

JASON STANLEY

ON THE LINGUISTIC BASIS FOR CONTEXTUALISM*

Contextualism in epistemology is the doctrine that the proposition expressed by a knowledge attribution relative to a context is determined in part by the standards of justification salient in that context. The (non-skeptical) contextualist allows that in some context *c*, a speaker may truly attribute knowledge at a time of a proposition *p* to Hannah, despite her possession of only weak inductive evidence for the truth of that proposition. Relative to another context, someone may make the very same knowledge attribution to Hannah, yet be speaking falsely, because the epistemic standards in that context are higher. The reason this is possible, according to the contextualist, is that the two knowledge attributions express different propositions.

The main advocates of contextualism have used the semantic doctrine of contextualism to defend a certain response to the problem of skepticism. According to it, the force of the skeptical paradoxes is due to the presence of unrecognized context-sensitivity in the language. When we are not discussing with skeptics, many of our ordinary knowledge attributions are true, because what is there at issue is a less demanding sense of knowledge. But the consequence of engagement with the skeptic is that the content of knowledge attributions shifts in a manner that is not recognized by the interlocutors. In such a “skeptical” context, knowledge attributions that may previously have expressed truths now express falsehoods.

The contextualist therefore seeks to explain the force of skeptical arguments by appeal to a feature of ordinary language. It is because knowledge-attributions are context-sensitive that we are fooled by skeptics into thinking that even in non-skeptical circumstances, many of our knowledge attributions are false. The contextualist claim is therefore not revisionary in character. The contextualist does not claim that we *ought* to use “know” as a context-sensitive expression, because that would somehow evade the skeptical para-



doxes; it is not clear what such a position would amount to. Rather, her dissolution of skepticism requires appeal to an *actual* feature of ordinary knowledge attributions. For according to her, the presence of this feature is what explains our actual inclination to accept skeptical arguments.

My purpose in this paper is to investigate whether the features contextualists impute to natural language knowledge ascriptions are in fact all present. My discussion will not be exhaustive; I aim simply to explore some of the apparent linguistic commitments of contextualism. I begin with a brief review of some of the intuitive linguistic evidence for contextualism. On the contextualist view, knowledge relations come in varying degrees of strength, depending upon the epistemic position of the knowledge-attributor. This suggests that the word “know”, like adjectives such as “tall” and “flat”, is *gradable*. My first task is to evaluate the claim that “know” is gradable, and assess the significance of the results. Secondly, I assess an argument for contextualism about “know”, due to Stewart Cohen, which involves the gradable expression “justified”. Third, I explore the background model of context sensitivity assumed by the contextualist. Finally, I discuss whether the contextualist can evade the problematic commitments by appeal to ambiguity or loose talk.

1. MOTIVATING CONTEXTUALISM

As Stewart Cohen and Keith DeRose have emphasized, the most powerful consideration in favor of contextualism is that it explains certain relatively clear intuitions about the truth-conditions of various knowledge-ascriptions. Here is one example of the kind of argument at issue, due to Keith DeRose:

Hannah and her husband are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. But as they drive past the bank, they notice that the lines inside are very long, as they often are on Friday afternoons. Thinking that it isn't very important that their paychecks are deposited right away, Hannah says “I know the bank will be open tomorrow, since I was there just two weeks ago on Saturday morning. So we can deposit them tomorrow morning.” But then Hannah's husband reminds her that a very important bill comes due on Monday, and that they have to have enough money in our account to cover it. He says, “Banks do change their hours. Are you certain

that's not what is going to happen tomorrow?" Hannah concedes, uttering "I guess I don't really know that the bank will be open tomorrow."

This sort of example supports contextualism about knowledge ascriptions, because it suggests that the propositions expressed by one and the same knowledge-attribution may differ with respect to two different contexts of use, even though the two contexts are identical in all respects relevant for fixing the values of obvious indexicals. Here is the example in more detail.

Consider a sentence like (1), as uttered by Hannah in the first situation:

- (1) I know that the bank is open tomorrow morning.

Before she realizes the importance of having a bank account flush with resources by Monday, she utters (1). What she utters expresses a proposition that seems perfectly true. The proposition concerns a particular time, namely the next morning. She is then informed about the pressing need for a full bank account. She then utters:

- (2) I guess I don't really know that the bank is open tomorrow morning.

Again, it looks like Mary has expressed a proposition that seems perfectly true, one that concerns the same time as the proposition expressed by her previous utterance of (1). But (2) looks to be the denial of (1). If we take these intuitions at face-value, we obtain a contradiction.

Here are the two basic options one has to respond to this problem:

- (a) One can reject the semantic significance of one of the two intuitions. For example, one could deny semantic significance to the intuition that the proposition semantically expressed by Hannah's utterance of (1) is true. Alternatively, one could deny semantic significance to the intuition that the proposition semantically expressed by Hannah's utterance of (2) is true.
- (b) One can deny that the proposition expressed by (2) is really the denial of the proposition expressed by (1).

Contextualism about knowledge is a version of strategy (b).

The contextualist maintains that the situation with (1) and (2) is exactly like a situation in which Hannah utters (3) and John utters (4):

- (3) I am tired.
- (4) I am not tired.

There is no threat of contradiction here, because the proposition semantically expressed by John's utterance of (4) is not the denial of the proposition semantically expressed by Hannah's utterance of (3), despite the fact that the same words are used. For the word 'I' is an indexical expression; it contributes potentially different semantic contents to the semantic content of sentences containing it relative to different contexts of use. Similarly, according to the contextualist, the word "know" is an indexical expression. Relative to different contexts of use, it expresses different relations between persons and propositions. So, (2) is not the denial of the proposition semantically expressed by (1), for the same reason that (4) is not the denial of the proposition semantically expressed by (3). The word "know" has a different content in Hannah's utterance of (1) than it does in Hannah's utterance of (2).

In particular, for the contextualist, the word "know" has a content that is a function of the epistemic standards in the context. When Hannah finds out that she must cash her check before the day is out, her evidence must satisfy a higher epistemic standard in order for her to know that the bank is open. The contextualist accounts for this by supposing that the word "know" changes its content in the new context. It expresses a relation that Hannah stands in to a proposition only if her evidence for that proposition satisfies this higher epistemic standard.

For the contextualist, then, knowledge relations come in higher or lower "strengths." Knowledge attributions are thus comparable to context-sensitive gradable adjectives, such as "tall" and "flat". An attribution of tallness is sensitive to a contextually salient scale of height, as is an attribution to flatness. If what is at issue are basketball players, then that brings in one rather high standard for "tall"; if what is at issue are fifth-graders, then that brings in a considerably lower standard for "tall". In this sense, one could speak of tallness relations coming in higher or lower "strengths" as well.

2. IS “KNOW” GRADABLE?

Contextualists typically tell us, when introducing the thesis, that it wouldn't be at all surprising if predicates such as “knows that Bush is president” turned out to be context-sensitive in the ways they describe. After all, we are told, many natural language predicates are context-sensitive. As Stewart Cohen (1999, p. 60) writes:

Many, if not most, predicates in natural language are such that the truth-value of sentences containing them depends on contextually determined standards, e.g. ‘flat’, ‘bald’, ‘rich’, ‘happy’, ‘sad’ . . . These are all predicates that can be satisfied to varying degrees and that can also be satisfied simpliciter. So, e.g., we can talk about one surface being flatter than another and we can talk about a surface being flat simpliciter. For predicates of this kind, context will determine the degree to which the predicate must be satisfied in order for the predicate to apply simpliciter. So the context will determine how flat a surface must be in order to be flat.

There is a great deal of evidence for Cohen's claim that there is a kind of predicate, of which ‘flat’, ‘bald’, ‘rich’, “happy”, and “sad” may occur as a constituent, the semantics of which involve degrees or scales. The reason that contextualists appeal to such predicates is so that, given their frequency in the language, the claim that “know” is a predicate of this kind will be unsurprising. How good is this *prima facie* case for contextualism about “know”?

The predicates mentioned by Cohen – the “kind” of which he speaks in his second to last sentence, are not a disjunctive sort. They are *gradable adjectives*. Most of the gradable adjectives listed by Cohen are (I believe) context-sensitive, such as “flat”, “tall”, and “rich”.¹ In talking about buildings, “is tall” may express a property that it doesn't express when talking about people. Furthermore, this sub-class of gradable adjectives are context-sensitive in just the way that Cohen and DeRose claim that knowledge-ascriptions are. According to Cohen and DeRose, knowledge-ascriptions come in varying degrees of strength. In other words, knowledge-ascriptions are intuitively *gradable*. Contextualists speak, as their theory suggests, of higher and lower standards for knowledge. Gradable adjectives are the model for gradable expressions. It is therefore no surprise that epistemologists since Unger (1975, Chapter 2) and Lewis (1986) have been exploiting the analogy between “know” and context-sensitive gradable adjectives such as “flat” and “tall”. But,

as I will argue in this section, the attempt to treat “know” as a gradable expression fails. This casts suspicion upon the contextualist semantics for knowledge ascriptions. First, it shows that one cannot appeal to the context-sensitivity of words like “tall”, “flat”, and “rich” to justify the context-dependence of knowledge-ascriptions. Secondly, it casts doubt upon the claim that knowledge comes in varying degrees of strength, a core claim of contextualism.

There are two linguistic tests for gradability. First, if an expression is gradable, it should allow for modifiers. For example, predicative uses of comparative adjectives allow for modification, as in:

- (5) a. That is very flat.
- b. That is really flat.
- c. John is very tall.
- d. John is really tall.

Secondly, if an expression is gradable, it should be conceptually related to a natural comparative construction. So, for “flat”, “tall”, and “small” we have “flatter than”, “taller than”, and “smaller than”. Both of these features are to be expected, if underlying the use of the relevant expression is a semantics involving degrees or intervals on a scale. For instance, the semantic effect of a modifier such as “very” on a word like “tall” is to increase the contextually salient degree on the scale of height that an object must exceed in order to satisfy the predicate.²

The claim that knowledge ascriptions are gradable fits elegantly into the contextualist explanation for, say, DeRose’s bank case discussed in the previous section. For in Hannah’s final utterance, she claims that she does not *really* know that the bank is open. It is natural to read “really” here as a degree modifier, as in the examples in (5). That is, it is natural to read this discourse as providing evidence for the gradability of “know”. Over the course of the discourse, Hannah asserts that she knows that the bank is open, but also asserts that she doesn’t *really* know that the bank is open. That is like someone asserting that Bill is tall, but conceding that Bill is not *really* tall.

But the explanation of the bank case that appeals to the gradability of knowledge is not correct. As the above example suggested,

negations of degree-modifier uses of “really” can be conjoined with assertions of the unmodified forms without inconsistency:

- (6) a. John is tall, but not really tall.
 b. Michigan is flat, but not really flat.

In contrast, the same facts do not hold of the use of “really” when appended to a knowledge-ascription (“#” expresses oddity):

- (7) # If the bank is open, then John knows that the bank is open, but doesn’t really know that the bank is open.

The sentences in (6) are perfectly natural. In contrast, (7) is very odd. This suggests that the “really” that occurs in the above description of the bank case is not a degree modifier.³

Indeed, *prima facie*, propositional knowledge ascriptions are not gradable. First, knowledge ascriptions do not seem to allow for modification:

- (8) a. *John very knows that penguins waddle.
 b. *John knows very much that penguins waddle.

Second, there is no natural comparative conceptually related to “know”. The following locutions are deeply strained:

- (9) a. ??John knows that Bush is president more than Sally knows it.
 b. ??Hannah knows that Bush is president more than she knows that Clinton was president.

If the semantics of “know” did involve scales, it would be mystery why there wouldn’t be a comparative form of “know” available to exploit the scale.

It has been noted before that ascriptions of propositional knowledge are not gradable (cf. Dretske, 1981, Chapter 5). However, the data surrounding knowledge ascriptions is more complex than these *prima facie* considerations suggest. There are several constructions that suggest that knowledge ascriptions are, despite initial appearances, gradable. In the remainder of this section, I provide a complete case that knowledge ascriptions are not gradable, and draw some morals for contextualism about “know”.

One might think that knowledge-ascriptions are gradable on the basis of the obvious felicity of the following sort of construction:

- (10) a. John knows Bill better than Mary does.

b. Hannah knows logic better than John does.

But in the sentences in (10), “know” does not express a relation between a person and a proposition. These sentences are not propositional knowledge ascriptions; rather the occurrences of “know” in them express the acquaintance relation, what would be expressed in German by “kennen” rather than “wissen”. It is only the gradeability of propositional knowledge ascriptions that is the issue in contextualism in epistemology.

However, “know” can marginally occur with “very much” or less marginally with “very well”, as in:

- (11) a. ?I very much know that Bush is president.
 b. I know very well that Bush is president.

But it is doubtful that these occurrences of “very much” and “very well” are genuine semantic modifiers of the knowing relation, rather than pragmatic indicators. In this sense, these constructions are similar to:

- (12) 2 is very much an even number.

Decisive evidence for this comes from several sources. First, note the unacceptability of negating the constructions in (13):

- (13) a. *I don’t know very much that Bush is president.
 b. *I don’t know very well that Bush is president.

The unacceptability of the sentences in (13) contrasts with the naturalness of negating the verb phrase in a case in which “very much” is clearly modifying the verb:

- (14) I don’t like Bill very much.

Secondly, “know” is only with great awkwardness combined with “very well” in non-assertoric speech acts. Contrast the sentences in (15) with (16):

- (15) a. ??Do you know very well that Bush is president?
 b. *Do you know very much that Bush is president?
 (16) Do you like Bush very much?

So, the sentences in (11) are clearly not cases where the degree of knowing is operated on by “very much” or “very well”.

Defenders of contextualism might hold that “better than” rather than “more” is the comparative relevant to “know”, as in:

- (17) Hannah knows better than anyone that she is poor.

But here again, the construction means that Hannah is familiar with the fact more than anyone else – e.g. she lives with the consequences.⁴ More importantly, “better than anyone” is idiomatic. For example, consider the oddity of:

- (18) a. ??Hannah knows better than three people that she is poor.
 b. *Hannah doesn't know better than anyone that she is poor.

So, “better than anyone” is an idiomatic construction, one from which we can infer little about the semantics of “know”.⁵

Furthermore, none of the non-philosopher informants I asked found the following acceptable, though they disagreed amongst themselves which was worst:

- (19) a. ??John knows that Bush is president better than Mary does.
 b. ??John knows that Bush is president better than Bill knows that Clinton is a Democrat.

Furthermore, all of my informants reported a strong difference in acceptability between these sentences, on the one hand, and the perfectly acceptable:

- (20) a. John likes Bill more than Mary does.
 b. John likes Bill more than Mary likes John.

So “better than” is not a natural way to express comparisons between levels of epistemic position with “know”. If the semantics of “know” did involve scales of epistemic strength, then there should be uncontroversially non-idiomatic comparisons and modifications.⁶

One might think that these facts about ‘know’ have syntactic rather than semantic explanations. Perhaps sentences like (8) and (9) and those in (13) are deviant because verbs that take sentential complements grammatically do not allow for comparisons or intensifiers. But consider “regret”, a factive verb in the same syntactic category as “know”:

- (21) a. Hannah very much regrets that she is unemployed.

- b. Hannah doesn't regret very much that she is unemployed.
- c. Hannah regrets very much that she is unemployed.
- d. Hannah regrets that she is unemployed very much.

Here, the degree of regret clearly seems to be modified by "very much". Furthermore, "regret" easily allows for comparisons:

- (22) Hannah regrets that she is unemployed more than she regrets that she is unpopular.

This shows that the lack of straightforward comparatives or degree modifiers has nothing to do with the syntax, or even the factivity, of "know". There are syntactically similar expressions whose link to degrees and scales is far more plausible.⁷

It is also worth mentioning that other expressions upon which one might be tempted to base the context-sensitivity of "know" are, unlike "know", also gradable. So consider *epistemic modals*.⁸ A contextualist might understandably wish to appeal to the apparent context-sensitivity of epistemic modals to justify the apparent context-sensitivity of knowledge-ascriptions. But one problem with this strategy is that epistemic modals, unlike knowledge-ascriptions, are intuitively gradable:

- (23) a. It is very likely that I will publish more papers on this topic.
 b. It is more possible that Hannah will become a philosopher than it is that she will become a mathematician.

So the analogy between epistemic modals and "know" is almost as strained as the analogy between context-sensitive gradable adjectives and "know".

The evidence concerning gradeability is more complicated when one considers the deverbal adjective "known". But even here, there does not appear to be a good case for a semantics involving a scale of *epistemic strength*. This adjective, unlike its verbal relative, does give rise to comparisons and modifications. But they are not of the relevant sort. So, for example, consider:

- (24) a. That broccoli is low-fat is better known than that broccoli prevents cancer.
 b. That broccoli is low-fat is well known.

(24a) does not mean that there is more evidence that broccoli is low-fat than that broccoli prevents cancer; rather, it means that the fact that broccoli is low-fat is more *widely known* than the fact that broccoli prevents cancer. Similarly, (24b) means, not that there is a lot of evidence that broccoli is low-fat, but that it is *widely known* that broccoli is low-fat. Evidence for this hypothesis comes from the fact that while (25a) is perfectly acceptable, (25b) sounds quite odd:

- (25) a. That broccoli prevents the flu is well-known, but ill-understood.
 b. ?That broccoli prevents the flu is well-known, though few people know it.

Furthermore, as Tamar Gendler has pointed out to me, instances of (26) are quite odd:

- (26) It is well known that p, and less well-known that q, but more people know that q than know that p.

This data is explicable on the assumption that the only available reading for “well-known” is widely known. So, while the data is more complex here, the adjectival relative of ‘know’, on the rare use of it where it expresses propositional knowledge, does not to be an obvious candidate for analysis via degrees on a scale of epistemic strength.

Another potential source of evidence for the gradability of “know” comes from its use in certain kinds of embedded questions. Consider, for example:

- (27) a. John knows how to swim well.
 b. John knows how to ride a bicycle better than Mary does.⁹
 c. Hannah knows where Texas is better than John does.

It is quite plausible that these are attributions of propositional knowledge.¹⁰ If so, one might think that this suggests that knowledge is gradable after all. However, in these cases, what is being compared are *answers to questions*. In each case, one person is said to have a better answer to a certain question than another; the answer Hannah has to the question “Where is Texas?” is better, or more complete, than the answer John has.¹¹ So, embedded questions do not provide evidence for the gradeability of knowledge claims.

Contextualists often write as if the *prima facie* case for the contextual sensitivity of “know” is strong, citing the relatively uncontroversial context-sensitivity of adjectives such as “tall”, “flat”, and “small”. Their purpose in so doing is to shift the burden of proof from their shoulders to their opponents; if “many if not most predicates of natural language are context-dependent”, then someone who claims that “know” is context-dependent does not suffer from a large burden of proof. My point in this section has been to emphasize that these arguments do not suffice to lift the burden of proof from the shoulders of the contextualist about “know”. The fact that the semantic contents of a subclass of gradable adjectives are sensitive to contextually salient standards is irrelevant to the claim that “know” has a similar context-sensitive semantics.

In fact, we may draw a stronger conclusion from the above discussion. Natural language expressions that are semantically linked to degrees on scales exploit this link in a variety of recognizable ways – by allowing for comparisons between degrees on the scale, and by allowing modifications of the contextually salient degree on the scale. If the semantic content of “know” were sensitive to contextually salient standards, and hence linked to a scale of epistemic strength (as “tall” is linked to a scale of height), then we should expect this link to be exploited in a host of different constructions, such as natural comparatives. The fact that we do not see such behavior should make us at the very least suspicious of the claim of such a semantic link. Thus, an investigation into the context-sensitivity of predicates such as “is tall”, “is small”, and their ilk adds to, rather than removes, the burden of proof on contextualists about “know”.

3. THE COHEN ARGUMENT FOR CONTEXTUALISM

Propositional knowledge ascriptions are not gradable. But some contextualists have claimed that it does not matter to contextualism whether or not they are.

Does knowledge come in degrees? Most people say no (though David Lewis says yes). But it doesn't really matter. For, on my view, justification, or having good reasons, is a component of knowledge, and justification certainly comes in degrees. So context will determine how justified a belief must be in order to be justified simpliciter.

This suggests a further argument for the truth of the contextualist's claim about knowledge. Since justification is a component of knowledge, an ascription of knowledge involves an ascription of justification. And for the reasons just indicated, ascriptions of justification are context-sensitive. (Cohen, 1999, p. 60)

In this section, I will argue that Cohen has not freed the contextualist from the thesis that "know" comes in degrees. Then, I will argue that Cohen's attempt to derive the context-sensitivity of "know" from the context sensitivity of "justified" fails.

Cohen's argument for the context-sensitivity of "know" has roughly the following structure:

COHEN

Premise 1. Gradable expressions are context-sensitive.

Premise 2. So, we can expect the gradable term "justified" to be context-sensitive.

Premise 3. "S knows P" means in part what is meant by "P is justified for S"

Conclusion. So the truth of "S knows P" depends upon context.

If this argument is sound, then it is difficult to see how Cohen is justified in claiming that "it doesn't really matter" whether "know" comes in degrees. For if this argument is sound, then surely so is:

*COHEN**

Premise 1. Gradable expressions are sensitive to contextually salient scales.

Premise 2. So, we can expect the gradable term "justified" to be sensitive to contextually salient scales.

Premise 3. "S knows P" means in part what is meant by "P is justified for S"

Conclusion. So the truth of "S knows P" is sensitive to contextually salient scales.

So it seems that Cohen has not discovered a way of freeing the contextualist from the view that "know" comes in degrees.

However, a more pressing task is to evaluate the soundness of Cohen's argument from the gradability of "justified" to the conclusion that "know" is context-sensitive.

The word "justified" is certainly gradable. We may say "My belief that Bush is a Republican is more justified than my belief that

I will eat Indian food tomorrow.” It does not, however, follow that “justified” is context-sensitive. Some expressions are gradable, yet not context-sensitive. Consider, for example, the expression “taller than six feet.” This expression is clearly not context-sensitive. Yet it is gradable; someone may be *much* taller than six feet. So, the expression “taller than six feet” is gradable, since it can be modified by a degree word such as “much”. But it is not context-sensitive. So gradability does not entail context-sensitivity, and premise 1 is false.

So the conclusion that “justified” is context-sensitive does not follow from the premise that it is gradable. My own view is that “justified” is gradable, but not context-sensitive, much like the expression “taller than six feet.” For a belief to be justified means that it is justified over the context-invariant degree of justification, just as to be taller than six feet tall is to be taller than the context-invariant degree of height of six feet.

The second crucial step in Cohen’s argument is premise 3. The claim accepted by most epistemologists is that, in the analysis of the knowledge relation, some justification property appears. That is:

The Knowledge-Justification Thesis

Some justification property is part of the conceptual analysis of the knowledge relation.

The Knowledge-Justification Thesis is an informal way of making explicit the thesis that knowledge is analyzed in terms of justification.

However, if premise 3 is another way of stating the Knowledge-Justification Thesis, then COHEN is not valid. For the Knowledge-Justification Thesis is not a claim about a relation between the word “know” and the word “justified.” It is rather the claim about the relation between the knowledge relation and some justification property. From such a claim, coupled with the thesis, that I will for the moment grant for the sake of argument, that the *word* “justified” is context-sensitive, nothing follows about the context-sensitivity of the word “know.”

One might think that there is *prima facie* case to be made, from the fact that a certain term *t* contains in the analysis of what it expresses a property that is expressed by a context-sensitive term *t'*, that *t* is therefore context-sensitive. But this does not seem in general to be true. Consider the term “vacuum”. A plausible analysis of the

notion of being a vacuum involves being completely empty. But the property of being empty is expressed by the context-sensitive word “empty.” This does not entail that “vacuum” is context-sensitive.¹² Another example is the expression “John’s enemy.” There is at least a prima facie case that this expression is not context-sensitive.¹³ But analyzing the notion of being John’s enemy involves appealing to the notion of being an enemy, which is expressed by the context-sensitive word “enemy” (in one context, it may mean an enemy of x, and in another context, an enemy of y).

So, if premise 3 is supposed to reflect the Knowledge-Justification Thesis, it does not allow Cohen to draw the conclusion that “know” is context-sensitive from the premise that “justified” is context-sensitive. One might think that, intuitively, an attribution of knowledge entails an attribution of justification. But even if this were true, it would not obviously help Cohen. For note that the following thesis is not sufficient for Cohen’s purposes:

For any context c, if what is expressed in c by a sentence of the form ‘x knows that p’ is true, then what is expressed in c by a sentence for the form ‘x is justified in believing that p’ is true.

This Thesis, together with Premises 1 and 2 of COHEN, does not entail that “know” is context-sensitive (granting the context-sensitivity of “justified”). For it may be (and is, according to the skeptic) that knowledge entails the highest kind of justification. If so, a true knowledge-ascription will always entail a true justification-ascription, even though “know” is not context-sensitive.

The thesis Cohen requires to move from the context-sensitivity of “justified” to the context-sensitivity of “know” is rather the following:

The Metalinguistic Knowledge-Justification Thesis

For any context c, the word “know” expresses a relation that, relative to that context, contains as a component the property expressed by the word “justified”, relative to c.

However, it is unclear how to argue for the Metalinguistic Knowledge-Justification Thesis. It does not appear to be a reformulation of an innocuous epistemological thesis. Furthermore, there are special reasons to think that the contextualist is in a difficult

position with respect to providing an argument for this thesis. For the contextualist account of skepticism involves the hypothesis that we are often mistaken about the semantically relevant features of the context. So the contextualist cannot argue for the Metalinguistic Knowledge-Justification Thesis by appeal to intuitions about cases.

So, there is no clearly persuasive route from the fact that “justified” is a gradable adjective to the context-sensitivity of “know”. First, the fact that “justified” is gradable does not entail that it is context-sensitive. Indeed, it is perfectly plausible to hold that “justified” is gradable and not context-sensitive. Secondly, even if “justified” is context-sensitive, and knowledge is analyzed in terms of justification, it does not follow that “know” is context-sensitive.¹⁴

4. CONTEXTUALISM AND THE DISCOURSE MODEL OF CONTEXT-DEPENDENCE

Contextualists typically speak as if there is one contextual standard in a context for all context-sensitive expressions in a discourse. So, for example, DeRose regularly speaks of “how high or low the standards for knowledge are set” in a conversational situation (1995). But this is not in general a good description of how context-sensitive expressions work. Rather, the context-sensitivity is usually linked to the term itself, rather than the whole discourse. If the discourse model of context-dependence is incorrect, this raises serious worries for the contextualist. In this section, I give some reasons to think that the discourse model is incorrect, and explain the problems that this raises for the contextualist.

In the case of virtually all terms that occur in semantically context-dependent constructions, the context-sensitivity is linked either to the term, or some element associated with that term. Let us begin with the case of context-sensitive gradable adjectives. Suppose John, who is very small for his age, identifies with small things. He has a picture on the wall in his bedroom of an elephant fighting off a much larger elephant. He also has a framed tiny butterfly on his wall. When he is asked why he has both things hung up, he says:

(28) That butterfly is small, and that elephant is small.

John in fact also has a fondness for flat things. On his wall is a picture of a field in Kansas, and on his desk is a rock. When asked why he has both, he replies:

(29) That field is flat, and this rock is flat.

Now imagine a picture of a butterfly that's surrounded by much smaller butterflies; it's huge for a butterfly. It's next to a picture of an elephant that's surrounded by much larger elephants. The following is a good description of the situation:

(30) That butterfly is large, but that elephant isn't large.

This sort of contextual shift is present in a variety of context-sensitive expressions other than gradable adjectives. For example, we see this behavior with demonstratives, context-sensitive determiners, quantified noun phrases, and modal auxiliaries. Here is a case with demonstratives (imagine two different demonstrations):

(31) This is larger than this.

Here is a case for the context-sensitive determiner "many":

(32) In Syracuse, there are many serial killers and many unemployed men.

In this case, the contextual determinants for the denotation of 'many' change within a clause. What counts as "many" for serial killers is not what counts as "many" for unemployed men. The same phenomenon occurs with quantified expressions, as (33) can express the proposition that every sailor on one ship waved to every sailor on another (Stanley and Williamson, 1995, p. 294):

(33) Every sailor waved to every sailor.

The same phenomenon occurs with modal auxiliaries, which can change interpretation within a clause:

(34) a. I couldn't bench press 100 pounds, though with regular physical training, I could.
b. John can lift the block that Mary says he can.

In (34a), the two different uses of "could" express a narrower and wider sense of physical possibility, respectively. (34b) has a reading according to which it expresses the proposition that John is physically able to lift the weight that Mary says he is permitted

to lift. That is, (34b) allows a reading in which the first “can” is physical possibility, and the second is permissibility. For all of these categories of expression, then, the discourse model of context-sensitivity is incorrect.

With a host of other context-dependent words, one detects similar behavior. Consider, for example, the word “nearby”. For a person with a car, a place can be nearby, that isn’t nearby for a person without such a rapid mode of transportation. Indeed, the following expresses a truth (considered as uttered in Ann Arbor):

- (35) If you have a car, Detroit is nearby, but if you’re on foot, Detroit isn’t nearby.

Similarly, consider a relational noun such as “enemy”; to be an enemy is to be an enemy of *x*, for some *x*. Now consider:

- (36) John has an enemy, and Bill has an enemy.

In the case of (36), what is said is that John has an enemy of John, and Bill has an enemy of Bill. So, in the case of both “nearby” and “enemy”, the interpretation of the relevant contextual parameter can change within a sentence. For such expressions too, the discourse model of context-dependency is incorrect.¹⁵

It is no surprise that context-sensitive expressions typically allow for standard-shifts within a clause. In each case, the context-sensitivity is linked not to the *discourse*, but to the particular context-sensitive *term*. So what one is speaking about when one speaks of the “standard of tallness” relevant for evaluating a particular use of “is tall” is simply the degree of tallness that is associated with the expression “tall” (or some element closely associated with it) by whatever semantic mechanism one exploits.

The discourse model of context-sensitivity is particularly inadequate in the case of the version of contextualism advocated in Lewis (1996). According to Lewis, the semantics of the word “know” invokes universal quantification over possibilities. Lewis then exploits facts about natural language universal quantification to motivate contextualism about “know”:

Finally, we must attend to the word “every”. What does it mean to say that every possibility in which not-P is not eliminated? An idiom of quantification, like “every”, is normally restricted to some limited domain. If I say that every glass is empty, so it’s time for another round, doubtless I and my audience are ignoring

most of all the glasses there are in the whole wide world throughout all of time. They are outside the domain. They are irrelevant to the truth of what was said.

Likewise, if I say that every uneliminated possibility is one in which P, or words to that effect, I am doubtless ignoring some of all the uneliminated alternative possibilities that there are. They are outside the domain. They are irrelevant to the truth of what was said. (*Ibid.*, p. 553)

So, Lewis deduces contextualism about “know” first from the claim that “know” involves universal quantification over possibilities, and secondly from the fact that natural language quantification is typically restricted. But (as (33) demonstrated) it is a well-established fact that different occurrences of the same quantified expression can be associated with different domains (Soames, 1986, p. 357; Stanley and Williamson, 1995, p. 294; Stanley and Szabo, 2000, p. 249). Given that Lewis’s contextualism flows from facts about natural language quantification, it should follow that two different occurrences of “know” within the same discourse should be able to be associated with different sets of possibilities (say, a set including quite remote possibilities, and a set only including quite close possibilities).

Why is it important that contextualists are operating with the wrong model of context-sensitivity? First, if this is the incorrect model for context-sensitivity, then it opens the contextualist up to a number of objections that she does not otherwise face. Secondly, if this is the incorrect model for context-sensitivity, some of what contextualists say about the virtues of their theories over other theories falls by the wayside. I will substantiate these points in turn.

If “know” behaves like context-sensitive gradable adjectives, quantifier phrases, context-sensitive determiners, or modals, then we would expect it to be smoothly acceptable to associate different standards of knowledge with different occurrences of “know”, just as we associate different degrees of height with different occurrences of “tall”. So, if contextualism were true, we should expect the following to be fine:

- (37) If there is an external world, many normal non-philosophers know that there is, but, by contrast, no epistemologists know that there is.

If “know” could have different standards, then one would expect an utterance of (37) to felicitously assertible and true, just as utterances

of (28)–(36) are. For “large” means one thing when predicated of a flea, and quite another when predicated of an elephant, and “many” means one thing when it occurs with “unemployed men”, and quite another when it occurs with “serial killers”. So, if “know” is context-sensitive in a similar manner, one would naturally expect “know” to have one content when predicated of non-epistemologists, and another when predicated of epistemologists.

The fact that the contextualist is operating with the incorrect model of context-sensitivity opens them up to objections that they would not otherwise face. But the situation is worse for the contextualist. If the model of context-sensitivity assumed by the contextualist is wrong, a great deal of what contextualists say about the virtues of their theories over other theories is vitiated.

If shifts in standards for knowledge can occur with a discourse, some of the paradigm sentences the infelicity of which supposedly motivates their accounts over rival accounts turn out to be felicitous and potentially true by contextualist lights. For example, if “know” can, like other context-sensitive expressions, change its interpretation within a clause, one would expect the following to be felicitous:

- (38) a. If Bill has hands, then Bill knows that he has hands, but Bill does not know that he is not a bodiless brain in a vat.
- b. If Bill has hands, Bill does not know that he is not a bodiless brain in a vat, but Bill knows he has hands.

Keith DeRose calls sentences such as these “Abominable Conjunctions” (1995, pp. 27–29), and takes the infelicity of utterances of the sentences in (38) to undermine Robert Nozick’s account of skepticism on the grounds that it allows for acceptable utterances of these sentences. But if “know” is context-sensitive, then it would be mysterious why the sentences in (38) wouldn’t be smoothly acceptable. Felicitous utterances of these sentences would simply involve “know” shifting its standards across a conjunction.

Contextualists do in general claim that there are special rules governing the context-sensitivity of “know.” In particular, once a skeptical possibility has been raised, they say, that has ramifications for the evaluation of future uses of “know” within that discourse. In particular, once standards have been raised, it is not possible to lower

them again.¹⁶ But (37) and (38a) involve standard shifting within a sentence that is tantamount to the raising of standards. So appeal to this feature of contextualism does not help the contextualist.

5. AMBIGUITY AND LOOSE USE

It is worthwhile briefly considering two options that may seem open to the contextualist; *ambiguity* and *loose use*. Contextualists such as Cohen, DeRose, and Lewis have, to their credit, not presented their views in these terms. In this section, I briefly explain some of the reasons why they are right not to have so presented their views.

There is really not much to be said for a version of contextualism according to which “know” is ambiguous. To ground an ambiguity claim linguistically, one would need to show that there are languages in which the different meanings are represented by different words. But such a claim is unlikely to be substantiated in the cases of interest to the contextualist. It is quite unlikely that there are languages in which the word for “know” in the epistemology classroom is different than the word for “know” outside the epistemology classroom (of course, stranger things have happened ...).

The more interesting case involves loose use. Suppose a contextualist were to present her claims in these terms. Then, on this view, when “standards are low”, one is using “know” loosely, and this is why it is acceptable to predicate a knowledge-state of someone who only has weak inductive evidence for the truth of one of her (true) beliefs. When “standards are high”, then “know” is used strictly, and it is not acceptable to attribute knowledge to that person in that situation.

Here are some typical cases of loose use. One may utter “France is hexagonal”, to describe the rough shape of France. Though intuitively France is not “strictly” hexagonal, such an utterance can be felicitously made. Similarly, suppose someone utters “It’s three o’clock” when asked the time at 3:03. Her response can be felicitous, even though it is not intuitively “strictly” true. Finally, suppose that in a village, there is a person Mary with a small amount of medical training, and no one else has any medical training at all. In context, it may be appropriate to utter “Mary is the doctor of the

village”, though it is intuitively not strictly true. These are all cases in which, intuitively, language is used “loosely”.¹⁷

There are two kinds of accounts one could envisage of such cases. According to the first, and by far the most natural, utterances of such sentences express propositions that are literally false, but the utterances are nevertheless pragmatically acceptable. On the second account, utterances of such sentences express propositions that are literally true.

Suppose the contextualist took the first account, according to which loose use involved the expression of literally false propositions despite pragmatic felicity. On such an account, ordinary knowledge attributions that are generally held to be true would be pragmatically felicitous, but would express false propositions. Only when someone satisfies strict epistemological standards could an atomic knowledge attribution to that person be true. On this view, we generally speak falsely when we attribute knowledge to our fellow epistemic citizens, though we are not thereby violating conversational norms.

But this is obviously an unsatisfactory position way to present the contextualist view. For this position is simply *epistemological skepticism*. Ordinary knowledge attributions are generally false. But skepticism was precisely the view that the contextualist position was supposed to help us evade. Therefore, the contextualist position cannot be presented as the view that knowledge-ascriptions are used loosely, where loose use is understood on the pragmatic model.

Therefore, if contextualism about knowledge-attributions is supposed to be understood on the model of loose use, then its proponents must have a semantic account of loose use. I do not pretend to understand what would motivate a semantic account of loose use. After all, the common feature to cases of loose use is that they seem to be cases in which it is felicitous to describe something as satisfying a predicate, despite the fact that the thing in question does not satisfy the predicate, when it is taken literally. A semantic account of loose use is *prima facie* odd, since one might think that in order for something to satisfy a predicate semantically, it must at least satisfy the literal meaning of that predicate. To attempt to capture loose use semantically seems therefore to miss the phenomenon being described.

Be that as it may, suppose one adopted a semantic account of loose use, and couched contextualism about “know” as the thesis that “know”, like paradigm cases of loose use, can be used loosely. Then an ordinary atomic knowledge attribution to someone with only weak inductive evidence for one of her true beliefs can express a truth proposition, because it is used loosely, and loose talk is a semantic, rather than a solely pragmatic phenomenon.

But what would it mean to present contextualism along the lines of a semantic account of loose use? It is not sufficient to say that a semantic account of loose use is one according to which “know” expresses different properties in different contexts. For this is true of any context-dependent account. Constructions involving context-sensitive gradable adjectives such as “tall” or “flat” involve some expression having different contents in different contexts, but this does not mean that they are best modeled by appeal to “loose use.” It is not that “tall” is being used loosely, as in the paradigm cases above, when an eight year old child who is five feet tall is described as “tall.” It is rather that “tall” (or some expression associated with it) can have different contents in different contexts, consistent with it being used perfectly literally. A contextualist theory of knowledge attributions modeled along the lines of loose use (construed semantically) is therefore not just a view according to which knowledge attributions can change their contents across contexts. It must rather flow from some rather tight analogy between paradigm cases of loose use and knowledge attributions. On this account, it is because knowledge-attributions in ordinary contexts are *like* uttering at 3:03 p.m. the sentence “It is three o’clock” that they can be truthfully asserted. The point of the analogy between knowledge-attributions and loose talk would then be to free the contextualist from having to relate knowledge attributions to constructions such as comparative adjectives and other context-sensitive constructions that (unlike “know”) are gradable or can engage in sentence-internal context-shifts.

So, a contextualist “loose talk” account of knowledge is one according to which knowledge-attributions are context-sensitive “in the same sense” as it is permissible to describe 3:03 p.m. as satisfying the predicate “is three o’clock”, and France as satisfying the predicate “is hexagonal.” But now it seems that this version

of contextualism would face a similar charge as the pragmatic loose talk version of contextualism. According to this version of contextualism, Hannah only can satisfy the predicate “knows that she has hands” in the sense that France can satisfy the predicate “is hexagonal”, 3:03 p.m. can satisfy the predicate “is three o’clock”, and someone with rudimentary medical training can satisfy the predicate “is a doctor.” This is not a very satisfying way of “rescuing” ordinary knowledge-attributions. Indeed, one may wonder whether it has any advantages over skepticism at all.

So, contextualists have in general been correct not to employ either the ambiguity or “loose talk” versions of their theses. The former seems subject to empirical refutation, and the latter seems no different than the skeptical positions they rightly reject.¹⁸

6. CONCLUSION

In this paper, I have raised some worries that militate against the postulation of semantic context-dependence in ascriptions of propositional knowledge. First, propositional knowledge ascriptions are not gradable. Secondly, the general model for most context-sensitive expressions is that the relevant context-sensitivity is linked to a term, and not the sentence or the discourse. This feature is lacking in the case of the alleged context-sensitivity of knowledge-ascriptions. Finally, I have discussed alternative strategies to establishing the context-sensitivity of knowledge-ascriptions; first, Cohen’s argument from the alleged context-sensitivity of “justified” to the context-sensitivity of knowledge-ascriptions, and secondly the strategy of modeling the alleged context-sensitivity of knowledge-ascriptions upon cases of loose use. In each case, I have found the relevant strategy problematic.

But if ascriptions of propositional knowledge are not semantically context-sensitive, how might one capture the intuitive evidence that Cohen, DeRose, Lewis, and others have brought to our attention? My own favored strategy is to employ an account of knowledge according to which a subject’s knowledge can vary over time, as a function of the changing interests of the subject (see Stanley, forthcoming). Such a view is not contextualist, because the interests of the knowledge-attributor are irrelevant to the truth

of the knowledge-claim (unless the attributor is the subject of the ascription).¹⁹ This view has the virtue of providing an explanation of how changing goals and interests can result in fluctuations in what we know, without the problematic linguistic commitments of contextualism.

NOTES

* The paper I gave at the conference has subsequently split into two papers. The other descendant of the original paper (Stanley, forthcoming) focuses on developing a non-contextualist account of knowledge that captures the intuitive data as well as contextualism. Discussion with the participants at the conference at the University of Massachusetts was very helpful. I should single out John Hawthorne and my commentator Barbara Partee for special mention; e-mails with Stewart Cohen since then have also been invaluable. I am also indebted for discussion to Herman Cappelen, Keith DeRose, Tamar Gendler, Richard Heck, Jim Joyce, Chris Kennedy, Jeffrey King, Ernie Lepore, Peter Ludlow, Robert Stalnaker, and Timothy Williamson.

¹ I, along with many philosophers and linguists, am convinced that they are context-sensitive. Herman Cappelen and Ernie Lepore have, however, argued that they are not context sensitive, and that there is just one property e.g. of smallness, that both buildings and people can have. Ken Taylor (p.c.) also has this view.

² Another feature of gradable expressions is that they can typically occur with measure phrases. For instance, “tall”, “wide”, and “old”, co-occur with measure phrases, as in “Five feet tall”, “two feet wide”, and “30 years old”. There is no natural measure phrase with ‘know’, even though it doesn’t seem to be like the ‘negative’ comparative adjective, like ‘flat’, ‘small’, and ‘young’. I will not pursue this disanalogy in what follows.

³ It appears that the use of “really” in the description of the bank case is rather a *hedge*, in the sense of Paul Kay (1997). A hedge is some expression the linguistic function of which is to comment on the appropriateness of asserting the embedded sentence (as in uses of meta-linguistic negation such as “John isn’t happy, he’s ecstatic”). One such hedge, according to Kay (*ibid.*, pp. 140ff.), is the expression “technically”, as in “Technically, that isn’t water.” The occurrence of “really” in the bank case appears to be a hedge in this sense – at the end of the discourse Hannah is conceding the infelicity of asserting that she knows that the bank is open. Note that this is consistent with it being perfectly true throughout that she knows that the bank is open.

⁴ As Fred Dretske (2000, note 1) writes, concerning such examples, “I take such constructions to be describing not better knowledge, but more direct, more compelling kinds of evidence.” Of course, if evidence is just knowledge, as Williamson (2000, Chapter 9) has argued, Dretske’s explanation is less compelling.

⁵ As Nunberg, Sag, and Wasow (1994) emphasize, idiomaticity is a matter of degree, ranging from completely frozen idioms such as “kick the bucket” to somewhat less frozen idioms. “Knows better than anyone” is intermediate on the scale of idiomaticity.

⁶ Lewis (1996, p. 562), assuming that “know” is gradable, writes: “Take the far-fetched possibility that Possum has somehow managed to get into a closed drawer of the desk – maybe he jumped in when it was open, then I closed it without noticing him. That possibility could be eliminated by opening the drawer and making a thorough examination. But if uneliminated, it may nevertheless be ignored, and in many contexts that ignoring would be proper. If I look all around the study, but without checking the closed drawers of the desk, I may truly be said to know that Possum is not in the study . . . But if I did check all the closed drawers, then I would know *better* that Possum is not in the study.” Lewis does not seem to notice that the last sentence is in fact not grammatical.

⁷ Even though “regret” is gradable, it is not clear that it is context-sensitive. I see no reason to deny that regretting a proposition involves having over a certain context-invariant level of regret (though it is of course vague what this degree is). What this suggests is that the fact that an expression is gradable is independent of the question of whether it is context-sensitive (see below).

⁸ Thanks to John Hawthorne and Angelika Kratzer for pressing me on the analogy between epistemic modals and “know”.

⁹ I owe example (27a) to Jeff King and (27b) to Jamie Tappenden.

¹⁰ See Stanley and Williamson (2001).

¹¹ In the case of (27a), the “well” modifies the swimming, rather than the knowledge. The claim can be represented “John knows [how to swim well].” Sentences such as “John knows well how to swim”, where “well” is made to apply to the knowledge relation, are clearly not felicitous. Finally, “well”, in such uses, is not a degree modifier.

¹² Of course it is true that “vacuum”, like all words, can be used loosely. This is of no use to the contextualism (see section 4).

¹³ When a possessive is used with a relational noun such as “enemy”, it forces a reading in which the possession relation is determined by the nature of the noun (e.g. “John’s brother”). There are uses of such constructions in which the possession relation can be assigned a different reading (suppose John is one of a group of police officers each of whom is interviewing one of the four brothers’ of the suspect in a crime). But it is not clear that such cases are genuine readings of the possessive, rather than deferred reference of some kind.

¹⁴ I am especially indebted to numerous discussions with Stewart Cohen in writing this section.

¹⁵ Thanks to David Chalmers for pushing me to consider “nearby”, and Barbara Partee for pushing me to consider “enemy”.

¹⁶ I argue in my (forthcoming) that this feature of the contextualist view is considerably more ad hoc than has been previously recognized.

¹⁷ My own suspicion is that there is no uniform phenomenon behind “loose use”.

¹⁸ There are various other models and analogies I have not considered. For example, Stewart Cohen (p.c.) has suggested to me that “hit” is context-sensitive; in a society in which people regularly have strong friendly physical contact, one might have to strike someone with greater force than one actually does in order to count as “hitting” that person. So “hit” is context-sensitive. However, I do not find this persuasive. First, hitting is not just about the force of the contact; it is also about the intention behind the contact. It is not clear the force is all that relevant to whether or not a strike counts as a hit. Secondly, and more importantly, the context-sensitivity of a term is not well-established by appealing to societies with different customs. For one runs the risk of thereby describing a society in which that term has a slightly different meaning. A plausible account of Cohen’s example is that “hit” means something slightly different in the society he envisages. To model the alleged context-sensitivity of “know” upon this kind of possible variation for “hit” would then be to endorse the ambiguity model.

¹⁹ For discussion of the distinctions between contextualism and this sort of view, see also Cohen (1991), DeRose (1999) (who calls the latter sort of view “subject contextualism”), and Hawthorne (forthcoming, Chapter 4), who calls it “Subject-Sensitive Moderate Invariantism.”

REFERENCES

- Cohen, S. (1991): ‘Skepticism, Relevance, and Relativity’, in McLaughlin (ed.), *Dretske and His Critics*, Cambridge, MA: Blackwell.
- Cohen, S. (1999): ‘Contextualism, Skepticism, and the Structure of Reasons’, in *Philosophical Perspectives*, vol. 13, *Epistemology* (pp. 57–89), Blackwell Press.
- DeRose, K. (1995): ‘Solving the Skeptical Puzzle’, *Philosophical Review* 104, 1–52.
- DeRose, K. (1999): ‘Contextualism: An Explanation and Defense’, in J. Greco and E. Sosa (eds.), *The Blackwell Guide to Epistemology* (pp. 187–205), Oxford: Blackwell Press.
- Dretske, F. (1981): *Knowledge and the Flow of Information*, Cambridge: MIT Press.
- Dretske, F. (2000): ‘The Pragmatic Dimension of Knowledge’, in Dretske (ed.), *Perception, Knowledge, and Belief: Selected Essays*, Cambridge: Cambridge University Press.
- Hawthorne, J. (forthcoming): *Knowledge and Lotteries*, Oxford: Oxford University Press.
- Kay, P. (1997): *Words and The Grammar of Context*, Stanford: CSLI.
- Kennedy, C. (1999): *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*, New York: Garland Press.
- Klein, E. (1981): ‘A Semantics for Positive and Comparative Adjectives’, *Linguistics and Philosophy* 4, 1–45.
- Kratzer, A. (1977): ‘What “Must” and “Can” Must and Can Mean’, *Linguistics and Philosophy* 1, 337–355.

- Lewis, D. (1986): 'Scorekeeping in a Language Game', in Lewis (ed.), *Philosophical Papers*, vol. I (pp. 233–249), New York: Oxford University Press.
- Lewis, D. (1996): 'Elusive Knowledge', *Australasian Journal of Philosophy* 74, 549–567.
- Ludlow, P. (1989): 'Implicit Comparison Classes', *Linguistics and Philosophy* 12, 519–533.
- Nunberg, G., Sag, I. and Wasow, T. (1994): 'Idioms', *Language* 70, 491–538.
- Soames, S. (1986): 'Incomplete Definite Descriptions', *Notre Dame Journal of Formal Logic*, 349–375.
- Stanley, J. (forthcoming): 'Context, Interest-Relativity, and Knowledge'.
- Stanley, J. and Szabo, Z.: 'On Quantifier Domain Restriction', *Mind and Language* 15, 219–261.
- Stanley, J. and Williamson, T. (1995): 'Quantifiers and Context-Dependence', *Analysis* 55, 291–295.
- Stanley, J. and Williamson, T. (2001): 'Knowing How', *The Journal of Philosophy* 90.8, 411–444.
- Unger, P. (1975): *Ignorance: A Case for Skepticism*, New York: Oxford University Press.
- Williamson, T. (2000): *Knowledge and Its Limits*, Oxford: Oxford University Press.

Department of Philosophy
435 S. State Street
University of Michigan
Ann Arbor, MI 48109-1003
USA
E-mail: jasoncs@umich.edu