

RELATIVISM AND KNOWLEDGE ATTRIBUTIONS

John MacFarlane

Relativism, in the sense at issue here, is a view about the meaning of knowledge attributions—statements of the form “S knows that *p*.” Like contextualism, it holds that the truth of knowledge claims is sensitive to contextual factors, such as which alternatives are relevant at the context, or how high the stakes are. For the relativist, however, the relevant context is the context from which the knowledge claim is being *assessed*, not the context at which it was *made*.

The Relativist’s Position

One kind of relativist position can be defined by its acceptance of the following four theses:

Local Invariantism:

The relation expressed by “know” does not vary with context.

Dyadic Relation:

This relation is a two-place relation; it does not have implicit argument places that must be filled through hidden variables in the logical form or “free enrichment.”

Fancy Intentions:

Although “know” invariantly expresses the knowledge relation, this relation does not have an intension of the familiar sort—a function from possible worlds and times to truth values. That is, its extension is not determined by the state of the world at a time, but depends on something else in addition, which we will call an *epistemic standard*.

Assessment Sensitivity:

The accuracy of an assertion or belief depends on the epistemic standard that is relevant at the context of assessment. Thus, there is no “absolute” answer to the question whether such an assertion or belief is accurate; accuracy is an assessment-sensitive matter.

A few words of explanation are in order.

Local Invariantism is rejected by philosophers who take “know” to be an *indexical*. An indexical is a word whose content (its contribution to the propositions expressed by sentences of which it is a part) is determined in part by features of the context. A paradigm is “today,” which denotes the day on which it is uttered. Some epistemic contextualists (Cohen 1988: 97; DeRose 1996: 194 n.4) hold that “know” is indexical, meaning that there are many different knowledge relations, and the particular relation expressed by “know” on an occasion of use depends on contextual factors such as relevant alternatives or practical stakes.

One might accept Local Invariantism, however, while still holding that the contents of sentences of the form “S knows that *p*” depend on contextual factors of this kind. For one might reject Dyadic Relation and hold that “know” invariantly expresses a *three-place* relation between a person, a proposition, and something else—perhaps an epistemic standard, a set of alternatives, or a “question” (Schaffer 2004b). Although no part of the English sentence explicitly denotes the third relatum, it might be denoted by an “aphonic” element in the deep syntax, or the speaker might simply expect hearers to be able to fill in the blank using contextual cues. Let us call views with this shape *relational contextualism*.

Views that accept both Local Invariantism and Dyadic Relation take knowledge-attributing sentences not containing other indexical or demonstrative elements to have contents that are invariant over contexts of use. All such views, then, can be characterized as forms of invariantism. We can distinguish, however, between *standard invariantism*, which rejects both Fancy Intentions and Assessment Sensitivity; *nonindexical contextualism*, which accepts Fancy Intentions and rejects Assessment Sensitivity; and *truth relativism*, which accepts both Fancy Intentions and Assessment Sensitivity.

According to standard invariantism, “know” invariantly expresses a two-place relation between persons and propositions—the relation of knowing—and this relation has an intension of the standard sort: a function from worlds and times to extensions. There is still a lot of room for arguing about what that intension looks like. *Skeptical invariantists* say that *x* stands in the knowing relation to *p* at a world *w* and time *t* just in case *x*’s evidence for *p* at *w* and *t* is strong enough to rule out any possibility that *p* is false, while *dogmatic invariantists* propose a more relaxed condition, and *subject-sensitive invariantists* hold that the strength of evidence required for *x* to stand in the knowing relation to *p* at *w* and *t* depends on aspects of *x*’s practical situation at *w* and *t* (Hawthorne 2004: chs. 3–4). But although these different invariantist positions disagree about which worldly states of affairs suffice for a person to stand in the knowing relation to a proposition, they all assume that there is a definite, context-independent, answer to this question. That assumption is precisely what relativists and nonindexical contextualists deny in accepting Fancy Intentions.

According to Fancy Intentions, there is no answer to the question at issue between dogmatic, skeptical, and subject-sensitive invariantists, because the knowing relation does not have any particular extension at a world and a time. In order to get an extension, one must specify not just a world and a time, but also an *epistemic standard*, which determines how well placed a subject must be in order to stand in the knowing relation to a proposition. We can be neutral here about what an epistemic standard consists in. On some views, it will be defined by a threshold strength of evidence required for knowledge; on others, by a set of relevant alternatives that must be excluded if the subject is to have knowledge. In what follows we will talk of standards being “low” and

“high,” but we do not mean to imply that standards must be linearly ordered. Sets of relevant alternatives, for example, are only partially ordered.

In order to explain Assessment Sensitivity, which separates relativists from non-indexical contextualists, we must say a few words about what is meant here by “accuracy.” It is standard in semantics to think of contents as having their extensions relative to possible worlds, or in some cases worlds and times (see King 2003 for a nice discussion of some of the issues relevant to choosing between these two options). That means that propositions have truth values relative to worlds (or worlds and times). But in these frameworks, whether a speaker has gotten it right—whether her assertion is accurate—depends only on the truth value of the asserted content relative to the world and time of the context of use. Thus, even though the content of an assertion is something that has different truth values relative to different worlds (and perhaps times), whether a particular assertion is accurate is, in standard frameworks, an absolute matter.

The point is perhaps easiest to see if we work with “tensed” propositions, which have different truth values relative to different times. According to temporalists, in saying “Socrates is sitting,” one expresses the tensed proposition that Socrates is sitting—a proposition that is true relative to some times (noon) and false relative to others (midnight). (“Eternalists,” by contrast, hold that this sentence expresses different eternal propositions when used at different times.) Suppose Jake asserts this proposition at noon, when Socrates is sitting on a bench at the agora, and Sally asserts it at midnight, when Socrates is lying down at the symposium. Then, as both parties can recognize, Jake’s assertion is accurate, while Sally’s is not. Even though the truth of the proposition asserted is time-relative, the accuracy of an assertion of it is not; the accuracy of an assertion depends on the truth value of its content at the time the assertion was made.

The point can also be made with standard “eternalist” propositions. Suppose we are considering an assertion, in a possible world where diamonds can be found just about everywhere, of the proposition that diamonds are rare. Although this would be an assertion of a proposition that is true in the actual world, it would nonetheless be inaccurate, since diamonds are not rare in the world at which the assertion is made. Accuracy hangs on truth in the world in which the assertion is made.

But what about the new coordinate that Fancy Intentions adds to intentions—the epistemic standard? Which epistemic standard do we look at in determining whether a particular assertion or belief that *p* is accurate? A natural view is that, just as accuracy depends on the world and time of the context of use, so it depends on the epistemic standard that is relevant at the context of use. This view would resemble indexical contextualism in taking the extension of “know” (as used at a particular context) to depend not just on the worldly state of affairs but on a contextually relevant epistemic standard. It would depart from it, however, in denying that the *content* of “know” varies with the contextually relevant epistemic standard. MacFarlane (2005b, 2007a, 2007b, 2009) has dubbed such views forms of *nonindexical contextualism*. (Nonindexical contextualist accounts of “know” are defended by Kompa (2002) and Brogaard (2008).)

But this decision about how to treat the epistemic standard coordinate when evaluating uses of propositions is not the only possible one. One could, alternatively, take the accuracy of an assertion or belief that *p* to depend on the truth of *p* at the world and time of the context of use and the epistemic standard that is relevant at the context of *assessment*. On this view, there are no absolute facts about accuracy; a particular assertion or belief might be accurate as assessed from one context and inaccurate as assessed

from another, if different epistemic standards are relevant at the two contexts. It is this move—relativizing accuracy—that raises the philosophical questions that have traditionally been associated with “relative truth.” So this view deserves the label *truth relativism*. (Relativist accounts of “know” are defended by Richard (2004) and MacFarlane (2005a).)

Instead of relativizing truth, one might take the *contents* of knowledge attributions to be assessment-sensitive. According to *content relativism*, there is no assessment-independent fact of the matter about what proposition is asserted by an utterance of a knowledge-attributing sentence. This view accepts Assessment Sensitivity, because the accuracy of an assertion depends in part on what its content is, and according to content relativism, this depends on the context of assessment. It parts company with truth relativism, however, in rejecting Fancy Intentions and either Local Invariantism or Dyadic Relation. *Indexical content relativism* resembles indexical contextualism in rejecting Local Invariantism, while *relational content relativism* resembles relational contextualism in rejecting Dyadic Relation. (For a defense of a form of content relativism, see Cappelen 2008.)

We can categorize the positions we have discussed by looking at which of the four principles they accept (Table 49.1).

The Case for Relativism

Relativism about knowledge attributions can be seen as a kind of synthesis of contextualism and invariantism. Like the contextualist, the relativist holds that the truth of knowledge attributions is relative to contextual epistemic standards. And like the invariantist, the relativist holds that “know” expresses the same relation at every context of use. So a case for relativism can be made by cobbling together the best of the invariantists’ arguments against contextualism and the best of the contextualists’ arguments against invariantism. Relativism promises to retain what is right about both contextualism and invariantism while avoiding their flaws.

Against Contextualism

The literature contains quite a few different arguments against contextualism. Some of these apply only to indexical contextualism, while others apply also to relationalism and/or nonindexical contextualism.

Table 49.1

	<i>Local Invariantism</i>	<i>Dyadic Relation</i>	<i>Fancy Intentions</i>	<i>Assessment Sensitivity</i>
Indexical contextualism		✓		
Relational contextualism	✓			
Invariantism	✓	✓		
Nonindexical contextualism	✓	✓	✓	
Truth relativism	✓	✓	✓	✓
Indexical content relativism		✓		✓
Relational content relativism	✓			✓

One argument against indexical contextualism and relationalism is that competent speakers ought to know whether the contents of their knowledge claims are contextually sensitive. Yet, as Schiffer (1996: 326–7) observes, “no ordinary person who utters ‘I know that p,’ however articulate, would dream of telling you that what he meant and was implicitly stating was that he knew that p relative to such-and-such standard” (see also Feldman 2001: 74, 78–9; Hawthorne 2004: §2.7). (For a defense of the contextualist’s imputation of “semantic blindness” here, see DeRose 2006.)

Another argument (Hawthorne 2004: §2.7) is that “know” does not behave like paradigm context-sensitive terms inside attitude reports. If Joe, in a high-stakes context, says “I don’t know whether the bank will be open tomorrow,” Sarah can say the next day, in a low-stakes context, “Joe said he didn’t know whether the bank would be open today.” Sarah uses “today” instead of “tomorrow” in reporting the content of Joe’s claim, because “tomorrow” is indexical and would not have the same content if she were to use it that it did when Joe used it. But she does not find another word to replace “know,” and this suggests that she takes “know” as she uses it to express the same relation that it expressed when Joe used it the day before. However, this argument does not have any force against nonindexical contextualism, which holds that the *content* of “know” remains invariant between contexts. Nor does it rule out relationalist accounts of “know,” which allow that the reporter can “fill in” the implicit argument place in the same way as the original speaker did.

A third argument, which counts against all three forms of contextualism, is that contextualist views make faulty predictions about agreement and disagreement, and about proprieties for correction and retraction of assertions (Feldman 2001: 77; Rosenberg 2002: 164; Hawthorne 2004: 163; Richard 2004: 216–17; MacFarlane 2005a: §2.3; Stanley 2005: 52–6). Here there seems to be a real contrast between “know” and context-sensitive words like “tall.” If Joe says that Chiara is tall (meaning tall *for a fifth-grader*) and Sarah says that she is not tall (meaning tall *for an American female*), they have not disagreed, and (barring misunderstanding) Joe will not take Sarah’s claim to be any kind of challenge to his own. It would be positively bizarre for Joe to say to Sarah, “Yes, you’re right, she *isn’t* tall after all; I take back what I said.” Things are otherwise with “know.” If Joe (in a low-standards situation) says, “I know that the bank is open on Saturday,” and Sarah later says (in a high-standards situation), “You didn’t know that the bank was open,” Joe will naturally take Sarah’s claim as a challenge to his own, and either defend his claim or withdraw it. We do not expect him to say (as the contextualist account would suggest he should): “Yes, you’re right, I didn’t know. Still, what I said was true, and I stick by it. I only meant that I knew-by-low-standards.” Similarly, the skeptic regards himself as disagreeing with ordinary knowledge claims, but if the contextualist is right, this is just a confusion.

Thus the contextualist seems forced to say that ordinary speakers are *mistakenly* taking themselves to disagree (or to agree). But attributing this kind of error tends to undermine the positive case for contextualism, which rests largely on observations about speakers’ propensities to use “know” in various contexts. The more semantic and substantive error we attribute to speakers, the less their usage can tell us about the meanings of their words.

Against Invariantism

Invariantists have an easy time explaining speakers’ perceptions of agreement and disagreement, since they take them to be veridical. But they have a hard time explaining

the basic data that motivate contextualism. There *does* seem to be some variation in the strength of the epistemic position one must be in if one is to count as “knowing.” Contextualists explain that by saying that the extension of “know” is sensitive to contextual factors. How can invariantists explain it?

One strategy would be to attribute this variation to speaker error. Perhaps speakers systematically misjudge the strength of subjects’ epistemic positions, and that explains why they are readier to count someone as “knowing” in some situations than in others. Speakers certainly do make mistakes of this kind. However, this strategy is committed to positing a source of *systematic* error that precisely mirrors the contextual variation we see in the usage of “know.” This is a tall order.

A second strategy is to appeal to loose use or figurative uses of language. Perhaps knowledge demands a very high standard of evidence—so high that we rarely if ever meet it. If we wish to speak literal truth, then, we should not say that we know, but describe our evidence more precisely, acknowledging its limitations. In practice, though, this would often be pedantic and pointless. Just as it serves our purposes to say that it’s noon, when in fact it is one minute past noon, so it might serve our purposes to say that we know, when in fact we only approximate the epistemic position required for knowledge (Schaffer 2004a; Davis 2007). This strategy concedes that the skeptic is correct that, strictly speaking, knowledge is rare, and holds that our knowledge claims can be explained as loose use.

One problem with the loose use strategy is that speakers are normally conscious of their loose use. We might ask: “Strictly speaking, is it noon, or one minute past noon?” And the loose talker will say (with an exasperated grumble): “Okay, if you want to be precise, it’s not noon, it’s one minute past.” So the loose use strategy requires that speakers are normally aware that their knowledge claims are not strictly true. If that were so, however, skepticism would be universally accepted as true but uninteresting. And it isn’t.

A loose use theorist might hold, alternatively, that standards for knowledge are invariant but moderate, so that the skeptic’s denials of knowledge are false. Uses of “know” that seem to assume a stricter standard could perhaps be explained by appeal to implicatures (Rysiew 2001). One would also like some explanation, however, of why the standard for knowledge is what it is, and not something stricter or laxer.

A third invariantist strategy for explaining the apparent variation in standards for knowing is subject-sensitive invariantism (SSI) (Hawthorne 2004; Stanley 2005). According to SSI, the standard a subject must meet in order to stand in the knowledge relation to a proposition (at a particular possible world and time) depends on the subject’s practical situation (at that world and time). This is not a contextualist position, because the relevant epistemic standard is fixed by the *subject’s* situation, not the speaker’s. Nor does it require “fancy intensions”: since the world and time of evaluation fix a relevant standard for the subject, the intension of “know” does not need to be separately relativized to a standard.

SSI yields the same predictions as contextualism about present-tense, first-person knowledge attributions (“I know that p”), since for these the speaker’s context and the subject’s circumstances coincide. To distinguish the two views, we must look at third-person or past-tense knowledge attributions. Contextualism predicts that the accuracy of assertions of “Joe knew last Friday that the bank would be open on Saturday” can depend on aspects of the speaker’s current situation (for example, whether the stakes are high, or whether a given possibility is contextually salient), while SSI predicts that it

will depend only on Joe's situation last Friday. There is considerable disagreement about which of these predictions is better supported (see Hawthorne 2004; Stanley 2005; DeRose 2004, 2005). DeRose (2005: 189) notes that the contextualist can accommodate many cases that might seem to support SSI, because "sometimes speakers' own conversational purposes call for employing standards that are appropriate to the practical situation of the far-away subjects they are discussing." But, he observes, it is hard to see how SSI could accommodate the cases that seem to support contextualism.

Relatedly, SSI predicts that when the subject's circumstances are shifted by a temporal or modal embedding, there should be a corresponding shift in the standards the subject must meet in order to count as knowing. Thus, it predicts that sentences such as the following should come out true:

- (1) Joe doesn't know that the bank is open on Saturdays, but five minutes ago, before he learned that he would have to pay for emergency surgery on Sunday, he did know that it is open on Saturdays.
- (2) I don't know whether the bank is open on Saturdays, but if I didn't really need the money on Sunday I would know.

Even proponents of SSI are embarrassed by these predictions, and try to explain them away. Hawthorne (2004: 162–5) invokes a kind of error theory, arguing that we tend to "project" our current standards to other knowers, times, and circumstances (for criticism, see MacFarlane (2005a: §3.2.2)). Stanley (2005: ch. 6) argues that contextualist views have similar bad consequences for modal embeddings, and that SSI can handle temporal embeddings (for criticism, see Blome-Tillman 2009). In the end, Hawthorne and Stanley concede that temporal and modal embeddings are problematic for SSI, but that SSI should be accepted anyway because the problems facing contextualism and standard invariantism are worse. But this argument is weaker if there is a genuine relativist alternative that avoids both the embedding problem and the standard problems with contextualism.

Relativism as Synthesis

Many people find something compelling in both the arguments against contextualism and the arguments against invariantism. The relativist account provides a third option, removing the need to choose between two unpalatable alternatives.

As we have seen, invariantism faces difficulties accounting for the apparent contextual variation in the standards one must meet to be counted as "knowing." Relativism accounts for this variation straightforwardly, since it takes the epistemic standard relevant for evaluating instances of "know" to be fixed contextually. But, unlike contextualism, it takes this standard to be fixed by features of the context of assessment, rather than the context of use. Because of this, it avoids the problem of "lost disagreement" faced by all forms of contextualism. Suppose Joe says "Moore knows that he has hands," and René says "No, Moore doesn't know that he has hands." If, as the relativist holds, the accuracy of these assertions depends on the standards relevant at the context of assessment, then, although *which* of them is accurate might vary from one perspective to another, from *no* perspective will it be possible for *both* assertions to be accurate. This helps explain why we take these assertions to express a disagreement, even when they are made in very different contexts.

From the relativist's point of view, invariantism and contextualism each capture part of the truth about knowledge attributions. Invariantism is right that there is a single knowledge relation, and that the accuracy of knowledge ascriptions does not depend on which epistemic standard is relevant at the context of *use*. But contextualism is right that the accuracy of such ascriptions depends somehow on contextually relevant standards. Relativism seeks to synthesize these insights into a more satisfactory picture.

Questions for the Relativist

Although relativism does not share the problems of invariantism and contextualism, it faces philosophical difficulties of its own. Here are some questions the relativist needs to answer. (Answers are not attempted here.)

- (1) It would be odd if "know" were the only expression for which a relativist semantics was appropriate. Indeed, the relativist semantics would appear ad hoc if "know" were its only target. Are there other expressions for which a relativist treatment is needed? How does "know" relate to them? (See MacFarlane 2003; Richard 2004; Lasersohn 2005; MacFarlane forthcoming.)
- (2) Assuming "know" has a relativist semantics, can anything be said about *why* an expression with the role of "know" should work this way, or is this just a brute fact?
- (3) Can the relativist really vindicate the intuitions of disagreement that proved difficult for the contextualist (and even the nonindexical contextualist)? What is required for disagreement, in general? (See MacFarlane 2007a.)
- (4) Are there any operators that shift the epistemic standards coordinate of circumstances of evaluation, as modal operators bind the world parameter? If not, how can we motivate positing this coordinate? (See Stanley 2005: ch. 7; MacFarlane 2009: §6.)
- (5) Can we really make sense of the idea that there is no absolute answer to the question whether a particular assertion is accurate, but only a perspective-relative one? (See MacFarlane 2005b.) Even if we can make sense of a relativist linguistic practice, could it be rational to engage in such a practice? (See Zimmerman 2007.)

Whether the relativist synthesis is really an improvement on invariantist and contextualist views will depend on whether these (and other) questions can be answered adequately.

References

- Blome-Tillman, M. (2009) "Contextualism, Subject-Sensitive Invariantism, and the Interaction of 'Knowledge'-Ascriptions with Modal and Temporal Operators," *Philosophy and Phenomenological Research* 79: 315–31.
- Brogaard, B. (2008) "In Defence of a Perspectival Semantics for 'Know'," *Australasian Journal of Philosophy* 86: 439–59.
- Cappelen, H. (2008) "Content Relativism and Semantic Blindness," in M. García-Carpintero and M. Kölbel (eds.), *Relative Truth*, Oxford: Oxford University Press, ch. 12, 265–86.
- Cohen, S. (1988) "How To Be a Fallibilist," *Philosophical Perspectives* 2: 91–123.
- Davis, W. (2007) "Knowledge Claims and Context: Loose Use," *Philosophical Studies* 132: 395–438.
- DeRose, K. (1996) "Relevant Alternatives and the Content of Knowledge Attributions," *Philosophy and Phenomenological Research* 56: 193–197.

- DeRose, K. (2004) "The Problem with Subject-Sensitive Invariantism," *Philosophy and Phenomenological Research* 68: 346–50.
- DeRose, K. (2005) "The Ordinary Language Basis for Contextualism and the New Invariantism," *Philosophical Quarterly* 55: 172–98.
- DeRose, K. (2006) "'Bamboozled by Our Own Words': Semantic Blindness and Some Arguments Against Contextualism," *Philosophy and Phenomenological Research* 73: 316–38.
- Feldman, R. (2001) "Skeptical Problems, Contextualist Solutions," *Philosophical Studies* 103: 61–85.
- Hawthorne, J. (2004) *Knowledge and Lotteries*, Oxford: Oxford University Press.
- King, J. (2003) "Tense, Modality, and Semantic Values," *Philosophical Perspectives* 17: 195–247.
- Kompa, N. (2002) "The Context Sensitivity of Knowledge Ascriptions," *Grazer-Philosophische Studien* 64: 1–18.
- Lasersohn, P. (2005) "Context Dependence, Disagreement, and Predicates of Personal Taste," *Linguistics and Philosophy* 28: 643–86.
- MacFarlane, J. (2003) "Future Contingents and Relative Truth," *Philosophical Quarterly* 53: 321–36.
- MacFarlane, J. (2005a) "The Assessment Sensitivity of Knowledge Attributions," in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology* 1: 197–233. Reprinted in E. Sosa, J. Kim, J. Fantl, and M. McGrath (eds.), *Epistemology: An Anthology*, second edition, Oxford: Blackwell, 2008.
- MacFarlane, J. (2005b), "Making Sense of Relative Truth," *Proceedings of the Aristotelian Society* 105: 321–39.
- MacFarlane, J. (2007a) "Relativism and Disagreement," *Philosophical Studies* 132: 17–31.
- MacFarlane, J. (2007b) "Semantic Minimalism and Nonindexical Contextualism," in G. Preyer and G. Peter (eds.), *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*, Oxford: Oxford University Press, 2007, 240–50.
- MacFarlane, J. (2009) "Nonindexical Contextualism," *Synthese* 166: 231–50.
- MacFarlane, J. (forthcoming) "Epistemic Modals are Assessment Sensitive," in E. Egan and B. Weatherson (eds.), *Epistemic Modality*, Oxford: Oxford University Press.
- Richard, M. (2004) "Contextualism and Relativism," *Philosophical Studies* 119: 215–42.
- Rosenberg, J. F. (2002) *Thinking About Knowing*, Oxford: Oxford University Press.
- Rysiew, P. (2001) "The Context-Sensitivity of Knowledge Attributions," *Noûs* 35: 477–514.
- Schaffer, J. (2004a) "Skepticism, Contextualism, and Discrimination," *Philosophy and Phenomenological Research* 69: 138–55.
- Schaffer, J. (2004b) "From Contextualism to Contrastivism," *Philosophical Studies* 119: 73–103.
- Schiffer, S. (1996) "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society* 96: 317–33.
- Stanley, J. (2004) "On the Linguistic Basis for Contextualism," *Philosophical Studies* 119: 119–46.
- Stanley, J. (2005) *Knowledge and Practical Interests*, Oxford: Oxford University Press.
- Zimmerman, A. (2007) "Against Relativism," *Philosophical Studies* 133: 313–348.

EPISTEMIC MODALS

Josh Dever

1. Initial Scene-Setting

Let's begin with some truisms. First truism: one of the many things that philosophers want is a theory of epistemology, laying out generalizations regarding epistemological notions such as knowledge, justification, and evidence, and predicting (for example) under what conditions various agents will know various things. Second truism: philosophers want not just any theory of epistemology, but the right theory. That means we need a method of selecting among the vast array of possible epistemological theories. Such a method might deploy any number of considerations, such as theoretical elegance and integration with neighboring areas of philosophy, but surely (third truism) a central feature of any plausible method is conformity to epistemological data, in the form of facts of the form *so-and-so knows/doesn't know such-and-such under these circumstances*, or *so-and-so has/doesn't have good reason for believing such-and-such under these circumstances*.

Epistemologically minded philosophers thus have reason to collect epistemological data. There are many ways to gain such data, but one central method is via judgments about sentences that contain epistemological vocabulary. So, in the ideal (or, perhaps, oversimplified) case, a piece of epistemological theorizing might proceed through the following stages:

1. The sentence "Jones knows that the cat is on the mat" is true.
2. Therefore, Jones knows that the cat is on the mat.
3. Epistemological theory *T* entails that Jones does not know that the cat is on the mat.
4. Therefore, *T* is not the right epistemological theory.

Hence, a fourth and final truism: those interested in epistemology have reason to be interested in the linguistic devices for reporting epistemological facts.

The four truisms present a pretty picture of (a certain aspect of) epistemological theorizing, but in the real philosophical world, things inevitably become more complicated. Here are two complicating factors. First, not all epistemological vocabulary is transparently epistemological. No one will miss the fact that constructions of the form "A knows that *p*" or "That *p* is a reason to believe that *q*" are important sources of data for epistemological theorizing, but other constructions might only with some coaxing reveal their epistemological relevance. Indicative conditionals, for example, are often taken to depend on the evidential state of speakers—consider a view on which the truth