

KNOWLEDGE, CONTEXT, AND SOCIAL STANDARDS*

ABSTRACT. This paper defends the view that standards, which are typically social in nature, play a role in determining whether a subject has knowledge. While the argument focuses on standards that pertain to reasoning, I also consider whether there are similar standards for memory and perception.

Ultimately, I argue that the standards are context sensitive and, as such, we must view attributions of knowledge as indexical. I exploit similarities between this view and a version of the relevant alternatives reply to skepticism in order to defend this reply against the objection that it is ad hoc.

What consequences does the fact that we are social animals have for a theory of knowledge? Some philosophers have claimed that knowledge has a social component. Evidence one does not possess can undermine one's knowledge, when that evidence is possessed by a relevant social group to which one belongs.¹ While remaining neutral on this particular issue, I too, will claim that knowledge has a social component. However, I will argue that social factors influence whether evidence one *does* possess undermines one's knowledge. In the end, I will maintain that this phenomenon reveals that attributions of knowledge are context sensitive. One speaker may attribute knowledge to a subject, while another speaker denies knowledge to that same subject, without contradiction. Once we see this, we will obtain a better perspective from which to view skeptical arguments, and a means to resist their conclusions.

1.

We can begin by examining the concept of having good reasons for believing a proposition. I shall assume that such a concept is a fundamental constituent of knowledge.² Given this, under what conditions does (a subject) *S* have good reasons to believe that *q*? Having good reasons comprises both logical and psychological elements. I will argue that distinct concepts of having good reasons are revealed by different construals of the relation between these elements. The issue then will arise as to which of these concepts figure as constituents of

knowledge. Once this is settled we will be in a position to see some interesting features of the concept of knowledge.

As a first step we need to look more closely at the nature of epistemic reasons. A large class of epistemic reasons have a prima facie structure, i.e., they are defeasible. Thus, where r is an (epistemically) prima facie good reason to believe that q , there exists a defeater d such that (rd) is not a good reason for believing that q . We might say that S has good reasons *simpliciter* to believe q just in case S has prima facie good reasons for which he possesses no defeaters (or S has indefeasible good reasons).³ Following John Pollock⁴ we can say there are two types of defeaters. Where r is a prima facie good reason to believe q , d is a type I defeater just in case d is a good reason to believe that q is false even though r is true. For example, where $r =$ (the table appears to be red) and $q =$ (the table is red), d (a reliable sources says there are no red tables in the room) is a type I defeater. If d undercuts the connection between the truth of r and the truth of q , independently of being a reason to believe q is false, then d is a type II defeater, e.g., $d =$ (there is a red light shining on the table).

We can make use of this prima facie structure to display ambiguities in the concept of having good reasons. Suppose S knows that r where r is a prima facie good reason to believe that q . S also knows that d , where d is a defeater. If this is all that is relevant to whether S should believe that q , then when matters are proceeding optimally, S will refrain from coming to believe q . If S already believes q on the basis of his believing r , he will give up his belief that q upon coming to believe that d . Of course actual epistemic subjects often fall short of this ideal. Suppose S , in the face of d , believes q on the basis of r . Does S have good reasons (*simpliciter*) for believing that q ? A quick response is that S fails to have good reasons. After all, the case has been specified as one where S possesses a defeater of his prima facie reason r . However, the relevance of a defeater can range from obvious to very subtle. Suppose in the case imagined that d defeats through a very subtle line of reasoning that would escape all but the most acute intelligence. S believes q only because he is of normal intelligence and so fails to appreciate the defeating effect of d . With the details so specified, does S fail to have good reasons? Here we might hesitate, granting that one can have good reasons provided that one does not believe any obvious defeaters.⁵ (We might very well say S is *justified* in believing that q .) This does not entail that the original

judgment that the reasons are not good is incorrect. Rather it reveals an ambiguity in the concept of having good reasons. Our initial assessment reflects a conception of having good reasons where “good” means (something like) ideally correct. *S* has good reasons in this sense only if his reasons are undefeated by the evidence he possesses. Our subsequent assessment reflects a conception of having good reasons where “good” means (something like) permissible. In this latter sense, *S* can have good reasons even if those reasons are defeated (by evidence he possesses), provided it is still (epistemically) permissible for *S* to believe for those reasons, i.e., provided the relevance of *d* is not obvious.

Obviousness is a psychological notion. Epistemic connections between propositions are obvious or not, only relative to a particular level of reasoning ability. What is obvious to an acute reasoner will not necessarily be obvious to an inferior reasoner. When we specify that the relevance of *d* is or is not obvious, we need to specify further to whom it is or is not obvious. This reveals an additional ambiguity in the concept of having good reasons. Generally when we say something is obvious *simpliciter*, we presuppose a level of reasoning ability that is intersubjectively determined – the normal reasoning ability of a relevant social group to which we belong. Reasons can be permissible grounds for belief, relative to that standard, even though they are not ideally correct. But viewing reasons as permissible relative to a standard determined by reasoning ability suggests still a weaker notion of having good reasons. Suppose *S* believes *q* on the basis of *r* because while the relevance of *d* is obvious *intersubjectively*, it is not obvious to *S*. We might say that *S* has good reasons in the sense that given *his* level of reasoning ability, it is (epistemically) permissible for him to believe *q*.

We can introduce some terminology to help disambiguate the concept of having good reasons. Let us call a defeater whose relevance is obvious, relative to a standard determined by the normal reasoning ability of a social group, an “intersubjectively evident” defeater. A defeater that is not intersubjectively evident we can call “intersubjectively opaque”. A defeater whose relevance is obvious relative to a standard determined by the subject’s own reasoning ability, we can call a “subjectively evident” defeater. A defeater that is not subjectively evident we can call “subjectively opaque”. We might then say that where *S* has prima facie good reasons to believe

that q , S has ideally good reasons just in case S possesses no defeaters of those reasons. S has intersubjectively good reasons just in case S possesses no intersubjectively evident defeaters. S has subjectively good reasons just in case S possesses no subjectively evident defeaters.⁶

Although roughly capturing the intuitive aspects, these three definitions of having good reasons are inadequate.⁷ Prima facie reasons have a more complicated structure than the suggested definitions can accommodate. In addition to defeaters of prima facie reasons, there can be defeater defeaters, defeater defeater defeaters, etc. S may believe q (the table is red), on the basis of r (the table looks red), while possessing d_1 (the architect says red lights were installed in the building) and yet still have good reasons for believing q because he also possesses d_2 (the maintenance records show that the red lights were replaced by white lights). Proposition d_2 defeats the defeating effect of proposition d_1 , i.e., $(d_1 d_2 r)$ is a good reason for S to believe q . However, S may also learn d_3 (the janitor habitually falsifies the maintenance records), and so on...

These characterizations of having good reasons must be amended to account for meta-defeaters. Consider ideally good reasons. On the proposed definition, there will be cases where S 's reasons should count as ideally good, yet they will fail to meet the definition. Suppose in the example just discussed, d_3 was not part of S 's evidence. Although S would still possess the defeater d_1 , we should still say that he has ideally good reasons since he possesses d_2 .

Let us call defeaters that undermine prima facie reasons (e.g., d_1 and d_3 in the previous example) "undermining defeaters". Defeaters that restore prima facie reasons (e.g., d_2) we can call "restoring defeaters". We can then say:

S has *ideally good reasons* to believe q iff S has prima facie good reasons r to believe q , and for every undermining defeater of r as a reason to believe q possessed by S , there is a restoring defeater possessed by S .

Matters can be handled analogously for intersubjectively good reasons and subjectively good reasons. The intuitive idea is that intersubjectively good reasons are reasons to believe that are permissible relative to an intersubjectively determined standard of reasoning, while subjectively good reasons are reasons to believe that are per-

missible relative to a standard of reasoning defined by the subject's own reasoning abilities. Thus we can say:

S has *intersubjectively good reasons* to believe *q* iff *S* has prima facie good reasons *r* to believe *q*, and for every intersubjectively evident undermining defeater of *r* as a reason to believe *q* possessed by *S*, there is an intersubjectively evident restoring defeater possessed by *S*.

S has *subjectively good reasons* to believe *q* iff *S* has prima facie good reasons *r* to believe *q*, and for every subjectively evident undermining defeater of *r* as a reason to believe *q* possessed by *S*, there is a subjectively evident restoring defeater possessed by *S*.^{8,9}

2.

What relevance do these different concepts of having good reasons have to the concept of knowledge? Given our assumption that having knowledge entails having good reasons, which senses of having good reasons preserve this entailment? I will argue that consideration of this question reveals a social component of knowledge. Eventually I will claim that this social component is best seen as indicating that attributions of knowledge are context-sensitive.

We can begin by considering undermining defeaters. Returning to the previous case, suppose *S* believes the true proposition *q* (the table is red) on the basis of his knowing *r* (the table looks red). *S* also knows *d*₁ (the architect says red lights were installed in the building), although unbeknown to *S*, the architect is mistaken (i.e., *r* is true because *q* is true).¹⁰ While the defeating effect of *d*₁ is obvious to *S*, he persists in believing *q* because of a strong bias he has in favor of *q*. If we stipulate that *S* possesses no further relevant evidence, it is clear that *S* fails to know *q*. This suggests that if *S* possesses a subjectively evident undermining defeater of *r* as a reason to believe *q* (that is not itself defeated by a restoring defeater), *S* fails to know that *q* on the basis of *r*.

Now suppose *d*₁ is subjectively opaque for *S*. That is, *S* believes *q* on the basis of *r*, even though he knows *d*₁, because *S* is too dense to appreciate the relevance of *d*₁. (If it strains the imaginations of the reader to suppose that *d*₁ could be subjectively opaque, then we can

let d_1 be: it is 1985 and the janitor says that if it is 1985, the lights have been changed; and if the lights have been changed they are red. We can suppose that S has difficulty stringing together conditionals. (Of course the reader can add conditionals if this aids the imagination.) Surely S still fails to know q . Intuitively, if S believes q only because obtuseness prevents him from discerning the relevance of an obvious piece of defeating evidence, S does not thereby know that q . In this connection, we use the term “obvious” intersubjectively. *Ex hypothesi* the relevance of d_1 is *not* obvious to S . This suggests that if S possesses an intersubjectively evident undermining defeater of r as a reason to believe q (which is itself undefeated by a restoring defeater), then even if the undermining defeater is subjectively opaque, S fails to know q on the basis of r .¹¹

Since an intersubjectively evident undermining defeater is defined as a defeater whose relevance is obvious relative to a socially determined standard, we might conclude at this point that knowledge has an interesting social component: if S possesses a defeater that meets those standards, and fails to adjust his beliefs in accordance with the defeater, then S fails to know. But such a conclusion would be premature. For it may be that where S possesses *any* undermining defeater of his reasons (without possessing a restoring defeater), he fails to know. That is, it may be that S knows that q only if S has ideally good reasons for believing q .

Does knowledge entail ideally good reasons? The test case will be one where S possesses an intersubjectively opaque undermining defeater without possessing a restoring defeater. Intuitions might differ depending on how intersubjectively opaque the undermining defeater is. However, to test whether knowledge entails ideally good reasons rather than merely intersubjectively good reasons, we must consider undermining defeaters of arbitrary opacity. If knowledge is subverted by defeaters only up to a certain level of intersubjective opacity, it remains true that the level at which defeaters undermine knowledge is socially determined. Although the degree of opacity would not be fixed *at* intersubjective levels, it would still be fixed *by* intersubjective levels (in terms of how much beyond the intersubjective level it is).

Thus we should consider defeaters of arbitrary opacity. Suppose S possesses an undermining defeater of r (the table looks red) as a reason to believe q (the table is red) whose relevance escapes the

closest scrutiny of extremely acute subjects. We can imagine that only a super genius would see it. If *S* does not possess a restoring defeater, does he know that *q* on the basis of *r*? (Of course, I intend that in these cases, the reasoning ability of the subject does not exceed intersubjective standards.) When considering this case, we must remember that *S* believes the table is red on the basis of its looking red, and the table looks red because it is red. It seems that given this, *S* does not fail to know *q* (the table is red), simply because *S* possesses a misleading defeater whose relevance is discernible only through an inscrutable line of reasoning.

If knowledge does *not* entail ideally good reasons, then it looks as if whether a defeater undermines knowledge is in part socially determined. There is no doubt that intuitions will diverge as to whether knowledge does entail ideally good reasons. One could hold the view that for *S* to know that *q*, *S*'s total evidence must support *q* (which is true just in case *S* has ideally good reasons to believe *q*). Nonetheless, I do think that the intuition is much stronger that *S* fails to know when *S* possesses an intersubjectively evident defeater than when *S* possesses an intersubjectively opaque defeater. This in itself supports the view that socially determined standards play a role in determining whether *S* has knowledge.

On the assumption that knowledge entails ideally good reasons, the fact that an intersubjectively evident undermining defeater precludes knowledge turns out to be merely a specific instance of this general requirement. I will now argue that the social component of knowledge emerges even given that assumption. To see this, we need to consider the interaction of undermining defeaters with restoring defeaters. We can use our previous instantiations for *r*, *q*, and *d*₁, except that we will use the notation '*d*_u' instead of '*d*₁' to indicate that the defeater is undermining. Again *S* believes *q* on the basis of *r* while possessing *d*_u, but to meet the requirement that *S* have ideally good reasons, we suppose *S* possesses a restoring defeater *d*_r. Does that allow *S* to know *q* on the basis of *r*? This depends on facts concerning evidentness. Suppose *d*_u is subjectively evident but as in our original case *S* still believes *q* because of a bias. If *d*_r is subjectively opaque, although *S* has ideally good reasons, *S* fails to know.¹² Without *d*_r, *S* fails to know because of the obvious (given *S*'s own reasoning abilities) effect of *d*_u on his reasons. Although *d*_r restores *S*'s reasons from an ideal perspective, *S* still fails to know since the defeating (restoring) effect of *d*_r

is beyond his ken (d_r is subjectively opaque). This suggests that 'S knows that q ' entails that 'S has subjectively good reasons to believe q ' (S has prima facie good reasons and for every subjectively evident undermining defeater possessed by S, there is a subjectively evident restoring defeater possessed by S).

As we have noted, defeaters can be subjectively opaque but intersubjectively evident. If d_u has this status, what status must the restoring defeater have in order for S to know? Again we suppose that S knows that d_u , but being rather dense (by socially determined standards) he fails to appreciate its relevance, and so believes q on the basis of r . If S possesses a restoring defeater that is subjectively opaque (and we can suppose intersubjectively opaque as well), he will still have ideally good reasons for believing q . But, since he does not appreciate the effect of the restoring defeater, S fails to know that q . Although his obtuseness results in his failing to appreciate the relevance of d_u , surely it does not result in his knowing that q .

This strongly suggests that S knows that q only if for every intersubjectively evident undermining defeater possessed by S, there is a subjectively evident restoring defeater possessed by S.¹³ Thus it is not enough to require that S have ideally good reasons and subjectively good reasons in order to know. For S can have ideally good reasons and subjectively good reasons and yet still fail to know as a result of failing to meet intersubjective standards.

Again, we don't want to draw this conclusion too hastily. If S knows that q only where S possesses a subjectively evident restoring defeater for *any* undermining defeater he possesses, then intersubjective standards would be irrelevant. The above result would turn out to be a specific instance of this general requirement. But surely this general requirement is too strong. It entails that S fails to know q on the basis of r when:

- (1) r is a prima facie good reason to believe q .
- (2) The only undermining defeater S possesses is extremely intersubjectively (and subjectively) opaque (S has intersubjectively and subjectively good reasons).
- (3) The undermining defeater is itself defeated by a restoring defeater possessed by S, so that S's total evidence supports q (S has ideally good reasons).

Let us return to our example. *S* believes *q* (the table is red), because he knows *r* (the table looks red). Now we suppose that *S* has evidence *d_u* which supports a long, complicated line of reasoning to the conclusion that color vision is unreliable in region *m*, where the table in question is located in *m*. But *S*'s evidence *d_r* also supports, by an equally opaque line of reasoning, the conclusion that color vision is reliable in *m'*, where *m'* is a subregion of *m* and the table is located in *m'*. Here *S* possesses an (intersubjectively opaque) undermining defeater but he does not possess a subjectively evident restoring defeater. But surely *S* still knows that *q*.

Why does *S* know that *q* when an intersubjectively opaque undermining defeater he possesses is defeated by a subjectively opaque restoring defeater, but fail to know that *q* when the undermining defeater is intersubjectively evident? On the view we are entertaining, *S* knows *q* only if *S* has ideally good reasons to believe *q*. Both cases meet this requirement. The reason we balk at granting knowledge in the latter case is that *S*'s belief is, in a sense relevant to knowledge, impermissible. By believing *q* in the face of an obvious defeater of his reasons for *q*, without *seeing* that this defeater is itself defeated, *S* fails to meet certain intersubjective standards for belief revision. In the former case, no such violation of standards is involved – the defeater is too opaque.¹⁴

We can conclude that knowledge has a social component. Whether a person's reasons give him knowledge depends on intersubjective standards for discerning the effects of defeaters. If knowledge does not entail ideally good reasons, then the intersubjective standards determine the level of opacity up to which a possessed undermining defeater (that is itself undefeated) will undermine knowledge. If knowledge does entail ideally good reasons, then intersubjective standards will determine the level of opacity up to which a possessed undermining defeater will undermine knowledge without the possession of a subjectively evident restoring defeater.

3.

Does the point I have made concerning intersubjective standards for reasoning generalize to other cognitive faculties, viz., memory and perception? In the case of reasoning, I have claimed that a defeater possessed by *S* undermines his knowledge only if the defeater is

intersubjectively evident (on the simplifying assumption that knowledge does not entail ideally good reasons). Of course there are suppressed temporal indices in this formulation – a defeater possessed by *S* at (a time) *t* undermines his knowledge at *t*. However, we might consider the possibility that in assessments of knowledge, we should view a person's evidence diachronically. It may be that defeaters possessed by *S* prior to *t*, but no longer possessed by *S* at *t*, can undermine *S*'s knowledge at *t*. If so, intersubjective standards for memory may be relevant. Moreover, it has been claimed that evidence never possessed by *S* can undermine his knowledge.¹⁵ Here, intersubjective standards for perception may be relevant. We can begin by considering memory.

In the course of our lives we accumulate a vast amount of information. Being creatures with limited storage capacity, our memories fail to retain much of what we learn. Thus, we are often unable to adjust our beliefs at a time *t* in accordance with evidence we possessed prior to *t*. How does such evidence bear on what we know at *t*?

Consider a case where at time t_1 , *S* possesses *d*, where *d* is a defeater of *r* as a reason to believe *q*. At some later time t_2 , *S* comes to know *r* and so infers *q* (where *q* is true). Assuming *d* is subjectively evident, does *S* know that *q*? This depends on whether *S* remembers *d*. If he does, then he fails to know *q*. But suppose he has forgotten *d*. We can consider a concrete case:

- r*: The 1960 Rand McNally Atlas lists Austin as the capital of Texas.
- q*: Austin is the capital of Texas.
- d*: A reliable source says the 1960 Rand McNally Atlas contains numerous errors regarding state capitals.

If *S* has forgotten *d*, does he then know *q* on the basis of *r*? (We must assume that the reliable source is mistaken or *S* may fail to know for standard Gettier reasons.) This depends on the length of the interval t_1-t_2 . It is plausible to hold that if it is very short, *S* does not know. If *S* learns that *d* and then 30 seconds later comes to believe *q* on the basis of *r*, because he's forgotten *d*, then he does not know that *q*.

On the other hand, suppose *S* learns that *d* and 25 years later learns *r*. Because he has forgotten *d*, he infers *q* from *r*. Then, *S* does know that *q*. (At least the intuition that he does is considerably stronger than in the previous case.)

The disparity in our assessment of these two cases suggests the existence of intersubjective standards for memory. We could say that while in both cases *S*'s reasons (viewed diachronically) are subjectively good, only in the latter case are they intersubjectively good. Here the reasons fail to be intersubjectively good, not because of a failure to meet a standard determined by intersubjectively normal reasoning ability, but rather because of a failure to meet a standard determined by intersubjectively normal memory ability. The reasons are not permissible grounds for belief given those standards. As a result, *S* fails to know on the basis of those reasons.

Again, intuitions may diverge regarding these cases. I can see someone taking a hard line and denying knowledge to *S* whenever he once possessed an (undefeated) defeater regardless of the amount of time it has been since *S* possessed the defeater. Such a view is analogous to the view that *S* fails to know whenever he possesses an (undefeated) defeater, no matter how opaque. One can hold that knowledge entails ideally good reasons diachronically as well as at a particular time.

If we make this assumption for memory, then, as we did in the case of reasoning, we will have to consider the interaction of undermining defeaters and restoring defeaters in order to settle the question of whether there are intersubjective standards. I leave it to the reader to construct cases of this kind.

In his interesting paper 'How Do You Know?', Ernest Sosa provides an example that would seem to illustrate the same point for perception and unpossessed defeaters that I have argued for reasoning and possessed defeaters, and memory and once possessed defeaters. He discusses a case where

... despite his extensive experience with cable cars, Mr. Magoo [who is nearly blind and deaf] does not know that his cable car will arrive safely when, unknown to him, bombs are raining all around it.¹⁶

Sosa draws the general conclusion that a subject can fail to know if he has defective cognitive equipment which prevents him from acquiring defeaters "that a normal inquirer in the epistemic community would acquire in that situation". In the case of perceptual equipment, this is just to say that there are intersubjective standards for perception of the sort I have argued apply to reasoning and memory.

I am not sure whether the point can be extended to perception. Closer inspection reveals that the Magoo case does not show that it

can. Although the case elicits a strong intuition that Magoo fails to know, there are other factors besides Magoo's defective equipment that could explain this. Notice that Magoo would fail to know even if his equipment were normal but his view of the bombs were obstructed and the noise of the cable car were to mask the sound of the explosions. The mere fact that the bombs are raining all round is sufficient to undermine his knowledge. This can perhaps be explained in terms of an objective probability that Magoo's belief is false.¹⁷

Moreover, the intuition that Magoo fails to know may be elicited by the assumption that others witness the explosions and thereby doubt whether the car will arrive safely. If so, the case only illustrates the point that knowledge can be undermined by unpossessed evidence, if the evidence is possessed by a relevant social group.¹⁸

To test Sosa's conclusion, we should eliminate these elements from the case. Suppose that instead of actual bombs falling, there is an elaborate hoax involving powerful hidden speakers which project the sound of explosions and huge hidden fans creating dust clouds. Magoo is unaware of this misleading evidence due to his defective perceptual equipment. We must also suppose that no one else is in a position to witness the "explosions". Does Magoo fail to know the cable car will arrive safely? I think this is much less clear.

I think the case for intersubjective standards is strongest in the case of reasoning, less clear for memory, and less clear still for perception. This demonstrates the importance of the subject actually possessing the defeaters. In the case of reasoning, the defeaters are actually possessed. In the case of memory, the defeaters although once possessed, no longer are, and in the case of perception the defeaters are not and never were possessed.

4.

I have been arguing that whether *S* knows that *q* depends, in part, on intersubjectively determined standards. This raises the issue of how precisely those standards are determined. In particular, which social group sets the standards? A natural suggestion is that the relevant social group is the one to which the subject belongs. Thus, whether *S* knows that *q* in the face of an undermining defeater d_u will depend on whether d_u is obvious to *S*'s own social group. This proposal would seem too weak. If *S* possesses an undermining defeater whose rele-

vance is obvious to us, we would still judge *S* as failing to know even if he belongs to a society of morons to whom the undermining defeater is opaque. (If knowledge entails ideally good reasons, we can stipulate that *S* possesses a restoring defeater opaque to us and the morons.)

Perhaps, then, it is our own social group that sets the standards. But surely the moron society would not use our standards. And a genius society would surely make its epistemic assessments according to standards consonant with their superior reasoning powers. Is “knowledge” then ambiguous between various concepts each based on a different standard? This would entail an indefinite number of concepts of knowledge. It would also entail that, were our reasoning powers to improve or decline, our concept of knowledge would change.

A better way to view matters is to suppose that attributions (or denials) of knowledge are indexical or context sensitive. The standards that apply are determined by the context of attribution. The truth-value of a knowledge attribution will depend on the status of the defeater the subject possesses, relative to the standards that apply in the context of attribution.

In general, the standards in effect in a particular context are determined by the normal reasoning powers of the attributor’s social group. Thus, I may correctly deny knowledge to *S* where a member of the moron society correctly attributes knowledge to *S*. Similarly, I may correctly attribute knowledge to *S* where a member of a genius society correctly denies knowledge to *S*. Because the attributions are context sensitive, there is no contradiction.

This raises the question of *which* social group of the attributor (*A*) is the relevant one. Is it the society at large in which *A* lives?... his professional circles? Perhaps the standards that apply are determined by *A*’s own reasoning ability (in which case they are not intersubjective at all). I am not sure how to decide this. Presumably, the standards are variably determined in each of these ways, although it is unclear what mechanisms govern the shifts.¹⁹

In some contexts, the attributor *A* will have explicit intentions as to which standards apply. For example, in philosophical situations where cases are described (like the cases in sections 1 and 2), the specification of the abilities of the subject and the evidentness of the defeaters will naturally raise the issue of standards. In such cases, the intentions of *A* become a relevant feature of the context. I suspect

that as a purely psychological matter, *A* will generally intend standards that accord with the abilities of some social group to which he belongs. (Or if *A* is unique in his reasoning abilities, he may intend standards that reflect his own abilities.) This explains our tendency to deny knowledge in cases where it is specified that defeaters are possessed which are obvious to us. But nothing constrains the attributor from intending other standards. Where he does, those standards may become the ones that apply. Certainly, where it is specified that *S* possesses a defeater obvious to me but not to the morons, or obvious to the geniuses but opaque to me, I could agree with the morons or the geniuses in their knowledge attributions by adopting their standards for my knowledge attribution. If the intentions of the attributor were not a feature of some contexts, this would not be possible. Thus what I have been calling intersubjective or socially determined standards need not be such at all. The attributor may intend various standards (including idealized ones), and those intentions may determine which standards apply.²⁰ This does not mean that *A* thereby makes his attribution true, since the truth-value of his attribution will be a function of both the standards of evidentness he intends and the actual evidentness of the defeaters possessed by the subject of the attribution.

Compare “know” with a term like “flat”. The point has been made that the truth-values of attributions of flatness can vary depending on what standards are applied. According to some standards, certain things will count as bumps that would not so count according to relaxed standards.²¹ I may look out my window and claim that New Jersey is hilly while a giant may assess New Jersey as flat. Each claim can be correct. There is no contradiction, since the contexts of attribution yield different standards. Surely I would not want to claim that what the giant says is false owing to his distorted perspective or my claim that the road is flat would be subject to the same assessment by an ant-sized being. All of this is familiar enough. Again it is important to see that neither we nor the giant is constrained to use specific standards. I could adopt the giant’s standards and agree with him by truthfully stating that New Jersey is flat. This does not conflict with the psychological fact that routinely our own personal/social perspective determines which standards we intend.

This is just what I want to say about ‘know’ and standards of

evidentness for defeaters. Attributions of knowledge are relative to such standards, and the particular standards that apply for a given attribution are determined by context.^{22,23}

5.

Skepticism is a problem that any epistemology must face. In general, when we propose theories of knowledge, we presume that knowledge is something that we have in everyday situations. Skeptical arguments threaten this presumption. Such arguments, if sound, show that we fail to know all or most of what we ordinarily take ourselves to know.

One need only consider the numerous attempts to refute skeptical arguments, and the lack of agreement as to the success of the attempts, to see that skeptical arguments are very powerful. Some philosophers, intent on constructing theories of knowledge, yet unable to respond directly to skeptical arguments, settle for an appeal to common sense. Many find this approach at worst question-begging and at least unsatisfying. It would be nice to have an account of where the skeptical arguments go wrong.

I propose to discuss one such account that has occurred in various forms in the epistemological literature. One difficulty with this account is that it can appear unmotivated. I will argue that the results of section 4 provide a motivation.

We can begin by considering how skeptical worries arise. Suppose it is claimed that *S* knows that *q* (*S* sees a barn), on the basis of *r* (It looks to *S* as if *S* sees a barn). Let an alternative to *q* be a proposition incompatible with *q*. The skeptic points out that there are alternatives *h* to *q* that *S* is not in a position to rule out on the basis of *r* (e.g., *S* sees a papier-mache replica of a barn; *S* is hallucinating) or any other evidence (e.g., *S* is a brain-in-a-vat; *S* is deceived by a Cartesian demon). But if *S* is not in a position to know that not *h*, where *h* is an alternative to *q*, then *S* cannot be said to know that *q*.

Of course one could always insist that *S* does know that *q* on the basis of *r*, and since *q* entails not *h* (and *S* knows this), *S* knows that not *h*.²⁴ So it is best to state the skeptical problem as a paradox. A paradox comprises a set of inconsistent propositions, all of which have independent plausibility. Here is such a set (given that *S* knows that *q* entails not *h*):

- (1) S knows that q .
- (2) If S knows that q and S knows that q entails not h , then S knows that not h .²⁵
- (3) S does not know that not h .

The skeptic uses (3) and (2) to infer not (1). But as we have noted, one can resist the skeptical conclusion by using (1) and (2) to infer not (3). Of course it is also open to one to retain (1) and (3) by denying (2). But to simply choose one of these alternatives is not to resolve the paradox.

I want to focus on this last approach that denies proposition (2), the closure principle. An anti-skeptical strategy of this sort originates in the work of Dretske and variations of it occur in papers by several authors.²⁶ The fundamental idea of this approach is that S can know q on the basis of r even if he cannot rule out alternative h , provided h is not a *relevant* alternative. Skeptical alternatives, although not known to be false, fail to be relevant.

To give content to this approach we need a criterion of relevance. One that is suggested by some of these philosophers is objective probability.²⁷ An alternative h that cannot be ruled out is relevant just in case the conditional probability of h on r (and other features of the circumstances) is suitably high. Consider a case discussed by Goldman.²⁸ Under what circumstances is the inability of S to rule out the alternative that he sees a papier-mache barn replica relevant to whether he knows that he sees a barn on the basis of it appearing to him as if he sees a barn? (Under what circumstances does the failure of S to know that it is not the case that he sees a barn replica prevent him from knowing that he sees a barn on the basis of his perceptual evidence?) If there are numerous barn replicas in the immediate area, then this alternative is relevant. The probability that S sees a barn replica conditional on his evidence and the existence of the numerous barn replicas is suitably high. Thus, S fails to know that he sees a barn unless he is in a position to rule out this alternative. Where no such replicas exist, this alternative is not relevant (*ceterus paribus*), i.e., S can know without being in a position to rule it out.²⁹

Of course there is considerable vagueness here. How many barn replicas must there be? How circumscribed a region must they occur in?³⁰ But there is no reason why relevance cannot be a vague notion, provided that skeptical hypotheses are sufficiently remote in the circumstances to count as clear cases of irrelevant alternatives.³¹

The means for denying proposition (2) of the skeptical paradox are now available. The issue whether *S* knows that *q* on the basis of *r* is always to be assessed relative to probabilistic standards of relevance for alternatives (that cannot be ruled out). *S* knows that *q* on the basis of *r*, relative to standards that determine that skeptical alternative *h* is not relevant (in most everyday circumstances). That is, *S* knows that *q* relative to standards that allow him to know that *q* without knowing that not *h*, even though he may know that *q* entails not *h*. Thus we can affirm propositions (1) and (3) without inconsistency.³²

Is this a satisfactory resolution of the paradox? One defect is this: if knowledge attributions are always relative to these standards of relevance, why do we ever take the threat of skepticism seriously? It is true that we sometimes feel that the skeptical alternatives are too remote to threaten our knowledge claims, but it is very easy to begin to worry that this dismissal of skeptical alternatives is cavalier.

One way to explain our tendency to vacillate between skepticism and common sense is to suppose that the probabilistic standards in effect (and thus the relevance of alternatives) depend on the context of attribution. When *A* attributes (or denies) knowledge that *q* to *S*, the standards that operate in that context determine which alternatives are relevant (in combination, of course, with the actual probabilities in the circumstances).³³ So the truth of a particular knowledge attribution depends on the standards that operate in that context. In day to day contexts, when skeptical worries are not an issue, the standards that apply determine that skeptical alternatives are not relevant (unless of course the actual probabilities in the circumstances of evaluation are sufficiently high). This explains our confidence in the truth of our everyday attributions of knowledge (i.e., our acceptance of proposition (1)).

However, we are not constrained to use those standards in our knowledge attributions. Where *A* attributes knowledge that *q* to *S*, *A*'s intentions can determine which standards apply in the context. In normal, everyday contexts *A* will not have intentions regarding standards. However, when skeptical worries are raised, *A*'s intentions can lower the standards so that skeptical alternatives become relevant. Indeed, skeptical arguments are forceful precisely because they get us to consider skeptical alternatives as relevant. But again we are not constrained to use skeptical standards. If upon further consideration we decide that skeptical alternatives really are too remote, we may no

longer consider them relevant. Often we vacillate between considering skeptical alternatives relevant and dismissing them as irrelevant.

With this explanation of the appeal of skeptical arguments, do we have a satisfactory resolution of the skeptical paradox? No doubt some will feel it is ad hoc. Do we have any reason to think that this view is correct, independent of the fact that by taking it, we can preserve the integrity of our everyday knowledge attributions in the face of a seemingly powerful skeptical argument? Surely the skeptic will deny that knowledge attributions are context sensitive. He will insist that our confident, everyday knowledge attributions result from our failure to consider the possibility of skeptical alternatives. Moreover, he will attribute the fact that we sometimes continue to ignore them once he points them out to our psychology and not to our rationality.

Do we have a standoff then? Perhaps it is too much to expect a knockdown argument against skepticism. All we can hope for is a workable option that enables us to avoid skepticism. We avail ourselves of that option for no other reason than we want to avoid skepticism.

Still, it would be reassuring to have an independent motivation for the anti-skeptical view. If we were not motivated by skeptical worries, would we have any reason to think that a view of this kind is correct?

The results of section 4 provide us with just those reasons. There, we reached the conclusion that the truth-value of a knowledge attribution is relative to a context sensitive standard, independent of any skeptical worries. (There is no threat that our knowledge attributions will always be false as a result of our possessing undermining defeaters.) Closer inspection reveals even more specific similarities.

The anti-skeptical strategy consists in holding that the truth-value of an attribution of knowledge is relative to a context sensitive standard which determines, together with certain features of the circumstances of evaluation, the relevance of alternatives. The standard is a probability measure and the features of the circumstances are the actual probabilities of the alternatives.

Now consider the conclusion of section 4: the truth-value of an attribution of knowledge is relative to a context sensitive standard which determines, along with certain features in the circumstances of evaluation, which defeaters undermine knowledge. The standard is a measure of evidentness and the features of the circumstances are the actual evidentness of the defeaters possessed by the subject of the

attribution. So far, the two conclusions differ, insofar as in the former it is the relevance of alternatives that is at issue, whereas in the latter, it is the undermining effect of defeaters. However, we can view (type II) undermining defeaters as operating precisely by making certain alternatives relevant. For example, where *S* believes that *q* (*S* sees a barn) on the basis of his visual evidence *r*, *d* (there are numerous barn replicas in the area) is a defeater of *r* as a reason to believe *q*. That is, *S* possessing *d* (without possessing a restoring defeater) undermines his knowledge that *q* on the basis of *r*. The reason for this is that *S* possessing *d* makes relevant (relative to normal standards) the alternative *h* (*S* sees a barn replica). Where *S* possesses *d*, *S* cannot know that *q* on the basis of *r*, unless he knows that not *h*. Similarly where *S* possesses the defeater (*S* has ingested a hallucinogen), the alternative (*S* is hallucinating) becomes relevant to his knowing on the basis of his perceptual evidence.³⁴

With this in mind, we see that the conclusion of section 4 can be characterized thus: the truth-value of an attribution of knowledge is relative to a context sensitive standard (of evidentness) that together with certain features of the circumstances of evaluation (the actual evidentness of the defeaters possessed by *S*) determines the relevance of alternatives. This conclusion differs from the anti-skeptical view only with respect to the nature of the standard (and of course the feature of the circumstances governed by the standard). The existence of *this* difference simply reveals that there are two ways in which alternatives become relevant, viz., by the subject possessing a defeater or by the alternative being objectively probable. In either case, the relevance is governed by a context-sensitive standard. The conclusion that there is a context sensitive standard of *evidentness* that determines the relevance of alternatives lends credence to the view that there is a context sensitive standard of *probability* that determines the relevance of alternatives. The fact that the former view was reached independently of any skeptical worries alleviates the concern that the latter view is an ad hoc response to such worries. Both can be seen as instances of the same general phenomenon.

Is this a refutation of skepticism? Of course not. The skeptic will reject all the intuitions marshalled in support of both conclusions. But, of course, the skeptic cannot *refute* common sense. If we seek to defend common sense, the object is to develop an independently motivated framework that allows us to reject skepticism. And this is what I am proposing we have.

NOTES

* A shorter version of this paper, entitled 'Knowledge and Context', was presented as a symposium at the 1986 Eastern Division meeting of the American Philosophical Association and appeared in the *Journal of Philosophy* **83** (October 1986) 574–83. Because of space limitations, much of the material in sections 1 and 2 of this longer version was compressed into footnotes. Section 3 and many of the footnotes were deleted entirely and the argument of section 5 was presented in a condensed form. I am grateful to Frederick Schmitt and the editors of *Synthese* for the opportunity to present the paper in its original form.

¹ Gilbert Harman: 1980, 'Reasoning and Evidence One Does Not Possess', in Peter A. French, Theodore E. Uehling, Jr., Howard K. Wettstein (eds.), *Midwest Studies in Philosophy*, Vol. V, University of Minnesota Press, Minneapolis, pp. 163–82.

² My conclusion holds, as well, for reliability theories that deny that knowledge entails having good reasons. See note 22.

³ This is not quite correct. *S* can have good reasons *simpliciter* if he possesses defeaters which are themselves defeated. More of this later.

⁴ John Pollock: 1974, *Knowledge and Justification*, Princeton University Press, Princeton. In a later paper, Pollock calls type I and type II defeaters, rebutting and undercutting defeaters, respectively. See Pollock: 1984, 'Reliability and Justified Belief', *Canadian Journal of Philosophy*, pp. 409–22. I avoid this helpful terminology to preclude terminological confusion in my own paper.

The notion of defeaters of justification is discussed by Keith Lehrer and Thomas Paxson, Jr.: 1969, 'Knowledge: Undefeated Justified True Belief', *Journal of Philosophy* **66**, 225–37.

⁵ I use the term "obvious" here, because there is no epistemic requirement that subjects devote thought to searching for subtle defeaters in their evidence (in order to know). That would make knowledge practically unattainable. However, we do hold subjects responsible for obvious defeaters.

Of course "obvious" is a vague notion.

⁶ What is it to *possess* a defeater? We can view defeaters as propositions and, for the most part, define possession in terms of belief. In certain cases, defeaters may be possessed without being believed. So, if "it appears to *S* as if. . ." is a defeater, the mere fact that it does appear to *S* as if. . . , can constitute possession of the defeater even if *S* does not have a belief that "it appears to *S* as if. . .".

Does *S* possess a defeater *d*, when he possesses good reasons *r* to believe *d* but fails to believe *d*? In such cases, we can treat *r* itself as the defeater possessed by *S*. His failure to believe *d* may be explained by *r* being subjectively opaque, or perhaps by *S* having a bias against *d*.

It may be that where defeaters are believed, they must be believed with good reason in order to be possessed (see note 8).

⁷ If the reader does not find any intuitive basis for these distinctions, they can be viewed as stipulative. It will not matter for my argument.

⁸ It does not seem that an undermining defeater must be believed *with good reason* in order to be possessed. If *S* believes, *even irrationally*, that the room has red lighting, it would seem that the table appearing red to him is defeated as a reason to believe the table is red, i.e., *S* cannot know on the basis of this reason.

In the case of restoring defeaters, it looks more plausible to suppose that the defeaters

must be believed with good reason in order to be possessed. If the janitor tells *S* that the lights are red, and *S* arbitrarily believes that he is lying, *S*'s reasons are not restored, i.e., *S* cannot know on the basis of this reason. Issues like these may reveal further concepts of having good reasons.

These considerations introduce the much discussed regress of reasons. For our purposes we need not take a stand on this issue. The fact that defeaters and the prima facie reasons themselves are based on reasons allows for the possibility of hybrid reason chains.

Also, I assume, in these definitions, that the sequence of defeaters is finite. It is not clear how to characterize the reasons in a case where there is an infinite sequence alternating between undermining defeaters and restoring defeaters (if such a case is possible).

⁹ Additional concepts of having good reasons can be distinguished. For example, defeaters could operate through fallacious arguments, where the fallaciousness is (subjectively or intersubjectively) opaque. The prima facie reasons themselves can be (subjectively or intersubjectively) opaquely fallacious, or their correctness can be (subjectively or intersubjectively) opaque. None of these distinctions is relevant to my argument.

¹⁰ We must suppose this or else we would have a standard Gettier case.

¹¹ One might think that a subject as dense as *S* is incapable of knowing anything. This would make the explanation that he fails to know because he possesses an intersubjectively evident defeater otiose. But why should *S* fail to know the table is red when he sees it under perfectly good perceptual conditions, when he does not possess any defeaters? Surely not because his feeble reasoning powers would make it difficult for him to appreciate a defeater were he, *contrary to fact*, to possess one.

¹² Remember that "evidentness" is defined in terms of obvious relevance. If *S* should discover the relevance of *d*, after a period of reflection, then he would know. Of course we are to suppose that *S* does not do this.

¹³ This condition is stronger than "knowledge entails intersubjectively good reasons". For the latter condition only requires *intersubjective* evidentness for the restoring defeater. The stronger condition entails the latter, in cases where subjective reasoning powers do not exceed intersubjective powers.

My argument for intersubjective standards does not require the stronger conclusion. If it is claimed that where *S* possesses an intersubjectively evident undermining defeater, only an intersubjectively (rather than a subjectively) evident restoring defeater is required, then intersubjective standards would apply doubly – to the level of undermining defeaters that requires a restoring defeater of a specific status and to the nature of that status. (See what follows in the text.)

¹⁴ Since knowledge entails intersubjective standards, do the subjective standards turn out to be merely a consequence of this requirement? No, because it is possible that subjective reasoning powers could exceed intersubjective powers. In such a case subjectively good reasons would still be required.

¹⁵ Gilbert Harman: 1974, *Thought*, Princeton University Press, Princeton; Ernest Sosa: 1974, 'How Do You Know?', *American Philosophical Quarterly* **11**, 113–22.

¹⁶ Sosa, p. 16.

¹⁷ Cf. Harman: 'Reasoning and Evidence One Does Not Possess', *Midwest Studies in Philosophy*; Alvin Goldman: 1976, 'Discrimination and Perceptual Knowledge', *Journal of Philosophy* **73**, 771–91; Marshal Swain: 1978, 'Reasons, Causes, and Knowledge',

Journal of Philosophy 75, 229–49.

¹⁸ Harman: 1976, 'Reasoning and Evidence One Does Not Possess'.

¹⁹ One possibility is that the standards are determined by the attributor's intentions. Most plausibly, these intentions could be ascribed to attributors only in some implicit sense. For example, the standards could be fixed by counterfactuals concerning whether the attributor *A* would have attributed knowledge to the subject *S*, were *A* to have believed that *S* possessed defeaters of a certain kind.

Perhaps the context doesn't strictly determine the standards either because the attributor's intentions are vague or the general conventions of the language are indeterminate in this way. If so, then attributions of knowledge can be evaluated at different standards, or more precisely, at [world, time, standard] triples.

It should be noted that this lack of precision is no special problem for the claim that attributions of *knowledge* are context-sensitive. The mechanisms of context-sensitivity are not very well understood in general. Other (relatively) uncontroversial cases of predicates whose applications depend on context-sensitive standards face the same difficulty, e.g., flat. See the discussion that follows in the text.

²⁰ It is unclear just when *A*'s intentions can override the standards normally in effect, especially in conversational contexts. Is a *private* intention enough, or must the intention, in some way, be made apparent to the other participants in the conversation?

David Lewis: 1979, 'Scorekeeping in a Language Game', *Journal of Philosophical Logic* 8, 339–59 (reprinted in *Philosophical Papers* Vol. 1, Oxford, 1983), proposes some rules that determine how context-sensitive standards shift according to conversational dynamics. Lewis argues that in some conversational contexts the intentions of the speaker are to no avail.

²¹ Lewis: 1979, 'Scorekeeping in a Language Game', Fred Dretske: 1981, 'The Pragmatic Dimension of Knowledge', *Philosophical Studies* 40, 363–78. Both of these authors draw analogies between "flat" and epistemic terms.

Peter Unger: 1975, *Ignorance*, Oxford University Press, Oxford, pp. 65–68, takes a different line, arguing that (almost) nothing is flat. Drawing an analogy with epistemic terms like "certain", Unger argues for a skeptical conclusion.

²² Since reliability theories of knowledge generally eschew any requirement that *S* have good reasons, they may seem to be immune to the sort of context sensitivity I argue for. However, this is not true. Since a belief can be reliably produced, or a reliable indication of its truth, even where the subject possesses a defeater, reliability theories must have provisions that account for the effects of defeaters. And if I'm right about the relevance of defeaters being context sensitive, then any adequate reliability theory will have to allow for context sensitivity in the way it handles defeaters.

For example, Alvin Goldman: 1979, 'What is Justified Belief', in George Pappas (ed.), *Justification and Knowledge*, D. Reidel, Dordrecht, pp. 1–23, attempts to handle defeaters in terms of the subject's failure to use an available reliable process, viz., the proper use of evidence (p. 20). The context sensitivity will enter in to the conditions under which that process is available; e.g., is it available if the defeater is subjectively or intersubjectively opaque?

For a good discussion of Reliability theories and defeaters, see Arthur F. Walker: 1986, 'Justified Belief and Internal Acceptability', *Canadian Journal of Philosophy* 16, 493–502.

²³I am indebted to Scott Soames for many helpful discussions regarding the nature of context-sensitivity in general and its application to knowledge attributions.

²⁴ See John Pollock, *Knowledge and Justification*, p. 46.

For a modern version of the skeptical line, see Keith Lehrer: 1971, 'Why not Scepticism?', *The Philosophical Forum* 2, 283–98.

²⁵ This closure principle may be subject to certain quibbles. But, surely, something very close to it is very intuitive. For a discussion of this see Robert Nozick: 1983, *Philosophical Explanations*, Harvard University Press, Cambridge, pp. 205–06.

²⁶ Dretske: 1981, 'The Pragmatic Dimension of Knowledge', and 1976, 'Epistemic Operators', *Journal of Philosophy* 67, 1007–1023; Lewis: 1979, 'Scorekeeping in a Language Game'; Goldman: 1976, 'Discrimination and Perceptual Knowledge'; Harman: 1980, 'Reasoning and Evidence One Does Not Possess'; Nozick, *Philosophical Explanations*.

²⁷ Harman: 1980, 'Reasoning and Evidence One Does Not Possess'; Goldman: 1979, 'What is Justified Belief'; Dretske: 1981, 'The Pragmatic Dimension of Knowledge'; Swain: 1978, 'Reasons, Causes, and Knowledge'.

²⁸ Goldman: 1976, 'Discrimination and Perceptual Knowledge'.

²⁹ There is an intractable problem here. Probability $h/(r$ and other factors) will of course vary as we vary the other factors and there is no obvious way to be precise about what those other factors are. For example, the probability that S sees a barn replica given his evidence and his location in an area where there are many barn replicas and few barns, is high. But that same probability conditional on his evidence and his particular location in front of a real barn is quite low. How do we state in some general way what these other factors are?

At least this much is true. There is some standard according to which the alternative that S sees a barn replica is relevant if, e.g., there are many barn replicas in the immediate area, and not relevant if there is just one barn replica at the North Pole.

³⁰ Goldman refers to these considerations in his discussion of the case.

³¹ Under what circumstances would skeptical alternatives be relevant (relative to normal standards)? If frequently we hallucinate or are deceived by a demon, then these alternatives can be relevant. If it frequently happens that members of my society are kidnapped and turned into brains-in-a-vat, then that alternative can be relevant.

³² It may seem that given the standards of relevance, S does know that not h , since h fails to be relevant. On this view, closure is preserved and the relevant alternatives approach amounts to a denial of proposition (3) of the skeptical paradox. This way of proceeding has been endorsed by G. C. Stine: 1976, 'Skepticism, Relevant Alternatives and Deductive Closure', *Philosophical Studies* 29, 249–61.

Those who deny closure (e.g., Dretske) must be presupposing that the set of relevant alternatives (and thus the standards of relevance) is different for (1) than it is for (3). In particular, h is not relevant to the former, but it is relevant to the latter. Thus, while the inability of S to rule out h does not prevent him from knowing q , it does prevent him from knowing not h .

Since the set of relevant alternatives is not held fixed, the alleged failure of closure may depend on something like an equivocation. (For a discussion of this, see the paper by Stine.) The precise interpretation of the relevant alternatives approach is a controversial and complex matter. While these issues are important, they are not directly relevant to my

purpose in this paper, viz., motivating the general relevant alternatives framework as a defense against skepticism.

³³ The context sensitivity of the relevance of alternatives has been suggested by different philosophers in different ways. Goldman in 'Discrimination and Perceptual Knowledge', suggests that it may be the speaker's intentions directly concerning which alternatives are relevant that determine relevance. But as Harman points out in 'Reasoning and Evidence One Does Not Possess', the speaker himself may be ignorant of the barn replicas. Harman suggests that it may be the speaker's standards in the context that determine which alternatives in the circumstances are relevant (p. 181). I agree that this is sometimes true although it is unclear whether this suffices for a complete account because of worries about psychological reality. Do speakers, in general, intend standards?

Dretske, in 'The Pragmatic Dimension of Knowledge', says that knowledge "... exhibits a degree of contextual relativity in its ordinary use" (p. 365).

Lewis, 'Scorekeeping in a Language Game', argues that attributions of certainty and "infallible knowledge" are sensitive to conversational context.

³⁴ Notice that if an alternative is relevant on the basis of *S* possessing a defeater, it does not follow that *h* is relevant on the basis of the objective probability criterion. For example *h* (*S* sees a barn replica) may be relevant because *S* falsely believes *d* (there are numerous barn replicas in the area) on the basis of bad reasons or on the basis of misleading good reasons (e.g., mendacious testimony).

Department of Philosophy
Princeton University
Princeton, NJ 08540
U.S.A.