

Skepticism, Contextualism, and Discrimination*

JONATHAN SCHAFFER

University of Massachusetts, Amherst

The skeptic says that “knowledge” is an absolute term, whereas the contextualist says that ‘knowledge’ is a relationally absolute term. Which is the better hypothesis about “knowledge”? And what implications do these hypotheses about “knowledge” have for knowledge? I argue that the skeptic has the better hypothesis about “knowledge”, but that both hypotheses about “knowledge” have deeply anti-skeptical implications for knowledge, since both presuppose our capacity for epistemically salient discrimination.

The *skeptic* (following Peter Unger 1975) says that “knowledge” is an absolute term, requiring the elimination of all logical possibilities of falsehood. The *contextualist* (following David Lewis 1979, Fred Dretske 1981, *inter alia*) says that “knowledge” is a relationally absolute term, requiring the elimination of only the contextually relevant possibilities of falsehood. Which is the better hypothesis about the word “knowledge”? And what if any implications do these hypotheses about the word “knowledge” have for the relation of knowledge?

I argue that (i) it is the skeptic who has the better hypothesis about “knowledge”; but (ii) both the skeptical and contextualist hypotheses presuppose our capacity for epistemically salient *discrimination*, where (iii) such discrimination reveals why skeptical scenarios fail to undermine our epistemic standing. There is a deeply anti-skeptical morale buried in the skepticism-contextualism dispute.

1. Skepticism, Unger Style

The skeptic says that we know nothing, or at least far less than is usually supposed. The skeptic may target various domains and employ various arguments. I am interested in the sort of skeptic who targets knowledge about the external world, with the following argument:

* **Editor’s Note:** This paper won a Young Epistemologist Prize for the Rutgers Epistemology conference held in 2002.

- 1) One cannot know that one is not a brain-in-a-vat (or dreaming, or under the spell of Descartes's demon, or in some other skeptical scenario in which appearances deceive);
- 2) If one cannot know that one is not a brain-in-a-vat, then one cannot know that one has hands (or any other mundane proposition about the external world);
- 3) Therefore one cannot know that one has hands (or any other mundane proposition about the external world).

This argument might well be read into the concluding portion of Rene Descartes's First Meditation, so this sort of skeptic might well be called *the Cartesian skeptic*. When I speak of 'the skeptic', it is this sort of skeptic whom I have in mind.

Peter Unger (1971, 1975) is perhaps the most prominent recent skeptic of this sort.¹ According to Unger, knowledge implies *certainty*, which in turn requires the ability to eliminate every logical possibility of falsehood. So one's inability to eliminate being a (bodiless and *ipso facto* handless) brain-in-a-vat precludes one's knowing that one has hands. Unger supplements this argument with three related linguistic theses. First, Unger situates "knowledge" within a broader class of terms, the so-called *absolute terms*, including "flat", "straight", and "empty", which are said to denote a limit condition. Second, Unger offers a number of tests for whether a term is absolute, such as (i) the effect of modification by "very" ("very flat" actually seems to mean *less flat* than "flat"), and (ii) the term's behavior in comparative constructions. Third, Unger offers a *pragmatic explanation* for the assertibility of absolute terms in non-limit conditions, which is that in calling the table "flat" we convey that, for the purposes associated with the context of utterance, there is *no salient difference* between the table and a perfect plane.

Unger's pragmatic explanation for our use of "knowledge" marks an important advance for the skeptic. It enables the skeptic to explain away our linguistic intuitions about the acceptability of mundane knowledge ascriptions. That is, the skeptic is liable to the following objection (*the objection from use*): (i) skepticism entails that utterances such as "I know that I have hands" are false; (ii) if such utterances were false, then native speakers would not judge them acceptable; but (iii) native speakers do (typically enough) judge such utterances acceptable; thus (iv) skepticism is false. The Unger-

¹ Or at least Peter Unger in the 1970s. By 1984 Unger reached the 'meta-skeptical' conclusion that skepticism and contextualism are equally acceptable. By 1986 Unger turned contextualist. When I invoke Unger in the main text, it is the skeptical Unger of the 1970s whom I have in mind.

style skeptic has the resources to deny (ii), as follows. Utterances such as “I know that I have hands” are indeed false, but native speakers judge them acceptable since they are pragmatically *assertible*. If an adequate pragmatic explanation can be provided, then *Unger-style skepticism will predict all actual linguistic judgments*.

Unger’s own pragmatic explanation, nevertheless, needs updating.² Here is what Unger says:

It may be, rather, that the normal use is, in each [acceptable flatness ascription], to indicate that the object in question is close to enough to being flat, for the kind of thing it is, for the purposes which may be presumed in the case, and relative to other factors the context may provide. (1975, p. 69)

This much is fine, but more is needed. In particular, what is needed is: (i) a non-contentious and applicable precedent, and (ii) a connection to general and independently plausible pragmatic principles, such as H. P. Grice’s rules of conversational implicature.

I propose that the skeptic should treat acceptable knowledge ascriptions on the precedent of *hyperbole*. Non-contentious examples of hyperbole include “The airplane is a mile long”, and “I am dying of thirst”. What makes hyperbole an appropriate precedent is that, for the skeptic, “I know that I have hands” entails that I have eliminated every logical possibility of handlessness, which exaggerates the range of possibilities I can eliminate.³

Hyperbole may be understood in Gricean terms, as a *flouting* of the maxim of *Quality* (Grice 1967; Stephen Levinson 1983, pp. 109-10). The maxim of Quality enjoins one to speak truthfully (and within the limits of one’s evidence). Hyperbole, as with most tropes, flouts truthfulness. Given that the speaker is still assumed to be co-operative, his flouting of a maxim forces his audience to infer that he must have intended to convey something else, namely the nearest co-operative counterpart of his utterance. Thus the speaker who says, “The airplane is a mile long” will (if all goes well) be understood to mean that the airplane is large by the standards of the current context. Likewise the speaker who says, “I know that I have hands” will (if all goes well) be understood to mean that he can eliminate the possibilities of

² Linguistic pragmatics has made considerable improvement in the intervening 25 years since Unger wrote *Ignorance*. Skeptical theory has not kept pace, primarily because so few have felt motivated to uphold it.

³ Some hyperboles are more obvious than others. While the hyperbole in “The airplane is a mile long” is fairly obvious, the hyperbole in “I am dying of thirst” is less obvious, because ‘dying of thirst’ is more formulaic. Highly formulaic tropes are particularly non-obvious, since, as Kent Bach explains, “[T]he hearer’s inference to what the speaker means is short-circuited, compressed by precedent (though capable of being worked out if necessary), so that the literal content of the utterance is apparently bypassed.” (2000, p. 263; see also Bach 1975, and Lawrence Horn and Samuel Bayer 1984) Thus the fact that “I know that I have hands” is not obviously hyperbolic is no objection.

handlessness relevant in the current context. So developed, *Unger-style skepticism predicts exactly the same linguistic judgments as contextualism*.⁴

Hyperbole may also be understood in terms of Dan Sperber and Deirdre Wilson's (1986a; 1986b) theory of Relevance. Sperber and Wilson define Relevance as a measure of the optimality of the balance of effect and effort: one aims to convey the most information in the least time.⁵ Relevance provides what is perhaps the most natural explanation for such *loose talk* as hyperbole (Sperber and Wilson 1986b, §5). Saying "the airplane is a mile long" generates the entailment that the airplane is greater than x feet long (for all $x < 5280$). In some contexts it may be (i) informative to generate some of these entailments (the audience might well wonder whether the plane is greater than 100 feet long), and (ii) non-misleading to generate the others (the audience might be expected to appreciate that the plane is not greater than 1000 feet long). And it may be that in such a context, "The airplane is a mile long" is optimal with respect to balancing such effects with brevity of expression. Similarly, saying "I know that I have hands" generates the entailment that I have eliminated every logical possibility of handlessness. In some contexts it may be informative to generate some of these entailments (my audience might well wonder whether I have stumps or prostheses), and non-misleading to generate others (my audience might be expected to appreciate that I do not mean to address the brain-in-a-vat possibility). And it may be that in such a context, "I know that I have hands" is optimal with respect to balancing such effects with brevity of expression.

Why exactly is it, then, that "I know I have hands" is assertible? On the present skeptical proposal, the exact explanation is:

- 4) "Knowledge" is an absolute term, requiring the elimination of all logical possibilities of falsehood; so "I know that I have hands" entails that, for every logical possibility p in which I do not have hands, I can eliminate p ;
- 5) For some possibilities p in which I do not have hands (such as possibilities in which I have stumps or prostheses), I can eliminate p , and such is presumptively informative for my audience;

⁴ As John Hawthorne has pointed out to me, Grice's maxims are geared to the communicative exchange, whereas "knowledge" can be deployed even in silent reflection. This raises the question of whether the Gricean maxims are operative in reflection. My suspicion is that for the most part they are, but however this question is settled, it should not trouble the specific hyperbole-based pragmatic story told here, since it is evident that hyperbole is operative in reflection. Thus the timid customer may inwardly groan "This is taking forever", and the bored student may inwardly gripe "His lecturing style is lethal."

⁵ For a more formal definition of Relevance see Sperber and Wilson 1986a, pp. 381-2. Note that Sperber and Wilson's notion of Relevance is distinct from the contextualist's notion of relevance.

- 6) For the possibilities of handlessness p which I cannot eliminate (such as the possibility of being a brain-in-a-vat), my literal claim to have eliminated p is presumptively non-misleading for my audience.

If this proposal is right, then *Unger-style skepticism yields a precedented and principled pragmatic explanation for all actual linguistic judgments.*

2. Contextualism, Lewis-Dretske Style

The contextualist says that the truth-conditions for knowledge ascriptions are contextually variable. For instance, “Holmes knows that Professor Moriarty is a criminal mastermind” may be true uttered in the courtroom, but false uttered in the classroom. The contextualist is thereby committed to denying (4), since absolute terms yield context-invariant truth-conditions.

David Lewis (1979) and Fred Dretske (1981) are among the pioneer contextualists. Lewis and Dretske both say that Unger’s alleged ‘absolute terms’ are really *relationally absolute terms*, for which the limit condition is contextually variable. Here is how Lewis puts the point:

The right response to Unger, I suggest, is that he is changing the score on you. When he says that the desk is flatter than the pavement, what he says is acceptable only under raised standards of precision. Under the original standards the bumps on the pavement were too small to be relevant either to the question whether the pavement is flat or to the question whether the pavement is flatter than the desk. Since what he says requires raised standards, the standards accordingly rise. Then it is no longer true enough that the pavement is flat. That does not alter the fact that it *was* true enough *in its original context.* (1979, pp. 245-6)

And here is how Dretske puts much the same point: “Such concepts [as knowledge, flatness, and emptiness] are *relationally absolute*: absolute, yes, but only relative to a certain standard. We might put the point this way: to be empty is to be *devoid of all relevant things,...*” (1981, p. 367) And so Lewis and Dretske both conclude that “knowledge” implies the elimination, not of every logical possibility of falsehood, but only of *the contextually relevant* possibilities of falsehood (Dretske 1981, p. 367; Lewis 1996, p. 425).

Subsequent contextualists have preserved the idea of contextually variable standards for “knowledge”, while departing from the idea (shared by Unger, Lewis, and Dretske) that the standards concern which possibilities must be eliminated. Thus for Stewart Cohen (1988) the standards concern the degree of evidence needed for justification, while for Keith DeRose (1995) the standards concern the distance in logical space through which one must track truth. These contextualists share the following core ideas about “knowledge”:

- 7) Skeptical scenarios (such as being a brain-in-a-vat) are irrelevant in typical contexts, and in those contexts the standards for “knowledge” are low enough that “I know that I have hands” is true;
- 8) Skeptical scenarios become relevant in atypical contexts in which they are considered, and in those contexts the standards for “knowledge” are high enough that “I know that I have hands” is false.

For the sake of definiteness, and to preserve the parallel to Unger-style skepticism, I treat the standards as concerning which possibilities must be eliminated:

- 9) “Knowledge” is a relationally absolute term, requiring the elimination of only the contextually relevant possibilities of falsehood.

When I speak of ‘the contextualist’, it is the advocate of (7)-(9) whom I have in mind. Though what I say is meant to be general enough to cover other conceptions of the standards.⁶

The contextualist claims to avoid skepticism. *Yet contextualism does not entail that we know anything.*⁷ Contextualism does entail that “G. E. Moore knows that he has hands” is true in the typical context, but this is a meta-linguistic claim about the truth of an utterance in a context. In order to descend from this meta-linguistic claim about truth to an object-level claim about knowledge one needs *disquotation for truth*: “*p*” is true iff *p*. And context-variability precludes disquotation.⁸ So even though the contextualist can say that “Moore knows that he has hands” is true in some contexts, this cannot then be disquoted to establish that Moore knows that he has hands (Ernest Sosa 2000).

⁶ In order to extend what I say about the skepticism-contextualism dispute to other contextualist conceptions of the standards for “knowledge”, one must first recover the skeptical counter-proposal for that standard. For instance, there is a skeptical counter-proposal to Cohen’s view on which “knowledge” invariantly requires absolutely conclusive evidence, and a skeptical counter-proposal to DeRose’s view on which “knowledge” invariantly requires tracking truth through the entirety of logical space. The terms of the current debate between the Unger-style skeptic and the Lewis-Dretske-style contextualist may then be transposed onto a debate between the skeptical and contextualist proposals for that standard.

⁷ The reader should beware that contextualists often adopt the expository convention of conducting their discussion in terms of knowledge rather than “knowledge”, with just a brief note to mention this slide.

⁸ Context-variability precludes disquotation because it allows for the possibility that “*p*” is true in its context of utterance, while *p* is false in its context of disquotation. For example, suppose I love you unrequitedly, and tell you truly “I love you”. If you then disquote my true utterance, you will wind up assigning it the truth-condition of you loving me, which is (by sad supposition) false.

Contextualism may still help to defuse the skeptical argument of (1)-(3), albeit in a very partial and indirect way, as a complement to an antecedently non-skeptical epistemic pluralism. The sort of epistemic pluralism I have in mind is as follows:

- 10) There exists an at-least-two-membered set K of knowledge relations;
- 11) K can be partitioned into a non-skeptical cell, $[k_m]$, in which I know_m that I have hands, and a skeptical cell, $[k_s]$, in which I do not know_s that I am not a brain-in-a-vat;
- 12) Each member of K is closed under known implication (so since I know_m that I have hands, I know_m that I am not a brain-in-a-vat; and since I do not know_s that I am not a brain-in-a-vat, I do not know_s that I have hands).⁹

The bridge between the contextualism about “knowledge” of (7)-(9) and the epistemic pluralism of (10)-(12) is *reference shifting*:

- 13) “Knowledge” shifts reference over the members of K , and while (i) “knowledge” typically refers to some member of the non-skeptical cell $[k_m]$, (ii) consideration of skeptical scenarios makes “knowledge” refer to some member of the skeptical cell $[k_s]$.

Putting (7)-(13) together, knowledge ascriptions are context-variable because “knowledge” shifts its reference over a plurality of knowledge relations.

The epistemic pluralism of (10)-(12) provides the following direct reply to the skeptical argument of (1)-(3): (1) and (3) are both false for at least one knowledge relation (knowledge_m). Yet this reply does not even *mention* the context-variability of “knowledge”! *It is epistemic pluralism that is providing the direct reply to skepticism*, by directly postulating the existence of a non-skeptical knowledge_m relation.¹⁰

The only anti-skeptical work left for the contextualist theses (7)-(9) is the following. *Contextualism about “knowledge” (7)-(9), via reference shifting*

⁹ Dretske (1970) takes contextualism to license a denial of closure. The point that each member of K should still satisfy closure is due to G. C. Stine (1976). It has been accepted by virtually all subsequent contextualists (see especially Cohen 1988, §3, and DeRose 1995, §10).

¹⁰ DeRose (1995) speaks of ‘solving the skeptical argument’ by combining a tracking view of knowledge (as per Robert Nozick 1981) with contextualism about “knowledge”. DeRose’s epistemic pluralism concerns the various distances in logical space through which one might be required to track. His non-skeptical knowledge relations are those that only require relatively near-in tracking. Near-in tracking is thus what is providing the direct reply to skepticism. Contextualism about “knowledge” plays a far less direct role in DeRose’s solution than might be thought.

(13), can be invoked in an attempt to explain away three implausible consequences of epistemic pluralism (10)-(12). One implausible consequence of epistemic pluralism is that I get to know, in a sense (knowledge_m), that I am not a brain-in-a-vat. Contextualism can be invoked in an attempt to explain away this implausible consequence as true but *ineffable*. By (13), uttering “I know that I am not a brain-in-a-vat” makes “knowledge” refer to knowledge_s, when by (11) I do not know_s that I am not a brain-in-a-vat (DeRose 1995, p. 40).

A second implausible consequence of pluralism is that I do not get to know, in a sense (knowledge_s), that I have hands. Contextualism can be invoked in an attempt to explain away this implausible consequence as true but *falsity-associated*. By (13), the form of words “I do not know that I have hands” typically involves reference to knowledge_m, when by (11) I do know_m that I have hands. Thus we associate the form of words with falsity (DeRose 1995, pp. 40-1).

A third implausible consequence of pluralism is that there is no conflict between our ordinary knowledge claims as expressed by “I know that I have hands”, and skeptical doubts as expressed by “I do not know that I have hands”. Contextualism can be invoked in an attempt to explain away this implausible conjunction as true but *superficially contradictory*. The explanation is that we are prone to be blind to the subscripts which, by (13), would resolve the apparent contradiction (DeRose 1995, p. 5; Cohen 2001, p. 89).¹¹

In summary, the real anti-skeptical force is the epistemic pluralism of (10)-(12). Contextualism is merely serving as a fig leaf to cover up certain resulting implausibilities.

3. Skepticism Redux, Or: Why the Skeptic Has the Better Linguistic Hypothesis

Which is the better hypothesis about “knowledge”, that of the skeptic or that of the contextualist?¹² It is generally thought (even by those who dismiss

¹¹ All three of these explanatory attempts require further development. Ineffability, falsity-association, and superficial contradiction do not generally confuse competent speakers. For instance, (i) “I am not speaking now” is ineffable, but I am not thereby inclined to suppose that I must be constantly talking; (ii) “I am G. E. Moore” is falsity-associated, but I am not thereby inclined to deny it if it happens to be Moore who says it; and (iii) the conjunction of (a) “I am G. E. Moore” said truly by Moore, and (b) “I am not G. E. Moore” said truly by me, forms a superficial contradiction, but I am not thereby inclined to announce a paradox (Stephen Schiffer 1996). If the contextualist is to sustain any of these three explanations, she will need to say what distinguishes the ‘straightforward’ cases above from the ‘confusing’ cases, and say why “knowledge” cases are among the latter.

¹² There is also room for a ‘Moorean’ semantics for “knowledge” that is invariant but not absolutist. More precisely, hypotheses about “knowledge” can be divided into (i) variantist (/contextualist) versus (ii) invariantist views; the invariantist views may then be

contextualism as little help against skepticism, such as Sosa (2000), Hilary Kornblith (2000), and Richard Feldman (2001)) that the contextualist has the better semantic hypothesis. I maintain that *the skeptic has the better hypothesis*.

Two points of standard linguistic methodology are worth mentioning here, by way of prelude to the argument. The first point concerns the evidential status of linguistic intuitions:

- 14) Our linguistic intuitions provide evidence for acceptability, and do not discriminate between semantic and pragmatic sources.

Noam Chomsky is emphatic on this point:

We may make an intuitive judgment that some linguistic expression is odd or deviant. But we cannot in general know, pretheoretically, whether this deviance is a matter of syntax, semantics, pragmatics, belief, memory limitations, style, etc., or even whether these are appropriate categories for the interpretation of the judgment in question. It is an obvious and uncontroversial fact that informant judgments and other data do not fall neatly into clear categories: syntactic, semantic, etc. (1977, p. 4; *inter alia*)

The semantic and pragmatic levels, after all, are the sophisticated posits of a scientific theory of language, to which naïve intuitions cannot be expected to be sensitive. By (14) our linguistic intuitions about the acceptability of ordinary knowledge ascriptions cannot decide between the contextualist's semantic explanation and the skeptic's pragmatic explanation.¹³

The second point of standard linguistic methodology concerns the preference for theoretical simplicity:

subdivided into (iia) absolutist (/skeptical) versus (iib) nonabsolutist (/Moorean) views. In the language of epistemic pluralism, the contextualist allows "knowledge" to shift over knowledge_m and knowledge_s, the skeptic takes "knowledge" to refer invariantly to knowledge_s, while the Moorean takes "knowledge" to refer invariantly to knowledge_m. See note 19 for a start on the comparison between skeptical and Moorean semantics.

¹³ There are tests for semantic entailment versus conversational implicature, but these tests turn out unhelpful here. The best test is cancellability (Grice 1967). If Jones says, "I need gas", and Smith replies, "There is a station around the corner", then Smith's reply implicates (by Grice's maxim of Relevance) but does not entail that the station is open and sells gas. That this is an implicature but not an entailment can be seen by Smith's opportunity to cancel without contradiction ("—but it closed years ago", or "—but it doesn't sell gas; in fact, it's a police station"). The main problem with the cancellability test is that it fails to distinguish invariance-plus-implicature from context-variability, because the alleged cancellatory clause may instead be claimed to function as a context shifter (Jerrold Sadock 1978). This is exactly what happens here. To apply cancellability, one should consider: "I know that I have hands—and I do mean that I can even eliminate being a brain-in-a-vat". For the skeptic this ought to cancel the hyperbolic implicature that less was meant, and so my claim to knowledge ought to seem unacceptable, as indeed it seems. So cancellability might appear to favor the skeptic. But really the contextualist makes just the same prediction, since the contextualist will take the mention of the brain-in-a-vat possibility to force a context shift in just such a way as to now render the (previously true!) "I know that I have hands" false.

15) Linguistic machinery should not be complicated without necessity.

The application of (15) is worth illustrating with two cases. The first case is that of someone who claims that “and” is *ambiguous* between (i) &: mere conjunction and (ii) and-then: conjunction plus temporal order.¹⁴ The evidence for the alleged ambiguity is that “London is the capital of England and Paris is the capital of France” is *reversible* (as mere conjunction is) in that one could acceptably reverse the order of the conjuncts, while “The Lone Ranger jumped on his horse and rode into the sunset” is irreversible (as temporal order is). Yet linguists nearly universally reject this claim of ambiguity, primarily because (i) the postulation of a second sense of “and” complicates the semantics, while (ii) the unacceptability of “The Lone Ranger rode into the sunset and jumped on his horse” can already be explained away pragmatically, by Grice’s maxim of Manner, which enjoins one to speak orderly. Thus the simplest overall linguistic hypothesis is that “and” *univocally* means mere conjunction, but sometimes *implicates* temporal priority. This hypothesis simplifies the semantics without complicating the pragmatics.

The second case is that of someone who claims that “bachelor” means *man* (married or not) and then attempts to explain away the unacceptability of “Jones is a married bachelor” on pragmatic grounds.¹⁵ This “bachelor” hypothesis should be rejected because (i) it does not simplify the semantics, and (ii) it considerably complicates the pragmatics, since there are no appropriate precedents for, and no general and independently plausible pragmatic principles that would explain, the non-assertibility of “Jones is a married bachelor”. Thus the simplest overall linguistic hypothesis is that “bachelor” means unmarried man.

With methodological principles (14) and (15) in hand, I claim that the skeptic has the better hypothesis about “knowledge”. The reason is akin to that of the classic Gricean argument for the univocity of “and”:

16) The skeptic has the simpler semantics;

17) The skeptic does not complicate the pragmatics.

Starting with (16), the reason why the skeptic has the simpler semantics is that the contextualist needs to posit extra semantic rules of relevance.¹⁶ Lewis

¹⁴ This case is drawn from Grice 1981 (p. 186). Grice cites P. F. Strawson as a proponent of “and” ambiguity. The reasoning involved here is sometimes referred to as Grice’s Razor (Wayne Davis 1998).

¹⁵ This case is drawn from DeRose 1999 (p. 197). DeRose imagines this ‘crazed philosopher of language’ to appeal to the following alleged pragmatic principle: “There’s a conversational rule to the effect that you shouldn’t assert ‘S is a bachelor’ where S is married.” (1999, p. 199)

¹⁶ Does the contextualist multiply senses? Yes and no: to adopt the terminology of David Kaplan (1989), contextualism multiplies contents, not characters. The point where the

(1996), who offers by far the most sophisticated contextualist semantics to date, postulates seven semantic rules of relevance. On Lewis's account, the following possibilities are always relevant: (i) actuality, (ii) those the subject believes or ought to believe to be actual, (iii) those the attributor is attending to, and (iv) those that directly and saliently resemble those relevant by (i)-(iii); whereas the following possibilities are defeasibly irrelevant: (v) those concerning errors in reliable processes, (vi) those concerning errors in sampling and abduction, and (vii) those conventionally ignored. Lewis's rules are not derivable from general and independently plausible semantic rules. To recognize such rules is to complicate the semantics.¹⁷

Turning to (17), the reason why the skeptic does not complicate the pragmatics is that the pragmatic machinery needed to handle hyperbole is already in place. By (4)-(6), the skeptic explains how knowledge ascriptions are interpreted hyperbolically. Since both the skeptic and the contextualist should already agree that the pragmatic machinery needed to handle hyperbole exists, assimilating knowledge ascriptions to hyperbole does not complicate the pragmatics.

Indeed, *the assertibility conditions on hyperbole already look to generate pragmatic counterparts of Lewis's rules*. The assertibility conditions on hyperbole (5)-(6) depend on which entailments one's audience will presumptively take seriously (treat literally rather than dismiss as merely figurative). So the assertibility of "I know that I have hands" depends, *inter alia*, on whether the audience will presumptively take seriously the entailment that one knows that one is not a brain-in-a-vat. If so, then the utterance will be unacceptably misleading. If not, not. Lewis himself characterizes his rules as concerning which possibilities are *properly ignored*, and so it seems that pragmatic counterparts of these rules should be derivable from which entailments are *pragmatically presumed* to be ignorable. Here it only needs to be pointed out that one's audience will presumptively take seriously the actual, believed, and attended possibilities, as well as those that saliently resemble them; and that one's audience will presumptively ignore the reliable-process-breakdown, methodological-error, and conventionally ignored possibilities.¹⁸

contextualist multiplies semantic machinery is in the function by which character plus context determines content.

¹⁷ Of course there are some contextually variant semantic rules, such as govern "I". This does not rebut the point of (16), since the contextualist's rules for "knowledge" do not derive from the rules for "I" or any other independently plausible semantic rules. And this fits the methodology of (15), since the context-variability induced by "I" cannot be explained away on antecedently established pragmatic grounds—if "I" invariably referred to John Perry, no Gricean rule could explain away our use of it.

¹⁸ Since Lewis understands his rules in terms of presupposition, and since presupposition is now generally regarded as pragmatic (see Levinson 1983, ch. 4, and esp. pp. 199- 204; see also Robert Stalnaker 1974, 1999), Lewis himself might best be read as a skeptic. Indeed, Lewis 1979 begins by focusing on acceptability as related to presupposition, carefully disclaiming any commitment to whether such acceptability is semantically or

In summary, both the skeptic and the contextualist need a theory of *relevant alternatives*. The difference is that the skeptic relegates relevance to the pragmatics, and there finds antecedently established principles that generate the needed implicatures; whereas the contextualist imports relevance into the semantics, and is thereby forced to postulate extra semantic rules to capture the needed entailments. So the skeptic's hypothesis is preferable on grounds of overall theoretical simplicity and elegance.¹⁹

The primary objection that contextualists level at skepticism is that it is *skepticism*, and involves competent speakers in "systematic and widespread falsehood" (DeRose 1995, p. 44; also 1999, p. 202; see also Stine 1976, p. 254; Cohen 1999, p. 83).²⁰ But linguistic questions are empirical questions, and empirical linguistics is not beholden to philosophical prejudices or charitable inclinations. Rather, as (14) and (15) suggest, *empirical linguistics is an attempt to provide a systematic theory of language, and pragmatic explanations are every bit as good as semantic explanations in that regard*. Speakers of a language need not mean what philosophers hope them to. To prefer contextualism in the hopes of rebutting skepticism is merely wishful thinking.

pragmatically based (p. 234: this accords with my (14)). Yet when Lewis comes to discussing "knowledge" (p. 246), he suddenly focuses only on truth. All he should have focused on, given his careful beginnings, is acceptability; then Lewis and the skeptic would be in agreement.

¹⁹ As to the skeptic versus the Moorean (as per note 12), (i) both offer equally simple invariant semantics, but (ii) the Moorean seems to complicate the pragmatics. While the skeptic needs to explain away the acceptability of "I know that I have hands" in the typical context, the Moorean needs to explain away the unacceptability of "I know that I am not a brain-in-a-vat" in any context. While the skeptic seems able to do what she needs via general and independently plausible pragmatic rules (on the precedent of hyperbole), I cannot see how the Moorean can do what she needs without inventing new pragmatic rules—at least, none of Grice's rules look to help her here, and I see no plausible precedents. Though Rich Feldman has suggested to me that the Moorean may invoke the precedent of understatement. Alternatively, the Moorean can run some other sort of explanation. In this vein, Sosa (1999, pp. 147-8) offers a Moorean explanation via performance error: we are prone to confuse the genuine and non-skeptical safety requirement (Bp>p), with the spurious and skeptical sensitivity requirement (~p>~Bp), because we wrongly suppose that counterfactuals contrapose. (Assuming that causal reasoning involves counterfactuals, are we prone to any sort of comparable error there?) I must leave discussion of such issues, and the prospects for Mooreanism generally, to another place.

²⁰ DeRose (1995, §§ 15-16; 1999, §§ 8-11) rebuts a skeptic who simply relabels the contextualist's new semantic rules as new pragmatic rules. I agree with DeRose that that sort of skeptic does not have the better theory (she merely trades semantic complexity, with no overall theoretical gain). My sort of skeptic claims not to need any new pragmatic rules, and only to need the thoroughly general and independently plausible principles at work in hyperbole.

4. Discrimination, Or: The Anti-Skeptical Epistemic Morale

What implications do the skeptical and contextualist hypotheses about “knowledge” have for knowledge? It is generally thought that Unger-style skeptical invariantism is disastrously skeptical, and questioned whether Lewis-Dretske-style contextualism is much better. I maintain that *both the skeptic and the contextualist are committed to a deeply anti-skeptical view of our discriminatory capacities.*

To illustrate the notion of discrimination I have in mind,²¹ suppose (to extend a case from Dretske 1970) that Student, Assistant, and Professor are eyeing a zebra. Each attempts to form beliefs on the basis of their experience. Student is remarkably ignorant, and cannot even eliminate the possibility that he is seeing a normal mule. Assistant is halfway educated, and can eliminate the possibility that he is seeing a normal mule, though he cannot eliminate the possibility that he is seeing a cleverly painted mule. Professor is learned, and can even eliminate the possibility that he is seeing a cleverly painted mule, on the basis of anatomical features that no mere paint job could disguise. In such a case, I say that Student cannot discriminate a zebra from either a normal or a painted mule, Assistant can discriminate a zebra from a normal mule but not from a painted mule, and Professor can discriminate a zebra from either a normal or a painted mule. (And none can discriminate a zebra from a vat-image of a zebra.)

More precisely, s has the capacity to discriminate proposition p from proposition q iff (i) p and q are exclusive, (ii) s 's experience is compatible with p , and (iii) s 's experience is incompatible with q .²² And s discriminates p from q (exercises the capacity) iff s believes that p rather than q on the basis of s 's capacity to discriminate p from q .²³ The totality of s 's discriminations relative to p , or *discriminatory range R for p* , is the set $R = \{q/s \text{ discriminates } p \text{ from } q\}$. Let p be the proposition that the beast is a zebra, let $q1$ be the proposition that the beast is a normal mule, and let $q2$ be the proposi-

²¹ This notion of discrimination is from Alvin Goldman: “My emphasis on discrimination accords with a sense of the verb ‘know’ that has been neglected by philosophers. The O. E. D. lists one (early) sense of ‘know’ as ‘to distinguish (one thing) from (another),’ as in ‘I know a hawk from a handsaw’ (Hamlet)... I suggest that a person is said to know that p just in case he distinguishes or discriminates the truth of p from relevant alternatives.” (1976, p. 772) See Schaffer (manuscript) for an argument that much of our actual knowledge talk is at least implicitly of the form ‘ s knows that p rather than q ’, and is keyed to this discriminatory sense of “know”.

²² Here I am drawing on Lewis’s (1996) definition of elimination. Strictly speaking all this should be relativized to a time t . In the main text I suppress t for convenience.

²³ Thanks to Bill Brewer and Fred Dretske for questions which inspired modifications here. The requirement of based belief is also intended to overcome an objection I have leveled elsewhere against the classical version of the relevant alternatives theory (–the problem of the missed clue in Schaffer 2001).

tion that the beast is a painted mule.²⁴ Then Student's discriminatory range **R1** over p does not contain $q1$ or $q2$, Assistant's discriminatory range **R2** over p contains $q1$ but not $q2$, and Professor's discriminatory range **R3** over p contains both $q1$ and $q2$. (And none contains $q3$, that the apparent beast is really a vat-image.)

Since discrimination requires an experiential and a doxastic aspect in the basing relation, there are three ways that one can fail to discriminate p from q : (i) one can lack the requisite experience (blindness), (ii) one can lack the requisite beliefs (ignorance), or (iii) one can have the requisite experience and the beliefs but fail the basing relation (luck). Student, Assistant, and Professor all have the requisite experience to discriminate p from $q1$ and $q2$ but not from $q3$. But Student, Assistant, and Professor differ in their relative abilities to base beliefs on their experience: Student can extract comparatively little doxastic juice from his experience, Assistant can extract a bit more, while Professor can squeeze the most out of his experience.

Both the skeptic and the contextualist are committed to viewing our epistemic standing as grounded in our discriminatory range. Taking the contextualist first, by (7)-(9) the contextualist explicitly takes the epistemic standards to concern which possibilities must be eliminated. One's epistemic standing for a given proposition p then concerns which epistemic standards one meets, which is determined by the extent of one's discriminatory range **R** over p . If the propositions in **R** pertain to all the possibilities required under a given epistemic standard e , then s meets e ; if not, not. Since Student's discriminatory range **R1** over the proposition that the beast is a zebra is a proper subset of Assistant's range **R2**, which in turn is a proper subset of Professor's range **R3**, it follows that Student has a lesser epistemic standing than Assistant, who in turn has a lesser epistemic standing than Professor. (Though none have perfect epistemic standing.)

Turning to the skeptic, by (4)-(6) the skeptic implicitly takes the assertibility of a knowledge ascription to depend on which possibilities must be eliminated. One's epistemic standing for a given proposition p then concerns how close to absolute knowledge one comes, which is determined by the extent of one's discriminatory range **R** over p . If the propositions in **R** pertain to all the possibilities required under a given assertibility standard e , then s meets e ; if not, not. Thus the contexts in which "Student knows that the beast is a zebra" is assertible will be a proper subset of the contexts in which "Assistant knows that the beast is a zebra" is assertible, which in turn will be

²⁴ As Jonathan Weinberg and Timothy Williamson pointed out to me, p and the q s should, strictly speaking, be taken as shorthand for more complex propositions with an added conjunct to the effect that perception is functioning properly (s is not dreaming or hallucinating or ...). In the main text I suppress this complication.

a proper subset of the contexts in which “Professor knows that the beast is a zebra” will be assertible. (Though none is assertible in all contexts.)

It should not be surprising that the skeptic and the contextualist share a commitment to the same underlying picture of our epistemic standing. Consider the absolutist and relational absolutist about “flat”. Both agree on the underlying physical and geometric properties of the surface (degree of curvature, irregularities, etc.); both agree that such properties determine the use of flatness ascriptions; they only disagree as to how. Likewise both the skeptic and the contextualist agree on the underlying discriminatory capacities of subjects; both agree that such properties determine the use of knowledge ascriptions; they only disagree as to how (Kornblith 2000).

Putting this together, both the skeptic and the contextualist are committed to:

18) We have the capacity to discriminate; and

19) One’s epistemic standing is based on one’s discriminatory range.

The only disagreement is a verbal disagreement, concerning whether one’s epistemic standing is ever sufficiently strong to satisfy “knowledge”.

I maintain that (18) and (19) yield a deeply anti-skeptical insight. The insight is that epistemic standing based on discrimination cannot be undermined by skeptical scenarios. Skeptical scenarios concern possibilities that one cannot eliminate. They reveal the limits of our discriminatory range. But the existence of possibilities outside one’s discriminatory range does not imply the absence of any possibilities inside that range. Neither Student nor Assistant nor Professor can eliminate the possibility that the apparent zebra is merely a vat-image. Yet each retains the epistemic standing associated with the various possibilities they can eliminate. Thus the epistemic distinctions one needs to draw between Student, Assistant, and Professor are untouched by brain-in-a-vat scenarios. This is why skeptical doubts do not collapse our epistemic standing.²⁵

The epistemic standing provided by discrimination is, moreover, the right sort of thing to care about. Our ultimate epistemic interest is *truth*. One’s discriminatory range represents *one’s progress towards truth*, by representing the extent of the possibilities one can eliminate. (Student is not so far along, Assistant has made some progress, and Professor more still.)

²⁵ Here I am drawing on Mark Heller: “Even though neither my wife nor I can rule out the possibility of an evil genius deceiving us about where the leftovers are, she is in a better epistemic position than I am.” (1999, p. 119). Heller, however, objects that skepticism collapses the epistemic distinction between he and his wife. I am saying that skepticism (in its pragmatic explanation of the use of knowledge ascriptions) actually presupposes the existence of such distinctions, and only prevents them from being truly described in terms of “knowledge”.

Think of the epistemic project as a *game of inquiry*. Each question presents a slate of alternatives. One who has the capacity to discriminate has *the power to answer a question*.²⁶

Picture the epistemic project as a *homing project* in logical space. One aims to locate the actual world among the other possibilities. One who has the capacity to discriminate has *the ability to navigate a region of logical space*.²⁷

Putting these last points together, I think that there is the following deeply anti-skeptical morale buried in the skepticism-contextualism dispute:

20) Skeptical scenarios do not collapse one's discriminatory range; and

21) Discrimination is the critical epistemic property.

What then of the 'skeptical argument' (1)-(3)? I say: one may as well grant the whole argument, and the resulting failure of knowledge. Such a failure of knowledge would at worst yield a *shallow skepticism*, due to a conceptual defect (as it were) with "knowledge". By (20) and (21) such failure would not impact what we should care about. Discrimination would remain to support a *deep anti-skepticism*.

Do I know that I have hands? I say: I can discriminate hands from stumps, though I cannot discriminate hands from vat-images of hands. That is what matters. *So do I know that I have hands?* I say the answer might as well be: *no*.²⁸

REFERENCES

- Bach, Kent 1975. "Performatives are Statements Too", in *Philosophical Studies* v. 28, pp. 229-236.
- 2000. "Quantification, Qualification, and Context: A Reply to Stanley and Szabó", in *Mind & Language* v. 15, pp. 262-83.
- Castañeda, Hector-Neri 1980. "The Theory of Questions, Epistemic Powers, and the Indexical Theory of Knowledge", in *Midwest Studies in*

²⁶ Hector-Neri Castañeda (1980) maintains that, "knowledge involves essentially the non-doxastic component of a power to answer a question." (p. 194) Castañeda attributes this idea to Plato.

²⁷ Here it is worth recalling Lewis's insightful answer to "What is it all for?": "The serious business of science has to do not with knowledge per se; but rather, with the elimination of possibilities through the evidence of perception, memory, etc., and with the changes that one's belief system would (or might or should) undergo under the impact of such eliminations. Ascriptions of knowledge to yourself or others are a very sloppy way of conveying very incomplete information about the elimination of possibilities." (1996, p. 440)

²⁸ Thanks to the audience of the Rutgers Epistemology Conference for very useful feedback.

- Philosophy V*, eds. P. French, T. Uehling, Jr., and H. Wettstein, pp. 193-238. Minneapolis: University of Minnesota Press.
- Chomsky, Noam 1977. *Essays on Forms and Interpretation*. Amsterdam: North Holland.
- Cohen, Stewart 1988. "How to be a Fallibilist", in *Philosophical Perspectives* v. 2, pp. 91-123.
- 1999. "Contextualism, Skepticism, and the Structure of Reasons", in *Philosophical Perspectives* v. 13, pp. 57-89.
- 2001. "Contextualism Defended: Comments on Richard Feldman's 'Skeptical Problems, Contextualist Solutions'", in *Philosophical Studies* v. 103, pp. 87-98.
- Davis, Wayne 1998. *Implicature: Intention, Convention, and Principle in the Failure of Gricean Theory*. Cambridge: Cambridge University Press.
- DeRose, Keith 1995. "Solving the Skeptical Problem", in *Philosophical Review* v. 104, pp. 1-52.
- 1999. "Contextualism: An Explanation and Defense", in *The Blackwell Guide to Epistemology*, eds. J. Greco and E. Sosa, pp. 187-205. Oxford: Basil Blackwell.
- Dretske, Fred 1970. "Epistemic Operators", in *The Journal of Philosophy* v. 67, pp. 1007-23.
- 1981. "The Pragmatic Dimension of Knowledge", in *Philosophical Studies* v. 40, pp. 363-78.
- Feldman, Richard 2001. "Skeptical Problems, Contextualist Solutions", in *Philosophical Studies* v. 103, pp. 61-85.
- Goldman, Alvin 1976. "Discrimination and Perceptual Knowledge", in *Journal of Philosophy* v. 73, pp. 771-91.
- Grice, H. P. 1967. *Logic and Conversation*. William James Lectures, Harvard University.
- 1981. "Presupposition and Conversational Implicature", in *Radical Pragmatics*, ed. P. Cole, pp. 183-98. New York: Academic Press.
- Heller, Mark 1999. "Contextualism and Anti-Luck Epistemology", in *Philosophical Perspectives* v. 13, pp. 115-29.
- Horn, Lawrence and Bayer, Samuel 1984. "Short-Circuited Implicature: A Negative Contribution", in *Linguistics and Philosophy* v. 7, pp. 397-414.
- Kaplan, David 1989. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals", in *Themes from Kaplan*, eds. J. Almog, J. Perry, and H. Wettstein, pp. 481-563. New York: Oxford University Press.
- Kornblith, Hilary 2000. "The Contextualist Evasion of Epistemology", in *Philosophical Issues* v. 10, pp. 24-32.

- Levinson, Stephen 1983. *Pragmatics*. Cambridge: Cambridge University Press.
- Lewis, David 1979. "Scorekeeping in a Language Game", in *Journal of Philosophical Logic* v. 8, pp. 339-59.
- 1996. "Elusive Knowledge", in *Australasian Journal of Philosophy* v. 74, pp. 549-67.
- Nozick, Robert 1981. *Philosophical Explanations*. Cambridge: Harvard University Press.
- Sadock, Jerrold 1978. "On Testing for Conversational Implicature", in *Pragmatics: A Reader* (1991), ed. S. Davis, pp. 365-76. New York: Oxford University Press.
- Schaffer, Jonathan 2001. "Knowledge, Relevant Alternatives, and Missed Clues", in *Analysis* v. 61, pp. 202-8.
- manuscript. "Contrastive Knowledge".
- Schiffer, Stephen 1996. "Contextualist Solutions to Scepticism", in *Proceedings of the Aristotelian Society* v. 96, pp. 317-33.
- Sosa, Ernest 1999. "How to Defeat Opposition to Moore", in *Philosophical Perspectives* v. 13, pp. 141-53.
- 2000. "Skepticism and Contextualism", in *Philosophical Issues* v. 10, pp. 1-18.
- Sperber, Dan and Wilson, Deirdre 1986a. *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.
- 1986b. "Loose Talk", *Proceedings of the Aristotelian Society* v. 86, pp. 153-71.
- Stalnaker, Robert 1974. "Pragmatic Presuppositions", in *Context and Content: Essays on Intentionality in Speech and Thought* (1999), pp. 47-62. Oxford: Oxford University Press.
- 1999. "Introduction", in *Context and Content: Essays on Intentionality in Speech and Thought*, pp. 1-28. Oxford: Oxford University Press.
- Stine, G. C. 1976. "Skepticism, Relevant Alternatives, and Deductive Closure", in *Philosophical Studies* v. 29, pp. 249-61.
- Unger, Peter 1971. "A Defense of Skepticism", in *Philosophical Review* 80, pp. 198-218.
- 1975. *Ignorance: A Case for Scepticism*. Oxford: Clarendon Press.
- 1984. *Philosophical Relativity*. Minneapolis: University of Minnesota Press.
- 1986. "The Cone Model of Knowledge", in *Philosophical Topics* v. 14, pp. 125-78.