justified in taking for granted that the universe has existed for millions of years, so none of my beliefs which rest on that assumption are justified', just as she is liable to think 'Oh dear, I do not really know that the universe has existed for millions of years'. To give a contextualist explanation of the latter reaction but not of the former would be suspiciously *ad hoc*. But contextualism about 'justified' is open to a more direct form of the objection that it grants inappropriate autonomy to the speaker's context in judging the agent, in setting the standard for when she needs an excuse.

Quite generally, it is plausible that, for some readings of some terms of epistemic evaluation, the appropriate standard will be set by the agent's context, the speaker's context will lack autonomy, and, on those very readings, the sorts of conversational phenomena that are supposed to motivate contextualism will nevertheless occur. In those cases, the real explanation of the phenomena will be non-contextualist. Indeed, just as with disagreements over the application of 'wrong', such an explanation is not far to seek. When standards are low, little is at stake and the sceptic is not bothering us, few possibilities of error are psychologically salient: no wonder that we do not doubt that cases which are or seem similar to socially accepted paradigms of knowledge (perception, memory, testimony, ...) are indeed cases of knowledge (or whatever privileged status is at issue). When standards are high, much is at stake or the sceptic is bothering us, more possibilities of error are psychologically salient: no wonder that we give the newly salient cases more weight than before in assessing claims to knowledge (or whatever privileged status is at issue). In order to apply terms like 'know', 'justified' and 'wrong', we must often balance conflicting considerations, with no formula to tell us how to assign them comparative weights. It would be astonishing if considerations for or against did not weigh more heavily with us

when they are salient than when they are not. Much of the rhetorical work in sceptical arguments goes into making the chosen scenarios of error as vivid and pressing as possible; of course that has an effect on our judgement. On the hypothesis that epistemological contextualism is false, the phenomena that are supposed to support it remain entirely predictable. Why invoke contextualism to explain them?²¹

Notes

- 1 Despite the conveniently gendered pronouns, the parties to the disputes may be social or intellectual groups rather than single individuals.
- 2 Lewis (1979) expounds his general contextualist approach.
- 3 Seminal (noncontextualist) discussions of the nonclosure problem are Dretske (1970) and Nozick (1981). Important examples of contextualist approaches include Cohen (1987, 1988, 1999), DeRose (1992, 1995, 200?), Lewis (1996), Stine (1976) and Unger (1986).
- 4 Hawthorne (2003) and Jason Stanley in unpublished work develop such an alternative to contextualism. The main arguments of this paper do not tell against it (except for a potential threat suggested by the final paragraph). For a line of criticism that does apply to both contextualism and such alternatives to it see Williamson (200?).
- 5 See Williamson (2001) in this connection for a way in which contextualism comes closer to scepticism than is often realized.
- 6 A pioneering attempt to use sorites paradoxes link vagueness and context-dependence is Kamp (1981); many authors have subsequently argued for similar claims. For a critique of such views see Stanley (2003).
- 7 See also Hawthorne (2003): 109-10.
- 8 The sentence 'It would not be wrong for me not to resign' is treated as the negation of 'It would be wrong for me not to resign' in order to avoid the clumsy 'It is not the case that it would be wrong for me not to resign'.

- 9 The ascription of truth to sentences involving moral terms such as ‘wrong’ is notoriously controversial but will not be defended here. No metaphysical account of truth is being assumed. Although contextualists about ‘wrong’ must reject a disquotational account of truth for sentences involving the word ‘wrong’, they can still accept that if one says that it would be wrong for Clare not to resign, one speaks truly if and only if it would be wrong for Clare not to resign, and falsely if and only if it would not be wrong for Clare not to resign (since ‘wrong’ is used rather than mentioned in such indirect speech reports). See Feldman (2001: 72-3) for a similar example to illustrate the implausibility of a contextualist treatment of ‘wrong’. Cohen (2001) replies to Feldman
- 10 We can consider such a context even if Clare does not in fact utter the sentence at that time. The convenient simplifying assumptions that Clare is a normal English speaker and that there is only one time when the action can be taken are not essential to the argument.
- 11 For convenience, the individuation of propositions is assumed to be coarse-grained, so that any difference in sense between ‘me’ (as used by Clare) and ‘Clare’ makes no difference to the propositions expressed by containing sentences. Since contextualism is supposed to resolve apparent disagreements by making a difference at the level of reference and truth-value, ignoring finer-grained distinctions is harmless for present purposes.
- 12 Bound variables or complex expressions containing bound variables may replace any of ‘S’, ‘*t*’ and ‘P’ to handle sentences such as ‘The ancient Egyptians knew many truths of mathematics’. In that case, consider *c* and *c** relative to the same assignment of values to variables.
- 13 Sentences are being evaluated with respect to the same possible world, the world of *c*, *c** and the agent’s context.
- 14 Classic examples of epistemological externalism include Armstrong (1973), Dretske (1981), Goldman (1986) and Nozick (1981).

15 For the sake of the example, we may assume that the politician is not using ‘It would be wrong for S to do A’ as equivalent to ‘S would call doing A “wrong”’. No single meta-linguistic reading will make adequate sense of cases in which he must take into account various opponents’ evaluations of another’s actions. In any case, the meta-linguistic reading involves an extreme lack of autonomy for the speaker’s context.

16 The example is of course a variant of Carl Ginet’s fake barns, described in Goldman (1976).

17 For the conception of knowledge as justifying belief see Williamson (2000). On the conception developed there, what justifies belief is the subject’s total evidence, which is simply the subject’s total knowledge; one’s degree of belief in a proposition should be its probability conditional on one’s total evidence. Thus, if one knows that P, one’s degree of belief that P should be 1. However, even if one does not know that P, the probability that P (or even that one knows that P) on one’s total evidence may be close to 1, so that one’s degree of belief should be close to 1. That is consistent with what is said in the text, for ‘degree of belief’ here does not mean degree of outright belief; if one has degree of belief 0.5 or 0.9 in a proposition, one has no outright belief in it whatsoever. To call belief in a proposition that has probability 0.5 on one’s evidence ‘half-justified’ is like calling a glass ‘half-full’; to call the belief ‘justified’ would still be straightforwardly false, just as calling the glass ‘full’ would be straightforwardly false. Awkwardly, since a probability distribution can give probability 1 to a false proposition conditional on some true propositions, even a false proposition may have probability 1 on one’s evidence, so that belief in it might count as justified to degree 1 without counting as justified. Any other theory of real-valued evidential probability is likely to encounter similar awkwardnesses, which arise mathematically from the slightly coarse-grained nature of real-valued measures on infinite probability spaces, independently of the theorist’s epistemological outlook. Discussion with Alexander Bird helped me to clarify my thinking about the relation between knowledge and justification. Sutton (200?) argues in detail that one’s belief that P is justified only if one knows that P. Williamson (2000) also sketches an analogy between justified belief and warranted assertion, and argues that one’s assertion that P is warranted only if one knows that P.

18 The connection between being known and being apt for use as a premise in practical reasoning is

explored in Hawthorne (2003).

19 For simplicity, 'P' is treated as context-invariant.

20 DeRose (200?a) argues for contextualism about 'know' *via* an account that connects the assertibility of 'P' with the truth of 'I know that P' in the context of assertion, although his arguments there do not obviously favour contextualist against the subject-sensitive invariantism suggested by Hawthorne and Stanley (see note 4). DeRose argues elsewhere (200?b) against subject-sensitive invariantism. The arguments of the present paper could be rephrased in terms of assertibility rather than the justifiability of belief.

21 Discussions with John Hawthorne and Jason Stanley played an important role in developing the ideas in this paper. Earlier versions of the material were presented in talks at a summer school on epistemology in Paris and at the following universities: Oxford, Queen's University (Ontario), Cornell (with Sydney Shoemaker as respondent), Nebraska (Lincoln), Minnesota, Manchester, East Anglia, Glasgow, York, Padua and Michigan; I thank the audiences and an anonymous referee for helpful comments.

Bibliography

- Armstrong, D. 1973. *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.
- Cohen, S. 1987. 'Knowledge, context, and social standards'. *Synthese*, 73: 3-26.
- Cohen, S. 1988. 'How to be a fallibilist'. *Philosophical Perspectives*, 2: 91-123.
- Cohen, S. 1999. 'Contextualism, skepticism, and the structure of reasons'. *Philosophical Perspectives*, 13: 57-89.
- Cohen, S. 2001. 'Contextualism defended: Comments on Richard Feldman's "Skeptical problems, contextualist solutions"'. *Philosophical Studies*, 103: 61-85.
- Craig, E. 1990. *Knowledge and the State of Nature: An Essay in Conceptual Synthesis*. Oxford: Clarendon Press.
- DeRose, K. 1992. 'Contextualism and knowledge attributions'. *Philosophy and Phenomenological Research*, 52: 913-29.
- DeRose, K. 1995. 'Solving the skeptical puzzle'. *Philosophical Review*, 104: 1-52.
- DeRose, K. 2007a. 'Assertion, knowledge, and context'. *Philosophical Review*, forthcoming.
- DeRose, K. 2007b. 'The problem with subject-sensitive invariantism'. *Philosophy and Phenomenological Research*, forthcoming.
- Dretske, F. 1970. 'Epistemic operators'. *Journal of Philosophy*, 67: 1007-23.
- Dretske, F. 1981. *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.
- Feldman, R. 2001. 'Skeptical problems, contextualist solutions'. *Philosophical Studies*, 103: 87-98.
- Goldman, A. 1976. 'Discrimination and perceptual knowledge'. *Journal of Philosophy*, 73: 771-91.
- Goldman, A. 1986. *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Hawthorne, J. 2003. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hofweber, T. 1999. 'Contextualism and the meaning-intention problem'. In K. Korta, E. Sosa and X. Arrazola (eds.), *Cognition, Agency and Rationality: Proceedings of the Fifth International Colloquium on Cognitive Science*. Dordrecht: Reidel.
- Kamp, H. 1981. 'The paradox of the heap'. In U. Mönnich (ed.), *Aspects of Philosophical Logic*. Dordrecht: Reidel.
- 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.
- Lewis, D. 1979. 'Scorekeeping in a language game'. *Journal of Philosophical Logic*, 8: 339-59.
- Lewis, D. 1996. 'Elusive knowledge'. *Australasian Journal of Philosophy*, 74: 549-67.
- Nozick, R. 1981. *Philosophical Explanations*. Oxford: Oxford University Press.
- Rysiew, P. 2001. 'The context-sensitivity of knowledge attributions'. *Noûs*, 35: 477-514.
- Schiffer, S. 1996. 'Contextualist solutions to scepticism'. *Proceedings of the Aristotelian Society*, 96: 317-33.
- Stanley, J. 2003. 'Context, interest-relativity and the sorites'. *Analysis*, 63: 269-80.
- Stine, G. 1976. 'Skepticism, relevant alternatives, and deductive closure'. *Philosophical Studies*, 29: 249-61.
- Sutton, J. 200?. 'Stick to what you know'. *Nous*, forthcoming.
- Unger, P. 1986. 'The cone model of knowledge'. *Philosophical Topics*, 14: 125-78.
- Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Williamson, T. 2001. 'Comments on Michael Williams' "Contextualism, externalism and epistemic standards"'. *Philosophical Studies*, 103: 25-33.
- Williamson, T. 200?. 'Contextualism, subject-sensitive invariantism and knowledge of knowledge'. *Philosophical Quarterly*: forthcoming.