A Noncontextualist Account of Contextualist Linguistic Data

Mylan Engel, Jr.

Northern Illinois University, DeKalb, IL, USA

The paper takes as its starting point the observation that people can be led to retract knowledge claims when presented with previously ignored error possibilities, but offers a noncontextualist explanation of the data. Fallibilist epistemologies are committed to the existence of two kinds of *Kp*-falsifying contingencies: (i) *Non-Ignorable contingencies* [NI-contingencies] and (ii) *Properly-Ignorable contingencies* [PI-contingencies]. For *S* to know that *p*, *S* must be in an epistemic position to rule out all NI-contingencies, but she need not be able to rule out the PI-contingencies. What is required vis-à-vis PI-contingencies is that they all be *false*. In mentioning PI-contingencies, an interlocutor can lead *S* mistakenly to think that these contingencies are NI-contingencies, when in fact they are not. Since *S* cannot rule out these newly mentioned contingencies and since she *mistakenly* takes them to be *NI-contingencies*, it is quite natural that she retract her earlier knowledge claim. In short, mentioning NI-contingencies creates a distortion effect. It makes *S* think that the standards for knowledge are higher than they actually are, which in turn explains why she mistakenly thinks she lacks knowledge. Conclusion: The primary linguistic data offered in support of contextualism can be explained without resorting to contextualism.

Keywords: contextualism, fallibilism, relevant alternatives, knowledge, error possibilities

> We ought likewise to take for granted, as first principles, things wherein we find a universal agreement, among the learned and unlearned, in the different nations and ages of the world. A consent of ages and nations, of the learned and vulgar, ought, at least, to have great authority, unless we can show some prejudice, as universal as that consent is, which might be the cause of it (Reid, *Essays*, 44).
>
> Thomas Reid

Contextualism,[1] as I use the term here, refers to an increasingly popular family of approaches to epistemology which share the following commitments:

(C1) There is no context-independent standard of knowledge.

(C2) Given C1, we should drop all talk about whether or not *S* knows that *p per se* and focus instead on whether sentences of the form '*S* knows that *p*' are true in some specified context of ascription.

(C3) It is the context of the ascriber, not the context of the epistemic subject (unless the subject happens to be self-ascribing knowledge), that determines the operant standards for 'knows' in the ascription in question.

(C4) Given C3, it is possible for an ascriber A_1 to truthfully assert '*S* knows that *p*' at *t*, while another ascriber A_2 truthfully asserts '*S* does not know that *p*' at *t*, provided, e.g., that A_1 is in a low standards context and A_2 is in a high standards context. Given the different epistemic standards operant in their respective contexts of utterance, the proposition expressed by A_2's utterance is not the negation of the proposition expressed by A_1's utterance. So, A_1 and A_2 need not be disagreeing.[2]

Contemporary proponents of contextualism have defended their view primarily by appealing to contextualism's supposedly unique ability to resolve various skeptical and lottery paradoxes.[3] However, as the contextualist "resolutions" of these paradoxes have come under greater scrutiny and greater criticism,[4] contextualists have increasingly sought support for their position elsewhere. In order to both motivate and defend their view today, contextualists cite various linguistic data—data concerning what ordinary competent speakers would *say* in contrasting conversational contexts—data which they contend either can *only* or can *best* be explained by contextualism. The ability to account for this data makes contextualism quite attractive and no doubt helps to explain the burgeoning interest it has recently engendered.

Despite its initial attractiveness, contextualism yields some rather counterintuitive and implausible results.[5] For example, contextualism is committed to a robust *semantic error theory* according to which, as Stephen Schiffer aptly puts it, "people uttering certain knowledge sentences in certain contexts systematically confound the propositions their utterances express with the propositions they would express by uttering those sentences in certain other contexts" (Schiffer 1996, 325). Schiffer finds such an error theory hard to accept:

> What's hard to see is how the hidden-indexical proposal can sustain the idea that fluent speakers systematically confound their contexts, so that even when they're in a context in which Tough is the induced standard occurring in the *false* proposition *they* have just asserted, they mistakenly think they've just asserted a *true* proposition, a proposition that evidently contains the standard Easy that would be induced by an utterance of the problematic sentence in a quite different context. It's as though a fluent, sane, and alert speaker, who knows where she is, were to assert the proposition that it's raining in London when she mistakenly thinks she's asserting that it's raining in Oxford (Schiffer 1996, 236).

Contextualism is also committed to the existence of *unspeakable* and *unthinkable* knowledge: According to contextualism, there are propositions that we know in ordinary contexts, like the proposition that we are not brains in vats, but any

attempt to *assert* that knowledge or *think* that we have that knowledge will create a high standards context in which the resulting assertion or thought is false (See my 2004, 212. Also see Davis 2004, 260.).[6] Finally, as Elke Brendel has shown, contextualism is a philosophical theory that is unknowable in a philosophical context: In a high standards philosophical context, one cannot truthfully claim to know that contextualism is true; one cannot even truthfully claim to know that one has low-standards knowledge that contextualism is true, in such a context (Brendel, present volume). Given these counterintuitive results, one shouldn't accept contextualism without a compelling reason to do so.

My aim in the present paper is to undermine one of the primary reasons offered in support of contextualism, namely, its supposedly unique ability to account for certain linguistic data. To achieve this goal, I shall provide a sketch of a noncontextualist relevant alternatives account of knowledge that explains equally well, if not better, the linguistic data contextualists marshal in support of their view. The significance of such an account is this: If we can account for the contextualists' linguistic data without appealing to contextualism, then we will, as of yet, have no good reason to embrace contextualism and its attendant counterintuitive consequences.

I will begin by presenting some of the most compelling and oft-cited contextualist linguistic data. I will then explain how contextualism accounts for this data, using the Cohen-Lewis Relevant Alternatives Approach to contextualism.[7] Along the way, I will point out some potential shortcomings with the Cohen-Lewis account. Finally, I will develop a noncontextualist relevant alternatives account, one with roots in the work of Gail Stine, that makes use of both internally relevant alternatives and externally relevant alternatives to explain the linguistic data in question.

1. The Data

The linguistic data to which contextualists appeal is obtained by examining our intuitions about what normal competent speakers would *say*, either about their own epistemic status or about the epistemic status of another subject, were they to find themselves in various contrasting conversational situations. Consider by way of illustration Fred Dretske's zebra cases and Keith DeRose's bank cases:

Zebra 1

Bob is at the zoo with his son Caleb and they are looking at the zebra exhibit, when Caleb asks with excitement, "Dad, do you know what kind of animal that is?" Bob responds, "Why yes I do, Caleb. I know that that's a zebra. See the pretty stripes."

Zebra 2

You and I are at the zoo and happen to be standing next to Bob and his son, while they are discussing an animal in the zebra exhibit. As in Zebra 1, Bob has just responded to his son's question with "Why yes I do, Caleb. I know that that's a zebra. See the pretty stripes." We overhear

this conversation. Being an epistemologist, I turn to you and ask whether Bob really knows that the animal in question is a zebra. You reply, "No, he doesn't know that that's a zebra. There's an article in today's newspaper accusing the zoo director of trying to cut costs by disguising mules to look like zebras. Since he doesn't know that that's not a cleverly disguised mule, he doesn't know that that's a zebra."[8]

Bank 1

Keith and his wife are driving home on a Friday afternoon. They plan to stop at the bank on the way home to deposit their paychecks. When they arrive at the bank, they notice that the lines are very long. Although they generally like to deposit their checks as soon as possible, it is not important in this case that they be deposited right away, so Keith suggests that they drive home and deposit their checks on Saturday morning. His wife replies, "Maybe the bank won't be open tomorrow. Lots of banks are closed on Saturdays." Keith reassures her, "No, I know it will be open. I was just there last Saturday, and it's open until noon."

Bank 2

Keith and his wife drive past the bank on a Friday afternoon, as in Bank 1, and notice the long lines. Again Keith claims to know that the bank will be open on Saturday, based on the fact that he was at the bank the previous Saturday, and it was open from 9:00 a.m. to noon. He then suggests that they deposit their checks on Saturday morning, so as to avoid the long lines. But in this case his wife responds as follows: "You know, banks sometimes change their hours, and if the money is not in our account by Monday morning, our house payment will bounce and the bank will foreclose on our mortgage. Are you sure that the bank will be open?" Keith then retracts his earlier knowledge claim and admits, "You're right, Honey. I don't know that it will be open. We'd better go in and check."[9]

What do contextualists say about these cases? They say that each of the above knowledge attributions and knowledge denials is appropriate, warranted, and *true* in the context in which it is made. In the next section, we'll examine why they think that all of these seemingly conflicting assertions are true.

2. How Contextualist Relevant Alternatives Theorists Account for the Data

In addition to being committed to C1–C4 above, those contextualists working within the relevant alternatives framework embrace three additional claims. Like all relevant alternatives theorists, they accept the following claim:

(C5) In order for *S* to know that *p*, *S* must be able to rule out all relevant alternatives to *p*.

What distinguishes *contextualist relevant alternatives theorists* from their noncontextualist counterparts is that they (unlike their counterparts) accept the following two claims:

(C6) Which alternatives to *p* are relevant is determined by the context of ascription. What makes a *p*-falsifying contingency relevant is its *salience* in the conversational context of the ascriber.

(C7) Since an ineliminable *p*-falsifying contingency (i.e. an alternative to *p* that *S* cannot rule out on the basis of her evidence for *p*) can be relevant in one context of ascription and irrelevant in another, the ascription '*S* knows that *p*' can be true in one context and false in another context, even when the respective ascribers are referring to the same *S* and *p*, at the same time.

We are now in a position to see how relevant alternatives contextualists [RA-contextualists] account for the linguistic data presented in section 1. Regarding the zebra cases, RA-contextualists maintain that, in Zebra 1, Bob's claim "I know that that is a zebra" is true, but they contend that, in Zebra 2, your claim "He doesn't know that that's a zebra" is also true. How can that be? Why aren't you and Bob contradicting each other? Because, according to RA-contextualists, you and he are in different conversational contexts. In Bob and Caleb's conversational context, the cleverly disguised mule possibility is not a relevant alternative, and hence, it's properly ignorable.

Things are different in our conversation. In our conversation, the newspaper article has made the cleverly disguised mule possibility salient, and in light of its salience, it has become a relevant *p*-falsifying contingency, a contingency that Bob is in no position to rule out. As a result, the semantic standards for 'knows' operant in *our* conversation are more stringent than those operant in Bob's conversation. For Bob's knowledge claim to be true, he only needs to be able to tell that the animal before him is not some other typical zoo inhabitant. If his evidence for the animal's being a zebra is sufficient to rule out the animal's being a lion or tiger or bear, etc., then his knowledge claim is true. Of course, in the case of zebras, this standard is easily met, since zebras have a quite distinct look and don't resemble any other typical zoo inhabitants. It is much harder to satisfy the standards for 'knows' operant in our conversation in Zebra 2. In light of the newspaper article, the cleverly disguised mule possibility is a relevant alternative to the animal's being a zebra. As a result, in order for the sentence "Bob knows that that animal is a zebra" to be true in *our* conversation (in Zebra 2), Bob would have to be able to rule out the cleverly disguised mule alternative, which he can't do solely on the basis of the visual evidence he has at hand. Since, in our conversation, there is a relevant alternative to the animal's being a zebra that Bob can't rule out, your assertion "Bob doesn't know that the animal is a zebra" is true in our conversational context. Thus, RA-contextualists are able to explain why Bob says what he does and why you say what you do. In Bob's conversation, his assertion is both appropriate and true, because he is in an epistemic position to rule out all the *p*-falsifying contingencies that are relevant in his conversational context. In our conversation, your knowledge denial is also appropriate and true, because there is a *p*-falsifying contingency relevant in our conversational context that Bob is in no position to rule out, and that is why you deny that Bob knows.

How do RA-contextualists account for the linguistic data reflected in the bank cases? They say that, in Bank 1, Keith's assertion "I know that the bank will be open"

is true, but that, in Bank 2, Keith's admission, "You're right, Honey. I don't know that the bank will be open" is also true, even though Keith's evidence for the bank's being open is identical in both cases. How can that be? The RA-contextualist answer is that since *so much more is at stake* in Bank 2 than in Bank 1, more error possibilities are salient in Bank 2 than in Bank 1. Given the newly salient *p*-falsifying contingencies made relevant in Bank 2, the standards for 'knows' operant in Bank 2 are much stricter than those operant in Bank 1. As a result, Keith's evidence is good enough to meet the low standards for 'knows' operant in Bank 1, but that same evidence is not good enough to satisfy the higher standards operant in Bank 2. Thus, once again, RA-contextualists are able to explain why Keith says what he does in each of these bank conversations. In Bank 1, there are no salient error possibilities that Keith can't rule out. Hence, his claim to know that the bank will be open is both appropriate and true. In Bank 2, his wife makes salient an error possibility that was previously ignored. Since Keith realizes that he is in no position to rule out that newly relevant alternative, his concession that he doesn't know the bank will be open is both appropriate and true in that context.

3. Troubling Cases for the RA-Contextualist Account

To its credit, RA-contextualism does provide a prima facie plausible explanation for why the various speakers in the cases just considered say what they do. But when we look at other cases, the RA-contextualist explanation starts to unravel. In the present section, I shall consider two cases where the RA-contextualist account looks decidedly less promising.

As just seen in section 2, RA-contextualists account for Keith's willingness to claim knowledge in Bank 1 and his unwillingness to claim knowledge in Bank 2 by stressing that certain error possibilities, which were irrelevant and properly ignored in Bank 1, are relevant and improperly ignored in Bank 2. What accounts for this difference in relevance? Why are the very same error possibilities that were irrelevant and properly ignored in Bank 1 relevant and improperly ignored in Bank 2? The answer lies in C6 above, namely, that what makes a *p*-falsifying contingency relevant is its *salience* in the conversational context of the ascriber. Cohen makes the point as follows:

> In the case of knowledge ascriptions, I think salience relations are particularly important. In particular, aspects of the context can make the chance of error salient. And when the chance of error is salient, it can lead attributors to intend, expect, presuppose, stricter standards (Cohen 2000, 98).

Why is the possibility that the bank has changed its hours salient in Bank 2, but not in Bank 1? According to Cohen, it is the fact that *so much is at stake* in Bank 2 [They will lose their house if the money is not in their account by Monday morning!] that makes the possibility that the bank has changed its hours salient.[10] If little were at stake, Keith might just as easily have said to his wife, "Don't be silly.

Banks rarely change their hours and almost never do so without ample warning. I know the bank will be open tomorrow, just as it was last Saturday." I have doubts about the "what's at stake" explanation of what makes a mentioned error possibility salient in a conversation. If what accounts for the salience of the error possibility mentioned by Keith's wife is the great deal at stake if he is mistaken, then we should suspect that error possibility to lose its new-found salience were it immediately discovered that nothing is at stake, after all. Imagine the bank case went like this:

Bank 3

As in Bank 2, Keith has claimed to know that the bank will be open on Saturday, his wife has just pointed out that banks sometimes change their hours and has reminded him of the grave consequences if their check to the mortgage company bounces, and Keith has just retracted his earlier knowledge claim, "You're right, Honey. I don't know that the bank will be open tomorrow." Just then, to their surprise, their precocious son pipes up from the back seat, "There's nothing to worry about. After reading about the newly instituted shorter floats at banks in this week's *Wall Street Journal*, I decided not to mail the mortgage check until after you both deposited your checks. So, it doesn't matter whether the money is in your account by Monday morning or not. We'll mail the mortgage check next week after your checks are duly deposited. The mortgage check isn't due for another week anyway." Realizing, once again, that nothing is a stake, Keith tells his wife, "See, just as I told you, I do know the bank will be open on Saturday. That settles it. We'll deposit our checks tomorrow."

Keith's final response in Bank 3 sounds forced, if not outright preposterous. After having just retracted a knowledge claim, it is extremely doubtful that anyone would reinstate it again so easily. But this is precisely what we should expect a speaker to do, if, as RA-contextualists maintain, the salience of the mentioned error possibility were a function of "what's at stake" considerations. I am not suggesting that RA-contextualists hold that *what's at stake* considerations are the only kinds of considerations that can make an error possibility salient, but they do hold that at least in some cases, like Bank 2, *what's at stake* considerations are what account for the salience of a mentioned error possibility.

My point is this: If what makes the possibility that the bank has changed its hours salient in Bank 2 is the fact that so much is at stake, then one would expect that possibility to immediately lose its salience once it is recognized that nothing is at stake. The very oddness of Keith's response at the end of Bank 3 suggests that this is not what one would say were one in Keith's position in Bank 3, but the RA-contextualist who accounts for salience relations in terms of *what's at stake* considerations can't explain why Keith wouldn't retract his retraction and reinstate his original knowledge claim, once it became clear that little is at stake. I will have more to say about Bank 3 later, but for now this much is clear: Bank 3 suggests that once one has retracted a knowledge claim in a given situation, one is typically loathe to reinstate it. Lewis (1979, 355) was perhaps the first to suggest that once

the semantic standards for 'knows' are raised in a given conversational context, they tend to stay raised.[11] The problem is that those RA-contextualists who tie salience relations to *what's at stake* considerations are hard put to explain why the standards remain high once it is discovered that nothing is at stake.

For a second case where RA-contextualism looks decidedly less promising, consider the following conversation between three speakers, low-standards Lisa, high-standards Hastings, and Conté the contextualist:

Bank 4

Arriving at DeRose's Bank late on a Friday afternoon, Lisa, Hastings, and Conté notice that the lines are very long. Hastings wonders aloud, "I wonder if the bank will be open tomorrow?" Lisa replies, "Yes, I know the bank will be open tomorrow. I was just here last Saturday, and it is open from 9:00 a.m. to noon." Hastings thinks a minute and then reminds Lisa that they have written a very large check which will bounce if the money is not in their account by Monday morning, in which case the bank will foreclose on their mortgage, or worse. Hastings then points out, "Perhaps the bank has changed its hours. With so much at stake, you can't be sure that that hasn't happened. So, you don't really *know* the bank will be open." To settle the matter once and for all, he smugly turns to Conté and asks him what he thinks. Realizing that the epistemic standards operant when Lisa made her knowledge claim were low, but also being sensitive to the fact that the importance of not bouncing the check has made Hastings's error possibility salient for the remainder of the conversation, Conté replies as follows: "Lisa's assertion, 'I know the bank will be open tomorrow' is true, but she doesn't know the bank will be open tomorrow."

Conté's assertion is decidedly odd, but according to contextualism, what he says is *true*. Moreover, as a contextualist, this is precisely what Conté should say, if he wishes to be maximally informative. After all, a contextualist like Conté should notice that the semantic standards governing low-standards Lisa's original assertion are different from the semantic standards governing high-standards Hastings's claims after he has mentioned and made salient error possibilities that Lisa was properly ignoring. Nikola Kompa has dubbed sentences like the one Conté asserts in Bank 4 "unpleasant sentences" (Kompa 2002, 5–6). We can see just how unpleasant Conté's assertion is by imagining what he might say next, given the high standards put in place by Hastings's remarks: "Lisa's assertion, 'I know the bank will be open tomorrow' is true, but she doesn't know the bank will be open tomorrow. We better go in and check." But why should they go in and check? The truth of Lisa's assertion entails that the bank will be open tomorrow (because even low standards knowledge requires truth). Since the first conjunct of Conté's reply entails that the bank will be open, even in the high standards context in which it is uttered, Lisa, Hastings and Conté should immediately be able to deduce and thus come to know that the bank will be open tomorrow on the basis of Conté's first conjunct, which then makes it decidedly odd for Conté to go on to claim that Lisa really doesn't know that the bank will be open tomorrow. All three of them should know that the bank will be open, even in the high standards context put in place by Hastings's remarks, given the truth of Conté's first conjunct. In the remainder of

the paper, I will sketch a noncontextualist relevant alternatives account that allows us to explain the linguistic data in section 1 without putting us in a situation where we have to embrace sentences as unpleasant as Conté's.

4. Back to the Zoo

Recall Zebra 1:

> Bob is at the zoo with his son Caleb and they are looking at the zebra exhibit, when Caleb asks with excitement, "Dad, do you know what kind of animal that is?" Bob responds, "Why yes I do, Caleb. I know that that's a zebra. See the pretty stripes."

Here, we are presented with a perfectly ordinary case of perceptual knowledge. Bob's claim to know that that animal is a zebra is correct. He knows that the animal is a zebra in these circumstances, because he has a perceptually justified true belief that the animal before him is a zebra in an environment where he is a reliable zebra discriminator.

Here are some of the things Bob didn't do to acquire his zebra knowledge: He didn't climb over the tall glass wall and poor paint thinner on the animal to see if it was a painted mule. He didn't take a tissue sample to look for signs of mule genes. He didn't search for evidence of holographic trickery. He didn't check to make sure he was awake and not dreaming about zebras. Nor did he check to make sure that the pimple-faced vendor didn't spike his soft drink with LSD. Were he seeing a painted mule or a hologram, dreaming of zebras, or having an LSD-induced zebra-hallucination, his knowledge claim would have been false, and yet, he didn't have to rule out these possibilities in order to know that the animal before him is a zebra. It's worth noting that each of these error possibilities is not only logically possible, each is physically possible as well, and yet Bob did not have to rule them out in order to know that the animal in front of him was a zebra. In fact, he couldn't have ruled out these possibilities on the basis of the perceptual evidence available to him at the time, but he nevertheless knew that the animal was a zebra. Such is the nature of ordinary knowledge: One can know that p [Kp] on the basis of evidence e, even though e is compatible with certain p-falsifying contingencies that one can't otherwise rule out.

In short, ordinary knowledge is fallible knowledge. Many philosophers have attempted to characterize fallibilism. Here's my attempt: Fallibilism is the thesis that S can know that p on the basis of non-entailing reasons for p. Fallibilism entails that S can know that p even though there are p-falsifying contingencies that S is in no epistemic position to rule out, i.e. even though there are p-falsifying contingencies compatible with S's evidence for p. If ordinary empirical knowledge required that one be in an epistemic position to rule out every physically possible chance of error, then ordinary empirical knowledge would be humanly unattainable. Simply put, ordinary knowledge is fallible so that it might be possible. Michael Roth has

aptly captured this important feature of ordinary knowledge with what he calls "The Fallibilist Assumption Governing Empirical Knowledge":

(FA) For every proposition of the form *Kp* (where *p* is an empirical proposition), there are certain contingencies such that (i) their obtaining is physically possible, (ii) if they were to obtain, then *Kp* would be false, and (iii) it is proper for *S* to ignore them when considering whether she is justified in believing that *p*. (Roth 1990, 143)[12]

Those *Kp*-falsifying contingencies satisfying (i), (ii), and (iii) of FA are *Properly-Ignorable contingencies* [PI-contingencies]. Of course, as Roth (1990, 146) rightly notes, not all *Kp*-falsifying contingencies are PI-contingencies. Some *Kp*-falsifying contingencies only satisfy (i) and (ii) of FA. These contingencies are *Non-Ignorable contingencies* [NI-contingencies]. In order for *S* to know that *p*, *S* must be in a position to rule out all NI-contingencies. An example will illustrate the point:

Cheetah 1

Bob and Caleb have moved on to the next exhibit and are now standing in front of a very large cat with black spots, and Bob claims to know that that animal is a cheetah. Bob is unaware of the subtle variations in appearance that differentiate cheetahs from leopards (and there is no sign labeling the exhibit "Cheetahs"). Moreover, like most normal adults, Bob knows that zoos with "large cat" exhibits often have lions and leopards and tigers on display, in addition to cheetahs. The animal's being a leopard is a possibility that he cannot properly ignore. Since his current visual evidence doesn't allow him to discriminate reliably between cheetahs and leopards, his claim to know that the animal is a cheetah is false (even if he happens to be looking at a cheetah).

This underscores the importance of distinguishing between PI-contingencies and NI-contingencies. The cleverly disguised mule possibility is a PI-contingency, whereas the leopard possibility is an NI-contingency. Bob need not be able to rule out the former possibility in order to know that the first animal is a zebra, but he must be able to rule out the latter possibility to know that the second animal is a cheetah.

Let us return to the cheetah exhibit. In the zoo, there are lots of large cats, including lions and leopards. Given their prevalence in the zoo, these are alternatives that must be ruled out in order for Bob to know that the animal is a cheetah [*c*]. In short, *the animal's being a leopard* [*l*] and *the animal's being a tiger* [*t*] are NI-contingencies, and Bob's justification/evidence for believing that *c* must be sufficiently strong to protect him from error with respect to these error possibilities, i.e. his justification for *c* must be sufficiently strong to rule out *l* and *t*. On the other hand, *the animal's being a cleverly painted puma* [*p*] and *it's being a hologram of a cheetah* [*h*] are far-fetched, extremely improbable error possibilities. The sheer remoteness of these latter error possibilities makes it perfectly proper to ignore

them. To be sure, were either of these latter possibilities actual, i.e. were either *p* or *h* true, then *c* and hence *Kc* would be false. Even so, Bob *doesn't* have to be able to *rule out p* and *h* in order to know that *c*. Since *p* and *h* are PI-contingencies, Bob's knowing that *c* only requires that *p* and *h be* false.

5. RATs

Let us call anyone who accepts the distinction between PI-contingencies and NI-contingencies and who thinks that the set of PI-contingencies is not empty a "RAT" [short for "Relevant Alternatives Theorist"].[13] We've just seen how a RAT would explain what's going on in Cheetah 1. Let's see how a RAT would deal with the following case:

Cheetah 2

Bob and Caleb have moved on to the next exhibit labeled "Cheetahs" and are now standing in front of a very large cat with black spots. Bob is an expert on large cats. Seeing the characteristic shape of the spots unique to cheetahs, Bob turns to Caleb and says, "I know that that cat is a cheetah." Unimpressed, his son replies, "Perhaps you're looking at a puma painted to look like a cheetah. If it were an appropriately painted puma, your evidence would be exactly the same. Since you don't know that you're not looking at a painted puma, you don't know that that cat is a cheetah."

What would Bob say at this point? What should he say? Perhaps he would respond as follows: (R1) "Don't be silly! There's no reason whatsoever to think that I am looking at a painted puma [*p*]. This isn't Hollywood, it's the Midwest. Here in Chicago at the Brookfield Zoo, we don't have to worry about such far-fetched nonsense. I know that that animal is a cheetah." But is (R1) what Bob *should* say? Here there is disagreement among the RATs. A noncontextualist RAT, like myself, will insist that (R1) *is* what Bob should say, if he wishes to speak truly. Noncontextualist RATs maintain that PI-contingencies remain PI-contingencies, *unless there is a reason to think that they are true*. Since Bob has no reason to think that *p* is true, this alternative remains properly ignorable, even after Caleb has mentioned it. According to the noncontextualist RAT, absent a reason to think that *p* is true, there is no epistemic difference between Zebra 1 and Cheetah 2.

Contextualist RATs, like Stewart Cohen and David Lewis, think (R1) is neither what Bob would say nor what he should say. Rather, they predict that, in light of the new error possibility made salient to him by Caleb in Cheetah 2, he would respond as follows: (R2) "Hmm. Now that you mention it, I guess I don't know that that's a cheetah." They also think that this is what Bob *should* say. As contextualist RATs see it, before painted pumas were mentioned, this alternative was not relevant, and so it was perfectly proper of Bob to ignore it. As a result, Bob's original assertion "I know that that cat is a cheetah" is true. But once the painted

puma possibility is made salient to Bob in conversation, it becomes relevant and is no longer properly ignorable. Lewis (1979) and John Carroll (2005) speak of a "boundary" that cordons off the NI-contingencies from the PI-contingencies. According to contextualist RATs, *mentioning* these PI-contingencies shifts the boundary outward, in effect, expanding the set of NI-contingencies. What were once PI-contingencies outside the boundary are now NI-contingencies inside the boundary. Since Bob's perceptual evidence is not sufficient to rule out the newly contextually relevant error possibility *p*, Bob's subsequent retraction "I guess I don't know that that's a cheetah" is also true, given the new higher standards context in which it is uttered. By mentioning the painted puma possibility, Caleb has created a context of ascription in which the standards for cheetah knowledge are so high that Bob fails to meet them.

6. Outward Boundaries

Contextualist RATs often speak of "raising epistemic standards." When they do so, they simply mean that more *Kp*-falsifying contingencies have become relevant in a conversational context than were previously relevant. Put in terms of the Lewis/Carroll boundary metaphor, the epistemic standards required for *S* to know that *p* are raised whenever the boundary is moved outward so that *Kp*-falsifying contingencies that were previously outside the boundary and properly ignorable are now inside the boundary and not ignorable. Once the boundary shifts outward, in order for an ascription of the form '*S* knows that *p*' to be true, *S*'s evidence for *p* must not only be good enough to rule out all of the NI-contingencies that were inside the original boundary, but also must be good enough to rule out all of the new NI-contingencies now within the newly expanded boundary.

As will become clear in what follows, the main dispute between contextualist and noncontextualist RATs is a dispute over how easy it is to shift the boundary outward. Contextualist RATs think that it is relatively easy to expand the boundary—a simple conversational mechanism will do. Noncontextualist RATs, on the other hand, think that it is relatively hard to move the boundary and that doing so requires evidence that the *Kp*-falsifying contingency in question is true. In the next section, I will explain how the boundary gets fixed on a noncontextualist RA-theory.

7. The Background

All RATs accept the distinction between PI-contingencies and NI-contingencies. But what accounts for the difference between these contingencies? How do we determine which *Kp*-falsifying contingencies are properly ignorable and which are not? How does the boundary get fixed in the first place? For contextualist RATs, as noted above, the boundary gets set by conversational mechanisms and

thus remains rather malleable. Those error possibilities that all parties to the conversation either *presuppose to be false* or *implicitly agree to ignore* are *properly* ignorable in that conversation.[14] However, a participant in that conversation can refuse to play along. She can refuse to ignore some *p*-falsifying contingency *q* that the other parties to the conversation are ignoring. In such a situation, by *mentioning q*, while perhaps also mentioning the dire consequences of being wrong with respect to *p*, she can make *q* salient. Once *q* becomes salient in the conversation, the boundary shifts outward, and *q* is no longer properly ignorable by those participating in the conversation.

The noncontextualist RA-theory that I endorse offers a different explanation of how the boundary gets fixed, namely, it gets fixed by epistemic considerations, not by conversational ones. We make our epistemic judgments and our corresponding knowledge ascriptions relative to a set of "commonsense" background beliefs that all "normal," "reasonable," "sane" people are expected to hold, beliefs that all "normal," "reasonable," "sane" people *take themselves to know*.[15] For example, all normal, reasonable, sane people believe the following commonsense propositions:

(P1) That there is an external physical world populated with people and other living beings, forests and oceans, houses and strip malls, tables and chairs, etc.

(P2) That the world existed before we were born and will continue to exist after we are dead.

(P3) That we possess physical bodies capable of experiencing pleasure and pain.

(P4) That people rarely engage in deceptive behavior *for no reason*.[16]

All reasonable sane people not only hold these beliefs, they *take themselves to know* that these beliefs are true. Such commonsense beliefs are by no means new. Over 200 years ago, Thomas Reid identified a number of similar commonsense beliefs, including beliefs in the following principles:

(P5) That those things did really happen which I distinctly remember (Reid, *Essays*, 474).

(P6) That those things do really exist, which we distinctly perceive by our senses, and are what we perceive them to be (Reid, *Essays*, 476).

(P7) That there is life and intelligence in our fellow-men with whom we converse (Reid, *Essays*, 482).

Reid points out that these principles have the "consent of ages and nations, of the learned and unlearned" (Reid, *Essays*, 464) and insists that when we find such a general agreement in principles that concern human life "this must have great authority with every sober mind that loves truth" (Reid, *Essays*, 464).

These commonsense beliefs and others like them constitute what I call "The Background". On the noncontextualist RA-theory I'm proposing, the boundary is fixed by The Background. Those error possibilities that are *likely to be true* on The

Background cannot be ignored, nor can those error possibilities that are *as likely as not* on The Background. These two kinds of possibilities are inside the boundary and are NI-contingencies. However, those error possibilities that are *false* on The Background,[17] as well as those that are *extremely improbable* on The Background, are outside the boundary, and it is perfectly proper to ignore them. And as Reid points out in the epigram with which I began, The Background should not be rejected, unless there is a *reason* to think that one of the propositions in The Background is false.[18] In the next section, we will see how the insights gleaned from the present section can be utilized to develop a viable noncontextualist RA-theory.

8. A Noncontextualist Relevant Alternatives Account[19]

Two preliminary remarks are in order. First, it is important to distinguish those *Kp*-falsifying alternatives that are internally relevant from those that are externally relevant. Let us start with the externally relevant alternatives. The set of *externally relevant alternatives* is the set of alternatives that *preclude S* from knowing that *p* iff they are true.[20] The set of *internally relevant alternatives* is the set of alternatives that *S* must be able to rule out or neutralize in order to know that *p*, regardless whether these alternatives are true or not (i.e. it is the set of NI-contingencies). Traditionally, RATs have focused exclusively on the internally relevant alternatives as just defined and have failed to take into account the epistemic significance of externally relevant alternatives.

Second, it is also important to distinguish the set of internally relevant alternatives from the set of contextually relevant alternatives. Let us define the set of *contextually relevant alternatives* as the set of *Kp*-falsifying contingencies that are salient in a given ascriber's conversational context. The contextualist RAT maintains that the set of internally relevant alternatives and the set of contextually relevant alternatives are identical. The noncontextualist RAT denies that these two sets are identical. As we shall see, according to the noncontextualist RA-theory, an alternative can be contextually relevant without being internally relevant, and an alternative can be internally relevant without being contextually relevant. To see why, we need an account that specifies which alternatives are internally relevant.

a. Internally Relevant Alternatives to *Kp*

Gail Stine has argued that "an alternative is relevant only if there is some reason to think that it is true" (Stine 1976, 252). The problem with Stine's proposal, as Ernest Sosa (2004, 37–38 and 51–53) has rightly observed, is that it fails to count those *p*-falsifying contingencies that are equipoised with *p* as internally relevant alternatives to *Kp*. After all, if a *p*-falsifying alternative *a* is equipoised with *p* (on *S*'s evidence), then there is no reason for *S* to think that *a* is true, and so, on Stine's account, it is not an internally relevant alternative to *p*. And yet, intuitively,

it seems that an alternative *a* that is equiprobable with *p* and entails ~*p* is precisely the kind of alternative that *S* must be able to rule out in order to know that *p*. That is, intuitively, such an alternative *is* an internally relevant alternative to *p*. Stine's mistake, as I see it, is that she was focusing exclusively on what is required to convert a PI-contingency to an NI-contingency, and in so doing, she ignored all of those NI-contingencies that are already inside the boundary, because they are *probable* or at least *as likely as not* on The Background. Nevertheless, Stine's fundamental insight—that in order for a PI-contingency to become an NI-contingency, there must be a reason to think that that PI-contingency is *true* (and not merely possible in some specified sense)—strikes me as right. I submit that the following account of internally relevant alternatives captures Stine's insight while addressing Sosa's worry:

(IRA) An alternative *a* to *Kp* is an *internally relevant alternative* for *S* in circumstances *C* iff:

(1) *a* entails ~*Kp* because either (i) *a* entails ~*p* or (ii) *a* is a defeater for *S*'s justification for *p*, and

(2) either (i) *a* is at least *as likely as not* on The Background or (ii) although *a* is *false* (or *improbable*) on The Background, *S* has a specific reason to think that *a* is true in *C*.

As noted above, RATs have focused primarily on internally relevant alternatives and have paid little attention to externally relevant alternatives. This oversight is significant since externally relevant alternatives play a crucial role in determining whether or not knowledge attributions are warranted. Since externally relevant alternatives directly affect the warrantedness of our knowledge ascriptions, it is important to have an account of such alternatives.

b. Externally Relevant Alternatives to *Kp*

Recall that a properly-ignorable *Kp*-falsifying contingency [PI-contingency] is any contingency outside the boundary whose truth entails ~*Kp*. As such, every PI-contingency is an externally relevant alternative to *Kp*. It is important to note that there are two types of properly-ignorable *Kp*-falsifying contingencies: *Type 1 PI-contingencies* are PI-contingencies such that were they to obtain, then *p* and hence *Kp* would be false. *Type 2 PI-contingencies* are PI-contingencies such that were they to obtain, *p* would still be true, but *S*'s belief that *p* would be accidental and hence *Kp* would be false.[21] Using this distinction, we can define externally relevant alternatives as follows:

(ERA) An alternative *a* to *Kp* is an *externally relevant alternative* for *S* in circumstances *C* iff *a* is either a Type 1 PI-contingency or a Type 2 PI-contingency in *C*.

The epistemic significance of these externally relevant alternatives is this: Any *true* Type 1 PI-contingency automatically precludes *Kp*, because it entails $\sim p$. Any *true* Type 2 PI-contingency is essentially a Gettier-style defeater. Hence, a true Type 2 PI-contingency *c* precludes *Kp* unless *S* has or acquires additional independent evidence for *p*, evidence that is not undermined by *c*. It being *true* that there are lots of painted mules in the "zebra" exhibit precludes *Kz*, simply on the basis of *S*'s visual experience, but if *S* climbs over the wall, takes a tissue sample, and determines that the animal in question has zebra genes, then *Kz* will be true, despite the truth of the Type 2 PI-contingency that there are lots of painted mules in the exhibit.

c. A Necessary Condition for Knowledge

Like all versions of the relevant alternatives theory, on my account a person knows that *p* only if she can rule out or neutralize all internally relevant alternatives to *Kp*, i.e. only if she can rule out or neutralize all the non-ignorable *Kp*-falsifying contingencies inside the boundary.[22] Of course, it must also be the case that there are no true, nonneutralized, externally relevant alternatives to *Kp*. Combining these insights, we get the following Necessary Condition for Knowledge:

(NCK) *S* knows that *p* only if (i) *S* is in an epistemic position to rule out or neutralize all internally relevant alternatives to *Kp*, and (ii) there are no true, nonneutralized, externally relevant alternatives to *Kp*.

Let us now apply this noncontextualist RA-theory to the cases with which we began.

9. How the Noncontextualist RA-Theory Accounts for the Contextualist Linguistic Data

We are now in a position to see how a noncontextualist RAT would account for the linguistic data identified in section 1 above. Recall that in the zebra cases, Bob says to his son "I know that *z*" [where *z* = that animal is a zebra], and you say to me, "Since he doesn't know that that's not a cleverly disguised mule, he doesn't know that *z*." On my noncontextualist RA-account, only one of you is right. If the newspaper's story is true and the zoo director has disguised many mules to look like zebras and put these mules in with the zebras, then Bob doesn't know that the animal in question is a zebra, even if he happens to be looking at a zebra. For in that situation, given that there is a *true* externally relevant alternative to *Kz*, it is just a matter of veritic luck that Bob is looking at a zebra rather than one of the indistinguishable mules and veritic luck is incompatible with knowledge.[23] In that case, you are right and Bob is wrong.

Why then does Bob claim to know that z? Because he is *warranted* in making such a claim. After all, he hasn't read the morning paper and, thus, he has *no reason* to think that the painted mule possibility is actual. Moreover, he is fully justified in ignoring that possibility, given its improbability on The Background. He claims to know that z, because (i) he is in a position to rule out all of those alternatives to z that are *at least as likely as not* on The Background and (ii) he has no reason to think that any externally relevant alternatives are true. In other words, he claims to know that z, because (i) he can rule out all of the internally relevant alternatives to z and (ii) he has no reason to think that there are any true externally relevant alternatives to Kz. In short, Bob claims to know that z because he is *warranted* in believing that he has satisfied both conditions required by (NCK).

Suppose, on the other hand, that the newspaper story is a fraud and the zoo director is a woman of impeccable integrity who would never allow such a deceitful practice to occur in her zoo. In such a situation, your claim that Bob doesn't know that z is simply false. In that case, Bob is right and you are wrong.[24] Why then do you assert that Bob doesn't know that z? Simple. On the basis of the newspaper story, you are *warranted* in believing that an externally relevant alternative to Bob's knowing that z is true, and since true externally relevant alternatives preclude knowledge, you are *warranted* in believing that Bob doesn't know that z. You assert what you do, because you are *warranted*, on the basis of the newspaper story, in believing that what you are asserting is true. Anyone in your position would assert just what you have asserted vis-à-vis Bob's epistemic status, but that does not make your assertion true. To think otherwise is to confuse warranted assertability with truth.[25]

What about DeRose's bank cases? How would a noncontextualist RAT like myself account for such cases? First, we've all had lots of experience with banks and know the following two propositions: (1) That banks rarely change their hours, and (2) that it is even more rare that they change their hours without posting those changes and without providing ample notification to their customers of the change of hours. These propositions are part of The Background. Consequently, *the bank has altered its hours* [a] is an internally relevant alternative to the bank's being open on Saturday *only if* there is a reason to think that the bank *has* changed its hours. Since there is no such reason in any of the Bank Cases we have considered, a is not an internally relevant alternative. Hence, Keith knows that the bank will be open on Saturday in both Bank 1 and Bank 2, and Lisa knows that the bank will be open on Saturday in Bank 4. As a result, Hastings is simply wrong in Bank 4 when he claims that Lisa lacks such knowledge. Conté is also wrong in Bank 4 when he claims that Lisa does not know that the bank will be open.

The puzzle, of course, is to explain why Keith retracts his knowledge claim in Bank 2, if it really is the case that he knows the bank will be open. I submit that his wife's mentioning the fact that banks sometimes change their hours creates a distortion effect. It makes Keith think that the alternative *the bank has altered its hours*

[*a*] is internally relevant when in fact isn't, because there is no reason whatsoever to think that their bank *has* changed its hours. Why then is Keith misled into taking *a* to be an internally relevant alternative? My suggestion is as follows: Typically, ordinary normal people only mention error possibilities that there is some reason to think are true. In polite ordinary discourse, competent speakers don't conjure up error possibilities willy-nilly, rather they only mention possibilities they think are true or highly likely. We might make the point as follows: In ordinary conversational practice, normal competent speakers seem to adhere the following Gricean-type rule:

The Rule of Appropriate Challenges

> When challenging a person's claim to know that *p*, *only* mention those *Kp*-falsifying contingencies that one has reason to think are true or highly probable.

Do I have any evidence that normal people adhere to such a rule and expect others to do so as well? Presumably, most of you reading this article are philosophers, and as philosophers, you've no doubt witnessed some philosopher or other, manifestly proud of his own cleverness, present an utterly preposterous, yet logically possible, error possibility in an ordinary conversation. You've also seen the attempt fall flat. Ordinary people don't feel that they have been refuted by such a far-fetched example, they simply think the person presenting the example is a jerk.

I submit that because most people adhere to The Rule of Appropriate Challenges, one's raising a challenge creates certain conversational implicatures. In particular, when one mentions an error possibility *e*, one conversationally implies that one has a reason to think that *e* is true (or likely to be true), and one conversationally implies that the truth of *e* would falsify the knowledge claim in question. Once a person thinks, of some PI-contingency (i.e. some externally relevant but internally irrelevant alternative) *x*, that there is a reason to think that *x* is true, that person will come to think that *x* is an internally relevant NI-contingency. In light of the aforementioned conversational implicatures, when a speaker mentions an externally relevant error possibility *e* (that she has no reason to think is true), the hearer will quite naturally think that she must have a reason to think that *e* is true or else she wouldn't have mentioned *e*. This, in turn, will lead the hearer to mistakenly think that *e* is an internally relevant alternative, and if the hearer is in no position to rule out *e*, then the hearer will mistakenly retract her original knowledge claim. If the speaker has no reason to think that *e* is true, she will have effectively distorted the standards required for knowing on the occasion in question.

I submit that something like this is what is going on in Bank 2. After all, the fact that banks sometimes change their hours is no reason to think that this particular bank *has* changed its hours. Of course, given the expectations created by The Rule of Appropriate Challenges, Keith would very naturally think to himself, "My wife

never would have mentioned the fact that banks sometimes change their hours unless she had a reason to think that this bank has changed its hours. After all, it is *this bank's hours* with which we are concerned." This, in turn, might well lead Keith to think that *this bank has altered its hours* is an internally relevant alternative and one he can't rule out or neutralize. In such a situation, he would be fully warranted in retracting his earlier knowledge claim. Of course, if there is no reason to think that *this bank* has changed its hours, then in point of fact, *a* is not an internally relevant alternative. Consequently, Keith's retraction, though warranted, is false.

This appeal to The Rule of Appropriate Challenges also helps explain why a person will continue to deny knowledge to himself on the basis of mentioned error possibilities even after it is discovered that little is at stake. The conversational implicature created by mentioning an error possibility *e*—namely, the implication that the speaker must have a reason to think the *e* is true and relevant to the proposition in question—is not cancelled simply by discovering that little is at stake. Once Keith thinks that his wife has a reason to think that the bank has changed its hours, he will continue to deny knowledge to himself regardless of how little is at stake.

Finally, The Rule of Appropriate Challenges would explain why Bob says what he does to Caleb in Cheetah 2. Bob realizes that the painted puma possibility is extremely improbable on The Background. He also realizes that Caleb has no reason to think that the animal in question is a painted puma. Thus, he realizes that the painted puma possibility is not an internally relevant alternative, despite its having been mentioned by his son. Here the conversational implicature that would normally be created by Caleb's assertion is cancelled by Bob's knowledge of what information Caleb possesses. As a result, rather than being misled into thinking that Caleb has a reason to think that the animal in question is a painted puma, Bob rightly chastises his son for violating The Rule of Appropriate Challenges.

10. Conclusion

The noncontextualist RA-account developed herein when supplemented by The Rule of Appropriate Challenges is able to account for the linguistic data contextualists typically offer in support of their theory, and it is able to do so without being committed to unpleasant sentences or any of the other counterintuitive consequences of contextualism. One of those consequences we observed is that contextualists are committed to a robust semantic error theory according to which normal, sane, competent speakers of the language frequently don't know what propositions are being expressed by the knowledge ascriptions they utter. On the noncontextualist RA-theory proposed herein, people do know what propositions their knowledge-ascribing utterances express. They get things wrong not because

they don't understand their own language, but because of a psychological distortion effect that is created when a speaker inappropriately mentions an externally relevant alternative *a* to *Kp* that she has no reason to think is true. Her doing so conversationally implies that she has a reason to think that *a* is true, which in turn misleads her hearers into thinking that *a* is an internally relevant alternative. In such a situation, if her hearers cannot rule out *a*, they will think they lack knowledge, which in fact they actually possess (assuming *a* is false). In such cases, their resulting knowledge retractions, though warranted, are false.[26]

Notes

1. There are numerous other contextualisms afloat in philosophy, e.g. contextualist accounts of "explanation" in the philosophy of science, contextualist accounts of "freedom" in metaphysics, contextualist accounts of "right" and "wrong" and "harm" and "benefit" in ethics. When I use the term 'contextualism' in the present paper, I am only referring to *epistemic contextualism* as defined by commitments (C1)–(C4).
2. This version of contextualism, whose principal defenders include Stewart Cohen, David Lewis, and Keith DeRose, has received many names in the literature including *conversational contextualism, ascriber contextualism, attributor contextualism,* and *semantic contextualism.* What name one settles on is of little significance. What is important is that we clearly distinguish this version of contextualism from the kind of *inferential contextualism* endorsed by Michael Williams. See Pritchard 2002, for a detailed discussion of these two different kinds of contextualism in epistemology. In this paper, I will be focusing exclusively on the Cohen-Lewis-DeRose-style contextualism that is defined by commitments (C1)–(C4).
3. See, e.g., Cohen 1988, 1999, and 2000; DeRose 1995 and 1996; and Lewis 1996.
4. Some of these criticisms have been spelled out in Schiffer 1996, Feldman 1999, Klein 2000, Sosa 2000, Hawthorne 2003, and my 2004.
5. A number of recent articles have highlighted some of these counterintuitive results. See, e.g., Schiffer 1996, Feldman 1999, Vogel 1999, Hawthorne 2000, Kompa 2002, Brendel 2003, and my 2004.
6. DeRose acknowledges this very point:

   > Thus, on our solution, we do know, for instance, that we're not BIVs, according to ordinary low standards for knowledge. But, though (1) [of the BIV argument, viz., 'I don't know that I'm not a BIV'] is false when evaluated according to those ordinary low standards, we're able to explain its plausibility, as we've seen, by means of the fact that the high standards at which (1) is true are precisely the standards that an assertion or denial of it put into play. Since attempts to assert (1) are bound to result in truth, and attempts to deny it are destined to produce falsehood, it's no surprise that we find it so plausible (DeRose 1995, 39f).

   DeRose is so eager to explain the plausibility of premise (1) of the BIV argument that he doesn't seem to notice how counterintuitive it is to maintain that people have all sorts of unspeakable and unthinkable knowledge.

7. I have elected not to focus on how the DeRose-style Subjunctive Conditional Approach to contextualism accounts for this data, because my primary concern is not with how contextualists account for their data, but rather with providing a noncontextualist account that can explain that data equally well. Moreover, since I will attempt to account for these data using a noncontextualist version of the relevant alternatives theory, it will be useful to have the contextualized version of the relevant alternatives account before us as a contrast.

8. I'm taking liberty here with Dretske's famous example. Dretske (1970) offered the zebra cases to illustrate that knowledge is not closed under known logical implication. Thus, while he would agree that Bob knows that the animal is a zebra in Zebra 1, he would reject my claim, in Zebra 2, that since Bob doesn't know it's not a cleverly disguised mule, he doesn't know that it's a zebra. Dretske's purpose in presenting these examples was to argue that that last conditional is false. Contextualists, however, have usurped Dretske's zebra examples and have argued that rather than showing that the principle of epistemic closure is false, they show that the semantic standards of 'knows' vary across conversational contexts, but that in any given context closure holds. Here, as far as linguistic data goes, I side with the contextualists. It seems to me that any normal speaker who took seriously the possibility that the animal was in fact a cleverly disguised mule would deny that Bob knows the animal is a zebra, just on the basis of its looking like a zebra.
9. Bank 1 and Bank 2 are close paraphrases of Keith DeRose's contrasting bank cases, Bank A and Bank B. See DeRose 1992, 913.
10. To my knowledge, Cohen hasn't specifically addressed DeRose's bank cases in print, but he has discussed the following very similar case:

    > Mary and John are at the L.A. airport contemplating taking a certain flight to New York. They want to know whether the flight has a layover in Chicago. They overhear someone ask a passenger Smith if he knows whether the flight stops in Chicago. Smith looks at the flight itinerary he got from the travel agent and responds, "Yes, I know—it does stop in Chicago." It turns out that Mary and John have a very important business contact they have to make at the Chicago airport. Mary says, "How reliable is that itinerary? It could contain a misprint. They could have changed the schedule at the last minute." Mary and John agree that Smith doesn't really *know* that the plane will stop in Chicago. They decide to check with the airline agent. (Cohen 1999, 58)

    What makes the possibility that the itinerary is inaccurate salient? Cohen's answer is: "In the case of John and Mary, it is the importance of the Chicago meeting that makes the chance of error salient" (Cohen 1999, 61). With such an important business meeting on the line, a great deal is at stake. It is the fact that so much is at stake that accounts for the salience of the error possibilities Mary mentions.

    Since Cohen maintains that it is *what's at stake* considerations that account for the salience of the error possibilities Mary raises, presumably, he would also hold that *what's at stake* considerations explain why the possibility that the bank has changed its hours is salient in Bank 2.

11. Lewis makes the point as follows:

    > For some reason, I know not what, the boundary readily shifts outward if what is said requires it, but does not so readily shift inward if what is said requires that. (Lewis 1979, 355)

12. This is a slightly modified paraphrase of Roth's (FA). See Roth 1990, 143.
13. With apologies to Jonathan Vogel who uses 'RAT' to refer the RA theory itself, rather than its proponents [See Vogel 1999.].
14. For an excellent discussion of the role of the *presupposition set* (i.e. the set of propositions presumed to be true by all parties to the conversations) in a contextualist semantics, see Carroll 2005. On this view, those error possibilities whose negations are members of the presupposition set are properly ignorable in that conversation, unless that presupposition is cancelled through conversational mechanisms.
15. Roth (1990, 144) considers a similar proposal.
16. Our belief in (P4) helps to explain why we form beliefs in accordance with what Reid calls the *principle of credulity*. As Reid puts it: "It is evident, that, in the matter of testimony, the balance of human judgment is by nature inclined to the side of belief; and turns to that side of itself, when there is nothing put into the opposite scale" (Reid, *Inquiry*, 194).

17. By "an error possibility that is *false* on The Background," I simply mean that The Background *entails* that the proposition expressing that error possibility is false. Of course, like any set of propositions, The Background can entail that a proposition is false, even though that proposition is in fact true, but this can only happen if one or more of the propositions constituting The Background are false. In claiming that those error possibilities that are *false* on The Background are PI-contingencies, I am claiming that they should be ignored unless there is a reason to think that some proposition in The Background is false.
18. It's worth noting that in most normal conversations, what I'm calling "The Background" and what Carroll calls the "*presupposition set*" will have the same members. The principal difference between my The Background and Carroll's *presupposition set* concerns how one can change the membership of these sets. On Carroll's view, conversational mechanisms can change the membership of the *presupposition set*. On my view, the only thing that can change the membership of The Background is evidence that one of the member positions is false (or highly likely to be false).
19. In the standard literature, relevant alternatives for *Kp* are generally thought to be ***p***-falsifying contingencies, i.e. they are thought to be alternatives to *p*! But there are relevant alternatives to *Kp* that are not *p*-falsifying contingencies. Rather, these latter alternatives are ***Kp***-falsifying contingencies because, if true, they defeat *S's* justification for believing *p*, without entailing $\sim p$. All relevant alternatives to *p* will, of course, be relevant alternatives to *Kp*, but not all relevant alternatives to *Kp* are relevant alternatives to *p*. In the present section, when necessary, I will clearly distinguish between those contingencies that are *p*-falsifying and those that are *Kp*-falsifying (without being *p*-falsifying).
20. Despite the "if and only if," this definition does not entail that *S* knows that *p* whenever all of the externally relevant alternatives are false, because *S* might fail to meet some other condition, like the belief condition, required for knowing that *p*. Externally relevant alternatives are defined in terms of epistemic preclusion. If an externally relevant alternative to *Kp* is true, then that alternative *precludes Kp*. On the other hand, if none of the externally relevant alternatives is true, then *Kp* is *not precluded* by these alternatives. That, of course, is consistent with *Kp* being precluded for some other reason.
21. An alternative way of understanding the difference between Type 1 and Type 2 PI-contingencies is as follows: Type 1 PI-contingencies are ***p*-falsifying** contingencies. When true, they preclude *Kp* by virtue of entailing $\sim p$. Type 2 PI-contingencies are ***Jp*-defeating** contingencies. When true, they preclude *Kp* by virtue of defeating *S*'s justification for believing that *p*. When a Type 2 PI-contingency is true, *S* fails to know that *p* unless either that defeating contingency is itself defeated or *S* has additional independent justification for *p* that is not undermined by the Type 2 PI-contingency in question.
22. To be in an epistemic position to rule out or neutralize an internally relevant alternative *a* to *p*, *S* must possess evidence sufficient to rule out or neutralize *a*. To be in an epistemic position to *rule out* an alternative *a*, *S* must possess evidence sufficient to justify the belief that $\sim a$. To be in an epistemic position to *neutralize* an alternative *a*, *S* must be justified in believing some neutralizing proposition *n* such that *n* undercuts *a*'s force as a defeater for *p*. For an excellent discussion of defeater neutralization, see Lehrer 1990, especially Chapters 7 and 9.
23. Veritic luck can be defined as follows: A person *S* is *veritically lucky* in believing that *p* in circumstances *C iff*, given *S*'s evidence in *C*, it is just a matter of luck that *S*'s belief is true in *C*. For a detailed discussion of the nature of veritic luck and a demonstration of its incompatibility with knowledge, see my 1992.
24. Assuming, of course, that no other externally relevant alternative is true and nonneutralized.
25. Contextualist RA-accounts collapse the distinction between warranted assertability and truth. It is a virtue of the present noncontextualist RA-account that it can retain this plausible distinction.

26. An earlier version of this paper was presented at the 2004 Bled Conference on Contextualism. I would like to thank those in attendance for their helpful comments and criticisms. In addition, I have benefited greatly from conversations with John Carroll, Alastair Norcross, Bruce Russell, and Timothy Williamson concerning contextualism. Special thanks to Elke Brendel and Matthias Steup, both of whom provided detailed comments and suggestions on the penultimate draft. The present paper is much improved as a result of their suggestions.

References

Brendel, E.: 2005, 'Why Contextualists Cannot Know They Are Right: Self-Refuting Implications of Contextualism', *Acta Analytica* 20(2), 38–55.

Carroll, J.: 2005 'Boundary in Context', *Acta Analytica* 20(1), 43–54.

Cohen, S.: 1988, 'How to Be a Fallibilist', *Philosophical Perspectives, 2, Issue Epistemology*, 91–123.

Cohen, S.: 1999, 'Contextualism, Skepticism, and the Structure of Reasons', *Philosophical Perspectives, 13, Epistemology*, 57–89.

Cohen, S.: 2000, 'Contextualism and Skepticism', *Philosophical Issues, 10, Skepticism*, 94–107.

Davis, W.: 2004, 'Are Knowledge Claims Indexical?', *Erkenntnis* 61, 257–281.

DeRose, K.: 1999, 'Contextualism: An Explanation and Defense' in: Greco, J. and Sosa, E. (eds.), *The Blackwell Guide to Epistemology*, Blackwell, Malden, MA and Oxford, 187–205.

DeRose, K.: 1996, 'Knowledge, Assertion and Lotteries', *Australasian Journal of Philosophy* 74, 568–579.

DeRose, K.: 1995, 'Solving the Skeptical Problem', *The Philosophical Review* 104, 17–52.

DeRose, K.: 1992, 'Contextualism and Knowledge Attributions', *Philosophy and Phenomenological Research* 52, 913–929.

DeRose, K.: 1991, 'Epistemic Possibilities', *The Philosophical Review* 100, 581–605.

Dretske, F.: 1970, 'Epistemic Operators', *Journal of Philosophy* 67, 1007–1023.

Engel, M.: 2004, 'What's Wrong with Contextualism, and a Noncontextualist Resolution of the Skeptical Paradox', *Erkenntnis* 61, 203–231.

Engel, M.: 1992, 'Is Epistemic Luck Compatible with Knowledge?', *Southern Journal of Philosophy* 30, 59–75.

Feldman, R.: 1999, 'Contextualism and Skepticism', *Philosophical Perspectives, 13, Epistemology*, 91–114.

Hawthorne, J.: 2004, *Knowledge and Lotteries* (Oxford: Oxford University Press).

Klein, P.: 2000, 'Contextualism and the Real Nature of Academic Skepticism', *Philosophical Issues, 10, Skepticism*, 108–116.

Kompa, N.: 2002, 'The Context Sensitivity of Knowledge Ascriptions', *Grazer Philosophische Studien* 64, 1–18.

Lehrer, K.: 1990, *Theory of Knowledge* (Boulder, CO: Westview 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.

Pritchard, D.: 2002, 'Two Forms of Epistemological Contextualism', *Grazer Philosophische Studien* 64, 19–55.

Reid, T.: *An Inquiry into the Human Mind: On the Principles of Common Sense*, ed. D. Brookes (University Park, PA: Pennsylvania State University Press, 1997).

Reid, T.: *Essays on the Intellectual Powers of Man*, ed. D. Brookes (University Park, PA: Pennsylvania State University Press, 2002).

Roth, M.: 1990, 'The Wall and the Shield: K-K Reconsidered', *Philosophical Studies* 59, 137–157.

Schiffer, S.: 1996, 'Contextualist Solutions to Scepticism,' *Proceedings of the Aristotelian Society*, 317-333.
Sosa, E.: 2004, 'Relevant Alternatives, Contextualism Included', *Philosophical Studies* 119, 35–65.
Sosa, E.: 2000, 'Skepticism and Contextualism', *Philosophical Issues, 10, Skepticism*, 1-18.
Stine, G.: 1976, 'Skepticism, Relevant Alternatives, and Deductive Closure' *Philosophical Studies* 29, 249–261.
Vogel, J.: 1999, 'The New Relevant Alternatives Theory', *Philosophical Perspectives, 13, Epistemology*, 155–180.

Received: December 2004

Mylan Engel, Jr.
Department of Philosophy
Northern Illinois University
DeKalb, IL 60115
USA
e-mail: mylan-engel@niu.edu