

Applying Pragmatics to Epistemology

KENT BACH

San Francisco State University

kbach@sfsu.edu

<http://online.sfsu.edu/~kbach/>

ABSTRACT

This paper offers a smattering of applications of pragmatics to epistemology. In most cases they concern recent epistemological claims that depend for their plausibility on mistaking something pragmatic for something semantic. After giving my formulation of the semantic/pragmatic distinction and explaining how seemingly semantic intuitions can be responsive to pragmatic factors, I take up the following topics:

1. Classic Examples of Confusing Meaning and Use
2. Pragmatic Implications of Hedging or Intensifying an Assertion
3. Belief Attributions
4. Knowledge-*wh*
5. The Knowledge Rule on Assertion
6. Testimony
7. Asserting and Thinking of Possibilities
8. Concessive Knowledge Attributions
9. "Pragmatic Encroachment"
10. Epistemic Contextualism

APPLYING PRAGMATICS TO EPISTEMOLOGY¹

KENT BACH

San Francisco State University

A number of recent and not so recent claims in epistemology depend for their plausibility on mistaking something pragmatic for something semantic. Or so I will suggest. After making several general observations about what makes the pragmatic pragmatic, I will survey these claims, some of which have become unduly influential, and explain just what sort of semantic/pragmatic confusion they rely on. This survey will also include a few other epistemological applications of pragmatics.

0. Introduction: Pragmatics and Semantic Intuitions

If speakers always meant precisely what they said, as determined by the meanings of the sentences they uttered, then, semantic ambiguity aside, making one's communicative intention evident and recognizing someone else's would always be perfectly straightforward. In fact, speakers generally do not mean precisely what they say. In leaving part of what they mean implicit or in meaning something in addition to what they say, they leave much to inference. However, there is a mutual understanding, collateral to any normal utterance, that the speaker is saying what he says with a certain identifiable communicative intention. By exploiting this Communicative Presumption, the listener can perform the job of figuring out, given that the speaker said what he said, what the speaker could have meant in saying it. In figuring this out, he takes into account mutually evident contextual information as well as the fact that he is intended to figure it out.

These elementary observations about the relationship between what a speaker says, what he means in saying it, and how he manages to be understood suggest a certain

¹ This is the text of a talk given at the Epistemology and Linguistics Conference, Aberdeen, Scotland, May 12-13, 2007. It will appear in *Philosophical Issues* 18.

conception of what makes the pragmatic pragmatic: the very fact that a sentence is uttered gives rise to distinctively pragmatic facts. As Paul Grice (1961) observed, it is the fact that a speaker utters a sentence with a certain semantic content (or even that sentence rather than another with the same content) that generates what Grice would later call a conversational implicature. His first example of this was an utterance of ‘Jones has beautiful handwriting and his English is grammatical’, made as an evaluation of Jones’s philosophical ability, to implicate that Jones is no good at philosophy. A different sort of illustration is provided by Moore’s paradox (so-called). If you say, “Snow is white, but I don’t believe it,” you are denying that you believe something you have just asserted. The contradiction here is pragmatic. That snow is white does not entail your believing it (nor vice versa), and there’s no contradiction in *my* saying, “Snow is white, but you don’t believe it.” The inconsistency arises not from what you are asserting but from the fact that you are asserting it. That’s what makes it a *pragmatic* contradiction.

In general, when a speaker assertively utters a sentence and, let’s assume, uses it literally, the fact that he utters it provides information over and above the semantic content of the sentence itself. He puts the hearer in the position of figuring out why he said that rather than something else. He may also be putting the hearer in the position of figuring out why he said what he said in the way he said it rather than in some other way.

Given the sentence actually uttered, the listener’s specific task is to figure out just what the speaker’s communicative intention is, partly on the basis that he is intended to do so. There is much that can be inferred from the fact that a speaker utters the sentence he does and thereby says what he says in uttering it. Part of this is what the speaker meant in saying what he said, and part of this is not meant. Of the latter, part is relevant to figuring out what the speaker meant and part is inferable from the fact that the speaker meant what he meant. For example, the hearer may infer that the speaker thinks it is worth asserting what he did, that he believes it, and that he intends him the hearer to believe it (or to continue to believe it). Otherwise, why would the speaker have bothered to assert it?

A special sort of case involves first-person utterances, where facts about what the speaker says about himself interact with the conditions for successfully and felicitously

doing what he does in saying it. This was already illustrated by Moore sentences. Consider that sincerity is a felicity condition on acts of communication in general, so that, in particular, if a speaker who asserts something is being sincere, he has the belief he is expressing. So if he then proceeds to deny that he believes it, he is pragmatically contradicting himself. Later we will see more examples of the interaction between the semantic content of a sentence and conditions for felicitously using it.

Intuitions about Meaning

Philosophers commonly appeal to their supposedly semantic intuitions, usually billed as “our” intuitions, to justify claims about the meanings of various philosophically interesting expressions. But it’s one thing to treat them as data that need to be taken into account and another to treat them as phenomena that semantic theories need to explain. Semantic intuitions are not the subject matter of semantic theories – semantic facts are. So these intuitions do not need to be explained if they can be explained away. And for good reason: there are serious problems with the reliability of such intuitions and with strategies that depend on relying on them.²

Consider the epistemologically central case of the word ‘know’. Philosophers generally assume that knowing entails justified (or some surrogate of being justified) true belief. But if you test people’s intuitions, you’ll find that none of these conditions is necessary, at least according to some people. For instance, people will say of something that struck hit them as true that they “knew it all along.” So does knowledge not entail belief? Some people insist that they “knew” things that they now acknowledge to be false. So does knowledge not even entail truth? There are college administrators who describe universities as repositories and transmitters of “knowledge,” regardless of how much of what passes for knowledge is not true (or not adequately justified, for that matter). There are cognitive psychologists concerned with the “representation of knowledge,” whether or not what is thus represented is true. And there are sociologists (of knowledge) who study how “knowledge” (true or not) is distributed and manipulated,

² Since I have addressed these problems elsewhere (Bach 2002, 2005a), I won’t dwell on them here. In the balance of this introduction, I will merely state a few claims I have defended before and mention a couple of places in epistemology where allegedly semantic intuitions have played a role.

and many of them don't even think there is such a thing as truth. Now epistemologists might well balk at these uses of the term 'knowledge' and insist that administrators, psychologists, and sociologists use the term loosely, as if it meant 'what passes for knowledge', which it doesn't. In so doing, they would be debunking the semantic intuitions of all those who use the term 'knowledge' in this allegedly loose way. But they still take their own intuitions seriously.

One problem with supposedly semantic intuitions is that they tend to be insensitive to the difference between linguistic meaning and pragmatic regularity. For example, some people think that 'and' has a temporal meaning or even a causal connection meaning because of cases like these:

- (1) Sam took off his shoes and climbed into bed.
- (2) Pierre ate some haggis and then threw up.

However, any suggestion of a particular temporal order or causal connection can be canceled without contradiction and, more importantly, when such a suggestion is part of what is meant, there is a straightforward Gricean explanation for it, via the maxims of relevance and manner. So, for example, it would be misleading for the speaker to say "Sam took off his shoes and climbed into bed" if what he meant was that Sam climbed into bed and then took off his shoes. But being misleading is not the same as speaking falsely. Moreover, if the temporal order is part of what the speaker meant, then what he meant was false insofar as it has the order backwards.

A key feature of any Gricean explanation is that it appeals not just to what was said but the fact *that* it was said. This is what makes the explanation pragmatic. So, as regards the above cases, even though conjunction is a symmetrical relation, putting the events in one order rather than the other is a conspicuous fact about what the speaker did, hence something that the hearer can assume he is intended to take into account in explaining what the speaker did. To explain it, in the way that he is intended to, he needs to understand why the speaker said that rather than the obvious alternative (of putting the conjuncts in the opposite order).

In epistemology, especially since the advent of contextualism, it is common to appeal to intuitions about truth-values of sentences (or of utterances of sentences) used to make

knowledge attributions. Appeals to intuitions about truth-values are dangerous. For one thing, they might be wrong. For another, they might be right, but not about what they're supposed to be about. They could concern what speakers are likely to mean rather than what speakers actually say. Their targets might not be the truth-value of the sentence (even taking any context-sensitivity into account) but rather what the speaker meant in uttering it, which, even if he is using all the expressions in it literally, could be some embellished or enriched version of the semantic content. This frequently happens in the case, for example, of statements of epistemic possibility, when the perspective relative to which the statement is being made, is not made explicit. Finally, intuitions about what a speaker is likely to mean in uttering a given sentence may implicitly assume a certain stereotypical context in which the sentence is likely to be used.

For these reasons it is easy to read things into the meaning of a sentence or into what a speaker says in uttering it that are really consequences of its being uttered under normal circumstances. Our seemingly semantic intuitions are responsive to *pragmatic regularities*. Pragmatic regularities include regularized uses of specific expressions and constructions that go beyond conventional meaning, for example generalized conversational implicatures, as well as general patterns of efficient communication, which involve streamlining stratagems on the part of speakers and inferential heuristics on the parts of listeners. These regularities are pragmatic because it is the speaker's act of uttering a given sentence, not the sentence itself, that carries the additional element of information.

To repeat, *that* a speaker says what he says rather than something else, or says it in one way rather than another, can contribute to what a speaker is likely to mean. Indeed, that he says it one way, by using certain words rather than certain others, can also contribute to what he is likely to mean in using those words. But people's intuitions (even philosophers') about what is said tend to be insensitive to these subtleties. Even so, people's actual reasoning in the course of figuring out what speakers mean takes these considerations into account. So, for example, a listener supposes that:

- If the speaker could have made a stronger statement, he would have.
- If the speaker needed to have been more specific, he would have.
- If there was something special about the situation, the speaker would have mentioned it.

However, if what the speaker actually said is obviously not strong or specific enough but it is obvious how he could have explicitly made it stronger or more specific, he does not need to do so. It is presumed that additional information, were it mentioned, would have needed to be mentioned. In both reasoning generally and rational communication in particular, we presume that things are not out of the ordinary unless there is reason to think that they are, and we presume that if there is such reason, this would occur to us (see Bach 1984). So, as Stephen Levinson puts it, “If [an] utterance is constructed using simple, brief, unmarked forms, this signals business as usual, that the described situation has the expected, stereotypical properties” (2000, 6). Otherwise, as he goes on to say, “If, in contrast, the utterance is constructed using marked, prolix, or unusual forms, this signals that the described situation is itself unusual or unexpected or has special properties.”³ So, when one can speak economically and exploit the default assumptions that apply in a given conversational situation, it would be misleading to make fully explicit what one means.

These considerations suggest that there is no need, at least in general, to expect semantics to explain what people would normally mean in uttering a given sentence. In particular, one shouldn't let features of stereotypical utterance contexts get incorporated into the semantics of a sentence. The speaker's act of uttering that sentence, with that content, provides or invokes additional information, such as mutually evident facts about the context, for the hearer to use in understanding an utterance of the sentence – there is no intermediate level of meaning, so-called “utterance content,” between the semantic content of the sentence and the speaker's communicative intention in uttering it.

In what follows I will present, in some cases very briefly (see the references for fuller discussion), various applications of pragmatics to epistemology. In some cases these concern

³ Levinson identifies three presumptions or “heuristics” that yield defeasible inferences as to a speaker's communicative intention (2000, 31-34):

- Q-heuristic: what isn't said, isn't.
- I-heuristic: What is simply described is stereotypically exemplified.
- R-heuristic: What's said in an abnormal way, isn't normal.

These are closely related, respectively, to Grice's (1989, 25-6) first maxim of quantity (“Make your contribution as informative as is required”, second maxim of quantity (“Do not make your contribution more informative than is required?”), and maxim of manner (“Be perspicuous,” and specifically to the submaxims “Avoid obscurity of expression” and “Avoid prolixity”).

epistemological claims whose plausibility depends on one sort of semantic/pragmatic confusion or another.

1. Classic Examples of Confusing Meaning and Use

Overlooking the semantic/pragmatic distinction led the so-called ordinary language philosophers to make some misguided claims about the meanings of such philosophically interesting terms as ‘tries’, ‘seems’, ‘good’, and ‘true’. They were right to ask what we would mean in uttering sentences containing such terms, but they were wrong to draw pragmatics-laden conclusions about the meanings of these terms. For example, it was thought that trying to do something entails some effort or difficulty in doing it and that something’s seeming to have a certain property entails doubt if not denial that it does have that property. These alleged entailments actually stem from conditions for appropriately and non-misleadingly using sentences containing ‘try’ or ‘seems’, not from their truth-conditions.

Ordinary language philosophers made similar claims about ‘believes’ and ‘knows’. For example, it was suggested that believing entails not knowing, on the grounds that “we would say” that someone believed something only if we thought they did not know it. Indeed, it was even suggested that knowing entails not believing, since if we said that someone knew something, we “would not say” that they believed it.

2. Pragmatic Implications of Hedging or Intensifying an Assertion

If believing that *p* is the sincerity condition on asserting that *p*, why would one ever insert ‘I believe’ into an assertion that *p*? Why, for example, would I utter (4) rather than (3)?

(3) Sheffield is north of Nottingham.

(4) Sheffield is, I believe, north of Nottingham.

The obvious reason is that I’m not sure that Sheffield is north of Nottingham and I want you to be aware of that. I think there’s a significant chance that I’m wrong, and I don’t want you to take my word for it without checking it out for yourself. So, by inserting the parenthetical ‘I believe’ I am weakening my assertion, or at least what I want you to infer from the fact that I made it. In particular, I’m implicating that I don’t know that Sheffield is north of Nottingham.

If I were to have uttered (5) instead of (4), I would still have been asserting that Sheffield is north of Nottingham, even though, taken literally I would be taken to be asserting merely that I believe that Sheffield is north of Nottingham.

(5) I believe that Sheffield is north of Nottingham.

There are other ways to hedge an assertion. For example, I could have uttered (6) or (7):

(6) Sheffield is north of Nottingham, or so I think.

(7) As I recall, Sheffield is north of Nottingham.

Using a sentence like (8) has a different import.

(8) As far as I know, Sheffield is north of Nottingham.

In this case I would be specifically implicating that my information could be faulty and that others know better.

Notice that in all of these cases part of what the hearer has to do is figure out why the speaker said what he said rather than the alternative, in this case a hedged sentence rather than an unhedged one.

Uttering ‘I know that p’ is stronger than simply uttering ‘p’, even though in asserting that p I am representing (though not asserting) that I know that p. Saying that I know indicates authority, confidence, and absence of doubt. However, frequently one utters ‘I know that’ to do something else, namely, to indicate that one already has the information. This is illustrated when a question is put to a group of people, say students, and someone blurts out, “I know (the answer).” I frequently get illustrations of this from my 5-year old granddaughter Sophia, who often responds to things I tell her by saying, “I know that.”

Now let’s look at the term ‘certain’, which is a bit tricky because we predicate certainty of both people and propositions. We can be certain of something, meaning we are convinced of it and have no doubt about it, and something can be certain if there is no chance it is not the case. We can be certain of something that isn’t certain, and we can be uncertain of something even if it is certain. The first can happen when we ignore possibilities that we should take into account, and the second can happen when we have doubts that are not warranted by the overwhelming evidence we have.

Despite this difference between being certain of something and something’s being certain, normally we would assert that something is certain only if we are certain of it.

There's a pragmatic explanation for this. If we are certain of something, we don't have any doubts about it, in which case we are prepared to assert that *it* is certain. Of course, we might be wrong, as pointed out above. And we wouldn't (sincerely) assert that something is certain unless we are certain of it. In sum, personal certainty is a matter of conviction and propositional certainty is a matter of being conclusively supported (by currently available information). The two are obviously distinct, but from one's own point of view they go together. Analogous points apply to doubtfulness.

3. Belief Attributions

It is arguable that a correct belief attribution need not fully specify something that the subject believes. Using Saul Kripke's (1979) well-known Paderewski case as an example, I have suggested (in Bach 1997) that one could believe that Paderewski had musical talent and disbelieve that Paderewski had musical talent without believing and disbelieving the same thing. And, I argue, every case is potentially a Paderewski case.

Also, the attribution can include language that characterizes the attributor's, not the subject's, way of thinking of a relevant object or property. The attribution's 'that'-clause may characterize the belief in a way that the subject would neither acknowledge nor endorse. For example, suppose that Vic enters an office where he sees a scruffy looking man, who happens to be the boss, sweeping the floor. It is easy to imagine circumstances in which we could correctly report Vic as believing that the boss is sweeping the floor, even though even though Vic himself, who mistakes the boss for the janitor, would deny that he believes anything of the sort.⁴

Finally, philosophers have been impressed by the fact that belief attributions can take 'believe-of' as well as 'believes-that' forms and sometimes label them, respectively, *de re* and *de dicto* (or descriptive) belief attributions. Many also suppose that these two forms correspond to two types of belief, *de re* and *de dicto*. However, just as the phrase 'French history teacher' is syntactically ambiguous, so is the phrase '*de re* (or *de dicto*) belief attribution'. Under the appropriate circumstances, either type of belief attribution can be used to ascribe either type of belief.

⁴ Both this and the next point are spelled out in (Bach 1987/1994: ch. 10).

4. Knowledge-*wh*

Epistemologists have generally focused on knowing-that, but there is also the case of knowledge-*wh*: knowing who, what, where, when, why, how, or whether. It is an interesting fact that you can know what the ingredients of haggis are but you can't believe what the ingredients of haggis are. I wish I had an explanation for that.⁵

One question about knowing-*wh* is whether it reduces to knowing that. Offhand it would seem to: if you had exhaustive knowledge-that about a certain topic, it seems that there would be no knowledge-*wh* left to for you to have. A related question is whether the same relation is involved. Consider these examples:

(9) George knows who is buried in Grant's Tomb.

(10) George knows that Grant is buried in Grant's Tomb.

I would argue that there is only one (relevant) knowing relation and that you can bear that relation to a *wh*-question as well as to a true proposition. As a first approximation, to bear that relation to a *wh*-question is know the propositional answer (or an answer, depending on the form of the question), where a propositional answer (as opposed to a phrasal answer) takes the logical form of the question but where the variable associated with the *wh*-term is replaced by an identifying noun phrase to yield a true proposition.

Here is a slightly formalized formulation of this idea:

Knowledge-*wh* is a relation to the question (= [*wh*_{*x*} [... *x* ...]_S]_Q),
but to be in that relation is to know the answer to the question:
To know [*wh*_{*x*} [... *x* ...]_S]_Q is to know that [... A ...]_S, where A is
the *x* or an *x* such that [... *x* ...]_S.

Multiple-choice questions require slightly special treatment:

To know [whether {*a, b, c, ...*}_{*x*} [... *x* ...]_S]_Q is to know that [... A ...]_S,
where A is the *x*: $x \in \{a, b, c, \dots\}$ such that [... *x* ...]_S, given that $\exists x$ [... *x* ...]_S.

This formulation is intentionally ambiguous as to whether the 'given'-clause is or is not within the scope of 'know'. That is, I don't want to settle, at least not here, whether knowing that the question has an answer, or merely its having an answer, is necessary for knowing the answer.

⁵ It can't be the factivity of 'know', since there are non-factive verbs, such as 'wonder' and 'guess,' that also can take *wh*-complements.

There is a special question about knowing how, construed as part of the ability to do something. Does it reduce to knowing that? Offhand, it seems that it does not. You can know what it takes to do something and in that sense know how to do it, that is, know how *one* does it, but this does not entail that *you* know how to do it. A similar point applies to the sort of knowing what and knowing when that is involved in complex skills, i.e. knowing what to do when during the course of exercising the skill.

Finally, there is a pragmatic side to attributions of knowledge-*wh* and to what counts as knowing the answer (or knowing an answer). The answer must be relevantly informative. So, for example, knowing that an American invented the zipper does not ordinarily count as knowing who invented the zipper, although in a trivial way that might be enough. However, there is a difference between knowing an answer and knowing a relevantly informative answer, and that depends on what the questioner knows and wants to know (see Braun 2006).

5. The Knowledge Rule on Assertion

Tim Williamson has argued for a knowledge rule on assertion and, indeed, maintains that this rule is somehow constitutive of the speech act of assertion (2000: ch. 11). Jonathan Sutton (2007) has recently invoked this knowledge rule as a basis for his startling claim that justification entails knowledge. Keith DeRose invoked it as the basis for what he advertises as a “powerful” new argument for contextualism and against so-called subject-sensitive invariantism. He sums up this argument as follows:

If the standards for when one is in a position to warrantably assert that P are the same as those that comprise a truth-condition for ‘I know that P’, then if the former vary with context, so do the latter. In short: The knowledge account of assertion together with the context-sensitivity of assertability yields contextualism about knowledge. (2002: 171)

Evidently, for ‘P’ to be warrantably assertible by someone is for that person to know that P, since that is the truth-condition for one’s utterance of ‘I know that P’. Unfortunately, this putatively powerful argument for contextualism applies only to first-person cases, in which attributor and subject are the same (as in DeRose’s many examples). Otherwise, what makes ‘P’ warrantably assertible by the attributor, that the *attributor* knows that P, is not the truth-condition for his utterance of ‘S knows that P’, if S is someone else. So

this argument does not discriminate between contextualism and the view DeRose is arguing against, subject-sensitive invariantism (or what Jason Stanley (2005) calls “interest-relative invariantism”).

But let me get to my main point about the knowledge rule on assertion. Williamson maintains that it is somehow constitutive of assertion. It seems to me, based on general considerations about speech acts, that the only relevant rule on assertion is belief, since an assertion essentially *is* the expression of a belief. This is just a special case of a general fact about communicative speech acts, that they are expressions of attitudes, propositional attitudes in most cases. Just as a request is an expression of a desire and an apology is an expression of regret, so an assertion is the expression of a belief.⁶ Of course, it is one thing to express an attitude and another thing actually to possess that attitude. A speaker is sincere if he has the attitude that he expresses. In the case of assertion, to be sincere is to hold the belief one expresses. So we can say that belief, not knowledge, is the rule on assertion.

However, it is plausible to hold that knowledge is the norm on belief. If that is correct (I won't try to defend it here), then we can derive the knowledge rule on assertion by combining the knowledge norm on belief with the belief rule of assertion. This suggests that there is really nothing special about the knowledge requirement on assertion. It has no independent significance but is, rather, the combination of the knowledge norm on belief and the belief rule on assertion.

6. Testimony

There are various questions concerning what you can come to know from what you are told: Can you come to know something that the speaker himself did not know? Does being told something provide you with evidence that the speaker himself does not have? Does being told something increase your justification for believing it beyond whatever justification you would have from knowing that the speaker believes it? I will not address these questions directly but will instead to make a few observations about the relevance of pragmatics to the discussion of such questions about testimony.

⁶ Bach and Harnish (1979: ch. 3) present a taxonomy of illocutionary acts in terms of different types of expressed attitudes.

The literature on testimony and interpersonal transmission of knowledge tends to neglect some basic platitudes of pragmatics and distinctions from speech act theory. Sometimes it gives the false impression that conveying information is simply a matter of putting something you know into words and that acquiring knowledge of that information is simply a matter of understanding those words. Writers on testimony tend not to take following pragmatic observations into account:

- The speaker might be saying one thing and asserting something else.
- The speaker might not believe what he is asserting.
- The point of his assertion that p might be something other than conveying the information that p.

Related to these observations are the following distinctions from speech act theory, which are often overlooked:

- between locutionary, illocutionary, and perlocutionary acts
- between speaking literally and nonliterally
- between speaking directly and indirectly
- between expressing a belief (or other attitude) and actually having it (being sincere)

When you take these observations and distinctions into account, you can easily see that the simplistic encoding-decoding model of communication is hopeless. And, when you combine them with a Gricean picture of communication as involving a certain sort of intention and correlative inference, you can see that the question of whether to believe something you are told is not altogether separate from figuring out what it is, if anything, that you are being told. The speaker may not mean what he says, and he may not believe what he is asserting. However, before you can consider believing what a speaker asserts or, for that matter, implies, you need to figure out what he is asserting and implying. And doing that requires coming up with something that he could plausibly have meant, given that he said what he said.

Consider one of the questions mentioned above: can being told something by someone give you any more reason for believing it than simply knowing, without the benefit of being told directly, that the person believes it? The null hypothesis is that being told cannot give you any additional reason – it gives you evidence that they believe the proposition but, beyond that, no further reason for believing it yourself. The only thing special about being told something is that this is the primary way, at least in everyday

life, that we find out what other people think (we generally don't overhear their conversations with others, sneak peaks at their diaries, or telepathically read their minds). Being told something does not give you special reason for believing it beyond whatever reason is given by the fact that the speaker believes it. However, the fact that the speaker tells you this may give you other information, such as about what they want to reveal, what they want to conceal, and how they may wish to influence what you think and do. But this is obvious, and trivial.

7. Asserting and Thinking of Possibilities

When one asserts a possibility, generally one isn't asserting that the possibility exists; rather, one is asserting that it is "open" or "live," not "idle," or "far-fetched." In so doing, one is implicating that it is worth considering. Otherwise, why mention it?

In making and evaluating statements of possibility, we do so relative to bodies of information or *perspectives*. Normally, we do so relative to information currently available to us, that is, relative to our own perspective at the time. Something is possible for us if it is not ruled out by that information. However, in asserting epistemic possibilities we commonly use bare, unqualified epistemic possibility (EP) sentences, sentences that do not indicate the perspective from which the possibility is being considered. I mean sentences like these,

(11) Maybe Dick Cheney believes what he says.

(12) Perhaps your keys are still in the front door.

as opposed to explicitly relativized ones like these:

(11R) As far as I know, Dick Cheney believes what he says.

(12R) As far as you know, your keys are still in the front door.

In most conversational situations we don't need to make the perspective explicit because it is evident what the relevant perspective is. Accordingly, we can effectively use bare EP sentences. Even so, such sentences fall short of fully expressing propositions, since they fail to mention (refer to, describe, or quantify over) perspectives and epistemic possibilities are relative to perspectives.

Proponents of currently popular views on epistemic possibility take for granted that bare EP sentences do express propositions. They implicitly take the fact that we have

strong intuitions about the truth or falsity of utterances of such sentences, given particular scenarios, as strong evidence that the sentences themselves are capable of being true or false. Their focus is on how truth-values can shift with a change in a person's perspective or with a shift from one person's perspective to another's. According to *contextualism*, epistemic possibility terms such as 'maybe' and 'perhaps' are like indexicals, so that bare EP sentences are context-sensitive. That is, the specific proposition such a sentence expresses depends on its context of use, which somehow determines the perspective included in the proposition. In contrast, *Relativism* does not treat epistemic modals as anything like indexicals. It maintains that a bare EP sentence expresses a fixed proposition (assuming the sentence is not context-sensitive in some other respect, by virtue of containing any genuine indexicals), but one whose truth or falsity depends on the perspective relative to which or from which it is considered.

Radical Invariantism, as I call my view (Bach forthcoming), denies that the contents of bare EP sentences vary with the context. But it also denies that their contents add up to propositions. EP sentences that are explicitly relativized to perspectives do express propositions, EP propositions that have perspectives as constituents, but bare EP sentences do not. They do not express full-fledged propositions, but only *propositional radicals*, and so are incapable of being true or false. Nevertheless, they are perfectly capable of being used to assert propositions and of being taken as so used, because utterances of them implicitly mention the perspective with respect to which the relevant possibilities are to be considered.

Radical Invariantism agrees with Contextualism on what EP propositions there are, namely ones with perspectives as constituents, but it denies that bare EP sentences semantically express propositions (contextualists holds that they do express propositions but only because they mistake propositional incompleteness for context sensitivity). And Radical Invariantism agrees with Relativism that the semantic contents of bare EP sentences are fixed, but it denies that these contents are fully propositional.

Ignoring as Evidence

The following observations (based on Bach 1985) are not strictly pragmatic, unless we recognize a kind of mental pragmatics in addition to pragmatics in communication. It is

plausible to suppose that thinking of a possibility is evidence that it is worth considering, at least in areas where you are knowledgeable. This means that not thinking of a possibility is (defeasible) evidence that it is not worth considering. However, at the time this can only be implicit evidence for one. Here's why. If you had to notice that you did not think of a certain possibility in order for your not thinking of it to be evidence that it is not worth considering, you would have to think of it. Moreover, not being able to get a possibility out of your mind is (defeasible) evidence that it needs to be taken seriously. Not being able to do so keeps you from firmly believing anything that precludes the possibility. So, in order to confidently believe such a thing, you must rule the possibility out. Successfully ruling it out not only requires coming up with good reason to reject it but keeping it from coming to mind any further. Lack of confident belief on account of that possibility is evidence (again defeasible) that the possibility needs to be ruled out.⁷

You don't have to be a contextualist to agree, with respect to a given bit of putative knowledge, about which counterpossibilities need to be considered and eliminated and which ones can be, in David Lewis's (1996) phrase, "properly ignored." As I was just suggesting, there is a common and commonly overlooked way in which the consideration or non-consideration of counterpossibilities is relevant to having knowledge. In forming beliefs and seeking knowledge, we implicitly rely on our reliability to think of and thereby consider counterpossibilities when and only when they are worth considering. Part of what makes a belief justified is that the cognitive processes whereby it is formed and sustained are sensitive to realistic counterpossibilities (so-called relevant alternatives). The very occurrence of the thought of a counterpossibility gives one prima facie reason to take it seriously. And the fact that a counterpossibility does not come to mind is evidence for its irrelevance. But that fact is evidence that one cannot explicitly consider, since to do so would be to bring the counterpossibility to mind. To the extent that we can trust our ability to know when there are no further counterpossibilities epistemically worth considering, we don't have to consider them in order to be justified in treating them as not worth considering.

⁷ It is defeasible evidence if only because of cases like obsessively thinking that one left the stove on or the front door unlocked. The unbidden recurrence of such thoughts is better evidence for one's inability to keep them from coming to mind.

8. Concessive Knowledge Attributions

Fallibilism is the view, now widely accepted, that something can be known without being certain. This view gives rise to a little puzzle about concessive knowledge attributions. Solving it combines considerations about knowledge attributions in general with considerations about statements of epistemic possibility. What is puzzling is that statements of the following forms do not seem capable of being true:

(13) Dave knows that Canberra is the capital of Australia, but it might be Sydney.

(14) I know that Canberra is the capital of Australia, but it might be Sydney.

The puzzle is based on the assumption that knowing something precludes its falsity. Of course, actual truth does not preclude possible falsity, but perhaps actual knowing does preclude it, at least in the epistemic sense. The puzzle, then, is to reconcile fallibilism with the apparent fact that statements like (13) and (14) cannot be true.

To solve this puzzle we should first note that fallibilism is an anti-skeptical doctrine, hence that the possibility of error it underwrites is general, not specific to particular knowledge claims. Moreover, it seems to me, this possibility is weaker than what speakers assert when making *might*-statements, which allude to particular counterpossibilities or sources of error specifically relevant to the knowledge claim in question.

Without going deeper into the question of fallibilism and related questions about knowledge, let me just make a couple of simple pragmatic observations about sentences like (13) and (14) and the statements they are used to make. Recall that epistemic possibility is perspective-relative and that, as suggested earlier, a *sentence* like ‘Sydney might be the capital of Australia’ does not manage to express a proposition and is not capable of being true or false. Accordingly, since the *might*-clauses in (13) and (14) are not explicitly relativized, (13) and (14) are neither true nor false as they stand. However, *utterances* of them can be true or false. That is, what a speaker asserts in uttering one of them can be true or false. Suppose you utter (13). It could be true that Dave knows that Canberra is the capital of Australia, even though as far as you know, Sydney might be the capital of Australia. Then what you assert would be true. However, you could not coherently assert it, since your reason for asserting that Sydney might be the capital of

Australia would be sufficient reason for you not to assert that Dave knows that Canberra is. The situation is different with an explicitly relativized version of the first-person (14), ‘I know that Canberra is the capital of Australia, but as far as I know it might be Sydney’, because the two clauses really do contradict each other. In contrast, 3rd-person concessive knowledge attributions like (13) are merely pragmatically contradictory. Asserting that someone else knows a certain thing presupposes that you yourself know it, but to raise the possibility of error is to raise doubt that the subject knows.

9. “Pragmatic Encroachment”

Recently it has been argued that whether or not one knows something is not a purely cognitive matter and that the stakes, the cost of being wrong, can affect what it takes to know (Fantl and McGrath 2002 and Stanley 2005). If this is right, then it is possible that of two people with the same evidence and the same degree of purely epistemic justification, one could know while the other does not.

What basis is there for this startling view? The more important it is to be right about whether or not *p*, the greater effort one should make to make sure that one isn’t wrong. One should guard against possibilities that normally one could justifiably ignore, check into things that one could normally take for granted, and double-check things that one has already checked into. Therefore, as the stakes go up it takes more to justifiably believe that *p*, hence more to know that *p*. What is good enough evidence when the stakes are low is not good enough when the stakes are higher.

It is hard to argue with the claim that as the stakes go up, then, everything else equal (for example, there is plenty of time and the cost of checking further and guarding against counterpossibilities is small compared to the stakes), one should take greater care in the ways indicated. But is this a matter of epistemic or merely practical rationality? Being practically rational requires risk management, and guarding against error is part of this. Since the size of the stakes increases the cost of being wrong, it also increases the range of possibilities one should guard against. It requires not reflexively acting on assumptions one would normally make and act on without question, not if there is a chance they are false and acting on them if they are false would be costly.

This much makes perfectly good sense. However, it does not follow that the standards of knowledge go up (or that the threshold of justification is higher or that the range of relevant alternatives to be ruled out is wider). Rather, the higher stakes raise the threshold of confident, doubt-free belief. One's practical reasons give one reason not to treat certain possibilities as closed, hence not to act as if they do not obtain. One's practical interest explains the rise in the threshold of confident, settled belief, and thoughts of counterpossibilities make it more difficult for this threshold to be crossed. Lack of doubt-free belief keeps one from meeting the doxastic condition on knowledge. Proponents of pragmatic encroachment assume that the justification condition on knowledge has gone up in cases of higher stakes. In my view, one doesn't know in these cases because the stakes have raised the threshold of confident, doubt-free belief.

This account of the import of high stakes straightforwardly applies to cases in which it is the agent's stakes that are high and the agent is aware of the relevant risks. But there are two other cases to consider, where the agent is not aware of the risks and where it is not the agent's stakes but the attributor's that are high. In the first case, where the agent is unaware of a certain risk, say that there is a significant chance that the food he has just been served contains an ingredient to which, unbeknownst to him, he and a few other people are highly allergic, it would seem that if most other people served such food would know it is safe, then so would he. On this view, the heightened risk to him does not affect his ability to know on the same grounds as anyone else. Alternatively, it could be argued that there is a relevant difference, namely that for each person served this food the relevant proposition is that this food is dangerous to that person. So, if you're a rare person for whom this food is dangerous (unbeknownst to you), the proposition that it is safe for you is more difficult to know to be true than the analogous proposition regarding someone else. In this case, it is not that your practical interest differentially affects what it takes for *you* to know the relevant proposition; rather, it affects what it takes for *anyone* to know that proposition. This suggests that important things are harder to know, regardless of who they are important to. I have my doubts about that, but this is not the view endorsed by proponents of pragmatic encroachment.

In the other sort of case, the *attributor's* practical interest keeps him from attributing knowledge to an agent without that interest. This is what happens in Stewart Cohen's well-known airport example (1999: 58). Because of the higher stakes and her correspondingly high threshold of confident belief, Mary is not willing to say "Smith knows that the plane will stop in Chicago" and goes so far as to assert its negation because of her own doxastic situation. Because she is not sure Smith's itinerary is reliable, she is not confident enough to believe that the plane will stop in Chicago. Accordingly, she can't coherently attribute knowledge of it to Smith, not if knowledge implies truth. In general, you can't coherently assert that someone else knows that p if you are not confident that p and think that it still needs to be verified. So Mary can't very well assert that Smith knows that the plane will stop in Chicago. Not only that, she has to deny that she knows it, since she thinks it is not yet established. And, since Smith has no evidence that she doesn't have, she must also deny that he knows it.

10. Epistemic Contextualism

According to this widely discussed and highly controversial doctrine, a sentence of the simple form 'S knows (at t) that p' can be true as uttered in one context and false as uttered in another, depending on the epistemic requirements (standards, level of justification, range of relevant alternatives) that govern the context.⁸ Contextualism says that the propositional content of such a sentence varies with the context. The requirements governing the context help determine which knowledge-attributing proposition the sentence expresses in that context.⁹ This means that there is no one knowledge relation and that the different propositions expressible by such a sentence involve different relations. However, even if contextualism were true and, contrary to invariantism, a given knowledge-ascribing *sentence* could express various propositions in various contexts, it would be less dramatic than it seems, for those *propositions* would not themselves be context-bound. Each such proposition could be expressed by a more

⁸ The best-known defenses of epistemic contextualism are Cohen 1999, DeRose 1995, and Lewis 1996. My opinion of this doctrine is evident from the title of Bach 2005b, "The Emperor's New 'Knows'." This section gives the gist of that paper.

⁹ Here I am going along with pretense that contextual variability is context dependence, as if it is context that determines the semantic value of a contextually variable expression. For more on the many uses and abuses of context, see Bach 2005a.

elaborate knowledge-ascribing sentence in which ‘knows’ is indexed or relativized to an epistemic standard. Each such proposition would be expressible in any context by the same context-insensitive sentence and would have a fixed truth-value.

The contextualist’s strategy for explaining the lure of skeptical arguments and resolving skeptical paradoxes requires that epistemologists not be cognizant of the previous point. If the relevant propositions are spelled out, there is nothing to be confused about. Also, contextualism imputes confusion to ordinary knowledge attributors: they talk as if the standards for knowledge vary with context and yet unwittingly take it be one and the same thing, knowledge, that is at issue from context to context. So contextualism is a strong error theory. Moreover, as I will suggest, the contextualist’s strategy for containing skepticism does not really explain why unsuspecting people can be duped by skeptical arguments. An alternative explanation is that the conflicting intuitions that give rise to skeptical paradoxes don’t really bear on the truth conditions of knowledge attributions but merely reflect vacillating responses to skeptical considerations.

Contextualism doesn’t really come to grips with skepticism anyway. In attempting to confine the plausibility of skeptical arguments to contexts in which far-fetched skeptical possibilities are raised, it concedes both too much to the skeptic and too little. Too much because it grants that in skeptical contexts skeptical arguments are sound. And too little because those arguments, however cogent or fallacious they may be, purport to show that *ordinary* knowledge attributions are generally false. Skeptics argue not merely that we don’t have empirical knowledge by the highest standards but that we don’t have it at all (or at least not very much of it), by *ordinary* standards. How assiduously people *apply* epistemic standards may vary from context to context, but the skeptic denies that the standards themselves come in various strengths.

One last question: does the invariantist need to resort to pragmatics to explain the relevant differences in these cases? In my view, invariantism, whether skeptical or moderate, need not rely on what Keith DeRose disparagingly calls “warranted assertibility maneuvers,” or WAMs. Here, for once, I do not seek a pragmatic way out of a philosophical predicament, or at least not that way out. I deal with examples like DeRose’s bank case and Stewart Cohen’s airport case with the same basic strategy as I

used to explain away the appearance of pragmatic encroachment. Contextualists appeal to intuitions about these cases, cases in which (1) the standards prevailing in the attributor's context are relatively high compared to normal, low-standards contexts, such as the subject's, (2) the attributor's and the subject's evidence is the same (identical or at least of equal strength), and (3) intuitively, (a) a 'knows'-attribution made of the subject in a normal, low-standards context is true, but (b) the one made of the subject by the attributor in her high-standards context is false. However, the intuitions about these examples that seem to support contextualism can be explained away. They do not support the view that 'know' is context-sensitive and that the truth conditions of knowledge-ascribing sentences can vary with the context in which they are uttered. What varies, rather, is the attributor's threshold of confidence. In the problem cases, I suggest, either a practical concern or an overly demanding epistemic consideration raises the attributor's threshold of confidence and leads the attributor to demand more evidence than knowledge requires. In contexts where special concerns arise, whether practical or skeptical, what varies is not the truth conditions of knowledge attributions but the knowledge attributions people are prepared to make. It is not the standards for the truth of knowledge attributions that go up but the attributor's threshold of confidence regarding the relevant proposition.¹⁰

Although I don't go in for WAMs here, I do think that certain pragmatic factors play a role in driving our intuitions about the truth-values of knowledge attributions. These intuitions may be responsive to something other than the explicit content of the attribution. I suggest that we often take a given knowledge attribution as implicitly conditional, in one of two ways. First, much of what passes for categorical knowledge

¹⁰ I will not take up epistemic relativism, best advocated by John MacFarlane (2005), but just mention it. Relativism is designed to account for contextualist intuitions. Like contextualists, relativists reject invariantism about 'knows'. They agree that a sentence of the simple form 'S knows (at t) that p' can be true relative to one context and false relative to another, but they deny that the propositional content of such a sentence varies with the context. Rather, it is the truth-value that varies. And the context in question is not limited to the context of utterance but can also be the context of assessment. Moreover, it is the same proposition whose truth-value can vary, say the proposition that Plato taught Aristotle. That proposition could be true from George's perspective and false from René's, even if George and René have exactly the same evidence. Clearly relativism requires introducing a special sort of proposition. It also invites the objection that it cannot account for genuine disagreement about whether a certain person knows a certain thing. Whereas contextualism is subject to the objection that it turns disagreement into changing the subject (since a different proposition is involved), relativism is subject to the objection that people who disagree can nevertheless both be right. At any rate, these problems can be avoided if we stick to invariantism, which in my view can account for contextualist intuitions.

(that p) may really be merely conditional in content (that p , if C , where C is a kind of *ceteris paribus* condition). For example, perhaps Smith doesn't categorically know that his plane will stop in Chicago. Perhaps what he knows is that his plane will stop in Chicago provided the itinerary is accurate, the flight won't be diverted from Chicago because of bad weather (the most likely counterpossibility), and the plane won't crash or be hijacked (etc.). Alternatively, the knowledge attribution, insofar as it can be deemed true, might best be interpreted as conditional on the implicit assumption that things are normal. For example, an utterance of 'Smith knows that the plane will stop in Chicago' is not categorically true but true only if understood as a loose, casual way of saying, 'Provided things are normal, Smith knows that the plane will stop in Chicago'.

References

- Bach, Kent (1984), "Default Reasoning: Jumping to Conclusions and Knowing When to Think Twice," *Pacific Philosophical Quarterly* 65: 37-58.
- _____ (1985), "A Rationale for Reliabilism," *The Monist* 62: 248-63. Reprinted in F. Dretske and Sven Bernecker (eds.), *Knowledge*, Oxford: Oxford University Press 2000, 199-213.
- _____ (1987/1994), *Thought and Reference*, expanded pbk. ed., Oxford: Oxford University Press.
- _____ (1997), "Do Belief Reports Report Beliefs?" *Pacific Philosophical Quarterly*, 78: 215-241.
- _____ (2002), "Seemingly Semantic Intuitions," in J. Campbell, M. O'Rourke, and D. Shier (eds.), *Meaning and Truth*, New York: Seven Bridges Press, 21-33.
- _____ (2005a), "Context *ex Machina*," in Z. Szabó (ed.), *Semantics vs. Pragmatics* (Oxford: Oxford University Press), 15-44.
- _____ (2005b), "The Emperor's New 'Knows'," in G. Preyer and G. Peter (eds.), *Contextualism in Philosophy: Epistemology, Language, and Truth*, Oxford: Oxford University Press, 51-89.
- _____ forthcoming, "Perspectives on Possibilities: Contextualism, Relativism, or What?"
- _____ and Robert M. Harnish (1979), *Linguistic Communication and Speech Acts*, Cambridge, MA: MIT Press.
- Braun, David (2006), "Now You Know Who Hong Oak Yun Is," *Philosophical Issues* 16: 24-42.
- Cohen, Stewart (1999), "Contextualism, Skepticism, and the Structure of Reasons," *Philosophical Perspectives* 13: 57-89.
- DeRose, Keith (1995), "Solving the Skeptical Problem," *Philosophical Review* 104: 1-52.
- _____ (2002), "Assertion, Knowledge and Context," *Philosophical Review* 111: 167-203.
- Fantl, Jeremy and Matthew McGrath (2002), "Evidence, Pragmatics, and Justification," *Philosophical Review* 111: 67-94.
- Grice, Paul (1961), "The Causal Theory of Perception," *Proceedings of the Aristotelian Society*, supp. vol. 35: 121-152.
- _____ (1989), *Studies in the Way of Words*, Cambridge, MA: Harvard University Press.
- Kripke, Saul (1979), "A Puzzle about Belief," in A. Margalit (ed.), *Meaning and Use*, Dordrecht: Reidel, 239-283.
- Levinson, Stephen C. (2000), *Presumptive Meanings: The Theory of Generalized Conversational Implicature*, Cambridge, MA: MIT Press.
- Lewis, David (1996), "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549-67.
- John MacFarlane (2005), "The Assessment Sensitivity of Knowledge Attributions," in T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Oxford: Oxford University Press, 197-233.
- Stanley, Jason (2005), *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Sutton, Jonathan (2007), *Without Justification*, Cambridge, MA: MIT Press.
- Williamson, Timothy (2000), *Knowledge and its Limits*. Oxford: Oxford University Press.