Philosophical Issues, 16, Philosophy of Language, 2006

TWO PERSPECTIVES ON KNOWLEDGE OF LANGUAGE

Stephen Schiffer
New York University

1. The Two Perspectives

The philosopher:

> To know a language is to *use*, or be able to use, it—to know how to speak it and to be able to understand what others are saying when they speak it.[1] Now, a language is an abstract object that may or may not be used by anyone (think of Esperanto), so the task of saying in what knowledge of a language consists divides into two subtasks: first, the task of saying what a language is, and second, the task of saying what relation a person must bear to a language in order for that language to be a language the person uses, or is at least able to use.

The linguist:

> To know a language is for one's language processing to employ an internally represented generative grammar of the language, where that is a finitely axiomatizable theory of the language whose assignments of structures and meanings yield for every meaningful expression of the language a theorem that assigns to that expression its meaning, or meanings, in the language. The language one knows—one's I-language, or idiolect—is whatever language the internally represented generative grammar is a theory of, and to know that language is just for one's language processing to deploy that grammar. An account of knowledge of language is an account of the nature of the internally represented generative grammars we humans deploy in our use of language.

Either of these two perspectives on knowledge of language may involve problematic presuppositions, but the two perspectives aren't in conflict. They are concerned with different questions. The philosopher wants to know in what using a language consists; but even if he's able to tell us what it is to

use a language, a further question will remain. This is the question of what psycholinguistic theory explains how we're able to use the languages we use. Although these are two distinct perspectives on knowledge of language, there is a question about how the two perspectives are related, and this is one of the issues I will address.

2. The Language in Knowledge of Language

So-called natural languages such as English, Japanese, and Arabic should not be regarded as "languages" in the sense relevant to either the philosopher's or the linguist's question. At best, they are clusters of related languages in the sense of language which needs to be at issue. But what is the notion of language which needs to be at issue? Neither the philosopher nor the linguist is trying to identify some pre-theoretic notion of language which she then wants to set about deploying in her theory. Each theorist wants instead to define a notion that will earn its theoretical keep in explaining the linguistic phenomena the theorist really has it in mind to explain when operating under the rubric "knowledge of language." So what are these underlying theoretical issues? There is of course a very complex skein of inextricably related issues. Still, we may isolate two dominant issues that will motivate defining language in a certain way in order to theorize about those issues.

A dominant issue for the philosopher is to account for what it is for each of infinitely many sentences to mean something for a person. Sentences don't have meaning absolutely: a sentence has a particular meaning *in* a language, or *for* a person or population of persons. For anyone who knows a language, there are infinitely many sentences that have a particular meaning for that person, even though he has never uttered, heard, or entertained those sentences. For example, even though you probably never before encountered the sentence 'Rebus monkeys rarely rumba', it means for you that rebus monkeys rarely rumba. Understanding what it is for a sentence to mean such-and-such for a person is crucial to understanding many of the other things about language philosophers and linguists want to understand.

One of those other things is a dominant issue for linguists. This is to account for how we're able to understand utterances of sentences we've never before encountered. Someone utters the sounds 'Hey, mister, you spilled coffee on my Gucci bag' and you straightway know that the speaker said that you spilled coffee on her Gucci bag, even though you never heard that sentence before. The linguist wants to know how you are able to do that.

Both dominant issues suggest that the notion of language we may need is that of a relation that holds between infinitely many finite sequences of sounds or marks or whatever—things we'll want to call the sentences of

the language—and an infinite class of things that may be taken to be the meanings of sentences. It is reasonable to suppose that any usable language will have its sentences generated by a compositional syntax which meshes with some sort of compositional meaning theory, but it is theoretically prudent to begin by building as little as possible into the notion of a language, only then bringing in whatever else is needed, as it is needed, in order to explain what it is for a language, as minimally defined, to be a person's language, or in order to explain how a person who uses a language is able to understand utterances of novel sentences of the language.

Elsewhere I have argued that insofar as we can make sense of there being things that are sentence meanings, they are best construed as things that constrain two things: the kind of speech acts that may be performed by one who utters the sentence on its own (i.e. not as part of a larger sentence), and the contents those speech acts may have (Schiffer 2003, ch. 3). I have also argued elsewhere that the things we assert and believe are what philosophers call *propositions*—abstract, mind- and language-independent entities that have truth conditions, and have those truth conditions both essentially and absolutely, i.e. without relativization to anything else (Schiffer 2003). For example, if you believe that eating carrots improves eyesight, then you believe the proposition that eating carrots improves eyesight. This proposition—that eating carrots improves eyesight:

- is *abstract* in that, like the number 7, it has no physical characteristics;
- is *mind- and language-independent* in that it wasn't brought into existence by anything anyone said or thought;
- has a truth condition in that it's true iff eating carrots improves eyesight;
- has its truth condition essentially in that it's a necessary truth that the proposition that eating carrots improves eyesight is true iff eating carrots improves eyesight;
- and has its truth condition absolutely in that it doesn't have its truth condition only in a language or for someone, but has it everywhere and everywhen.

Even if the things we believe and assert are propositions, there remains a further question about the nature of those propositions. Are they functions from possible worlds onto truth-values, as Robert Stalnaker would have it, or are they structured entities of a certain kind, as both Russellians and Fregeans would have it, or are they, as I have argued in *The Things We Mean*, still another kind of proposition? This paper doesn't require me to take a stand on the particular nature of the propositions we assert and believe.

As I said, I take the meaning of a sentence to be something that constrains both the kind of speech act a speaker performs in uttering the sentence on its own and the kind of content such a speech act must have. If, as I've suggested, the contents of speech acts are propositions, then we may for present purposes

take sentence meanings to be ordered pairs of the form $<A, P>$, where A is a kind of speech act and P is a kind of proposition. On this construal, we might, to a rough approximation, take the meaning of 'Is it raining?' to be the ordered pair <asking-whether, a proposition of the form *it's raining at place m at time m′*, where *m* identifies a place implicitly referred to by the speaker and *m′* identifies the time of the utterance>. The meaning of a sub-sentential expression would then also be a kind of propositional content that partially determines the propositional-content component of the meanings of the sentences that contain the expression; and we should expect the meaning of a complex expression to be determined by its syntax and the meanings of its component expressions.

We now have a provisional characterization of language primed to be used by both the philosopher and the linguist in at least their initial attempts to achieve their respective accounts of knowledge of language. This is the characterization of language as a relation whose domain is a class of infinitely many finite strings of sounds or marks or whatever, and whose range is a class of infinitely many ordered pairs of the form $<A, P>$, where A is a kind of speech act and P is a kind of proposition. I say that a language is a *relation*, to allow for a sentence's having more than one meaning, but it will simplify matters if we ignore ambiguity and say that a language is a *function* whose domain of arguments consists of infinitely many finite sequences of sounds or marks or whatever—the "sentences" of the language—and whose range of values consists of infinitely many ordered pairs of the kind just described—the "meanings" of the language's sentences. This sort of characterization of language even has Chomsky's imprimatur: the "characterization of a 'language' as a pairing of sound and meaning over an infinite domain is traditional and reasonable as a point of departure" (Chomsky 1980, p. 82). If L is a language and $L(\mathtt{s}) = <A, P>$, then we may say that s is a sentence of L and that $<A, P>$ is the meaning of s in L. The fact that s means $<A, P>$ in L is a necessary truth that has nothing to do with how anyone uses s or the words composing it.

Equipped with this provisional definition of a language, let's now look at how the philosopher might answer her question.

3. The Philosopher's Question and the Meaning-without-Use Problem

A person knows a language, in the philosopher's sense, if she uses it, or is able to use it. So the philosopher's question may be restated thus: What relation must obtain between a person and a language in order for the language to be used by the person? Let's say that the *language relation* is that relation which must hold between a language and a person in order for the language to be used by the person. The philosopher's account of knowledge of language is her account of the language relation.

Here's a first shot at an account of the language relation inspired by David Lewis's seminal work on that issue ('L' ranges over languages and 'x' over persons):

> (I) L is used by x iff x belongs to a group of communicators G such that there is a practice in G of communicating by uttering sentences of L wherein if $L(\mathtt{s}) = <A, P>$ and a member y of G utters $\mathtt{s}$, then in uttering $\mathtt{s}y$ performs a speech act of kind A whose content is a proposition of kind P.

This may seem at first glance to provide at least a sufficient condition for using a language, since it secures that x has a practice of uttering L sentences in conformity with their meanings. But we can sense trouble for the claim that (I) states a sufficient condition when we realize that the conditional embedded in the right-hand side ('if $L(\mathtt{s}) = <A, P>$ and a member y of G utters $\mathtt{s}$, then . . .') will be vacuously true for those infinitely many instances of its antecedent that never get uttered in G. For suppose English is Jane's language, and let English$^+$ be the same as English as regards every sentence that anyone is ever likely to utter but departs radically from English thereafter. Perhaps 'giraffe' means the same as 'grapefruit' in every English$^+$ sentence in which 'giraffe' occurs more than 100 times. This is a counterexample to the alleged sufficiency of (I) because English$^+$ isn't a language that Jane or anyone else uses, yet English$^+$ satisfies (I)'s alleged sufficient condition. English satisfies that condition because the uttered English sentences non-vacuously satisfy it while the unuttered ones vacuously satisfy it; and English$^+$ satisfies it in exactly the same way, since every English$^+$ sentence uttered in G is also a sentence of English.

The foregoing is a version of a problem I raised years ago against the account of the language relation that Lewis offered in the penultimate ms draft of his book *Convention*. Lewis subsequently dubbed the problem the *meaning-without-use problem*: it's the problem of accounting for how our use of language determines what language we use when all but a finite minority of the sentences have no chance of ever being used. This problem is appropriately called the meaning-without-use problem because if $\mathtt{s}$ means m in L and L is x's language, then $\mathtt{s}$ means m for x. Consequently, the problem of saying what it is for x to use L is the same as the problem of saying what it is for all of L's infinitely many sentences to mean for x what they mean in L. If L is x's language, then all but a finite minority of the infinitely many sentences of L have meaning-without-use for x. Whatever it is that secures that a particular language is your language also secures that most of its sentences have meaning-without-use for you. Whatever that securing fact is, it must surely be some actually obtaining, finitely statable fact. To solve the meaning-without-use problem, and thereby to say what makes it the case that a person x uses a language L, is to find some finitely statable condition pertaining to x's use of L whose satisfaction nails down all of L at once. What is this nailing-down condition?

After disowning his 1969 and 1975 attempts to solve the meaning-without-use problem, David Lewis proclaimed in 1992 that the meaning-without-use problem

> has an obvious solution: extrapolation. First, use somehow determines meaning for the fragment of the language that is actually used. These are rules of syntax and semantics that generate the right sentences with the right meanings within the used fragment. These rules also generate other, longer sentences, with meanings, outside the used fragment. Use determines some meanings, those meanings determine the rules, and the rules determine the rest of the meanings. Thus use determines meaning, in part directly and in part indirectly, for the entire language. (Lewis 1992, p. 149.)

But if that's the obvious solution to the meaning-without-use problem, then why did it take Lewis twenty-three years to find it? He says it's because he was

> scared off it by Kripkenstein's challenge (formerly Goodman's challenge). As follows: the used fragment does not determine the rules. There are many different systems of rules—different grammars—that yield just the same sentences with just the same meanings inside the used fragment, but that differ wildly when they go beyond it. Extrapolation, which means going on according to the same rules, is radically underdetermined. (Ibid., p. 150.)

So why the change of heart? Because, Lewis explains:

> [I] should not have been scared off. The obvious solution is right. True, there are many grammars. But they are not on equal terms. Some are 'straight' grammars; for example, any grammar that any linguist would actually propose. Others are 'bent', or 'gruesome', grammars; for example, what you get by starting with a straight grammar for English and adding one extra rule, which states that every expression with more than forty occurrences of the word 'cabbage' is a sentence meaning that God is great. . . . The notion of extrapolation presupposes the distinction between straight and bent. It means going on according to the same straight rules. (Ibid., pp. 150-1.)

Applied to the first shot (I), Lewis's proposal yields this revision:

(II) L is used by x iff x belongs to a group of communicators G such that there is a practice in G of communicating by uttering sentences of L wherein

(a) if $L(\mathtt{s}) = <A, P>$ and a member y of G utters $\mathtt{s}$, then in uttering $\mathtt{s}$ y performs a speech act of kind A whose content is a proposition of kind P; and

(b) the only straight grammar that fits the fragment of L used in G is a grammar of L.

One striking thing about the added condition (b)—Lewis's "obvious solution" to the meaning-without-use problem—is that it doesn't require the grammar that determines what language a person uses to play any role at all in the psycholinguistic explanation of the information processing that underlies, and thus accounts for, the person's ability to understand utterances in her language.

The striking fact invites an objection to Lewis's solution, one I made to him when in 1990 he showed me a draft of the paper from which I've been quoting. The objection is simply this. If we learned that the internally represented grammar implicated in the linguistic behavior of a certain population was in fact a bent grammar that determined a language L', then L' would be the language of that population, even if the only straight grammar that fit the used fragment determined a different language. Lewis considers my objection in a footnote:

> Maybe there is a grammar written into the brain. And conceivably it is a bent grammar, so that the language it generates differs, somewhere outside the used fragment, from the language we get by straight extrapolation. Schiffer has asked: does straight extrapolation give the right answers even then? I think so. If not, then whenever we resort to extrapolation to answer questions of syntax and semantics, we are engaged in risky speculation about the secret workings of the brain. That seems wrong. (Ibid., fn. 6, p. 151.)

I'm not convinced by Lewis's reply. We can see its inadequacy by appeal to two examples that satisfy the counterexample schema he attributes to me:

> Suppose that the infinite language L is in fact the language of group G. Let L_{frag} be the fragment of L that has actually been produced in G during, say, the past ten years (thereby securing that L_{frag} is "large and suitably varied"). We may assume that each member of G has a compositional understanding of L, one that somehow relies on a straight grammar of L, and that every straight grammar that generates L_{frag} is a grammar of L. Now it is certainly possible that there should be another group H which also uses L_{frag}, but with this difference: each sentence of L_{frag} is for them a *non-composite* utterance type. In other words, each sentence of L_{frag} is for them as a simple signal, such as a fire alarm, is for us: it has propositional meaning, but its meaning is not in any way determined by semantic features of its parts and structure. The members of H have prodigious memories, and they have learned the sentences of L_{frag} one by one. They know what their sentences mean because they have learned what each one means as a single fact, and they have no way of computing what a sentence means on the basis of their knowledge of its syntax and the meanings of its parts. As far as they're concerned, a sentence no more has a syntax and semantically relevant parts than a fire alarm has for us. Consequently, members of H have no way of understanding a sentence that belongs to L but not to L_{frag}; they have no way, in fact, of determining the meaning of any novel sentence. Clearly, the infinite

> language L is *not* used by the members of H. Yet every straight grammar that generates L_{frag}, the language that *is* used by them, is a grammar of L.[2]

> Suppose that the members of a secret society make up a language whose syntax and semantics they actually write down and formally adopt under a Mafia-like oath. Perhaps for some reason, perverse or not, the grammar they adopt is by their explicit design a bent grammar. Let's also suppose that an internal representation of the bent grammar plays an essential role in their processing of utterances in their invented language. It seems clear to me that the language of the society is the one their bent grammar describes, whether or not there's a straight grammar that fits the sentences the members of the society actually produce or are likely to produce.

These two counterexamples are examples where the language used is used in accordance with a "bent" grammar (the non-compositeness of H's language is one version of a bent grammar). But we don't need to rely on intuitions about recherché examples. There is an argument based on knowledge of meaning which also tells against the Lewisian solution to the meaning-without-use problem.

4. The Philosopher's Question and Knowledge of Meaning

Every speaker knows what numerous sentences mean. But what is it to know what a sentence means?

Knowing what a sentence means isn't having the propositional knowledge that it means such-and-such. While your Uncle Clyde knows what 'Is she there yet?' means, there is no completion of

> Clyde knows that 'Is she there yet?' means . . .

that is both true and entails that Clyde knows what 'Is she there yet?' means. Nor does there seem to be any other sort of proposition Clyde knows which entails that he knows what that sentence means. After all, four-year-old children know what 'Is she there yet?' means, yet it seems clear that there is no proposition p such that their knowing p entails that they know what that sentence means.

A popular answer to the knowledge of meaning question, whose advocates include Wittgenstein and J. L. Austin, agrees that knowledge of meaning isn't any kind of *propositional* knowledge, isn't a matter of knowing *that* such-and-such is the case, but is instead a kind of *knowing how*, a kind of ability or skill. This is the idea to which Wittgenstein was alluding when he advised not to ask for an expression's meaning, but for its use. But what kind of know-how is knowledge of meaning? A reasonable conjecture as regards sentence meaning might be some precisification of the rough idea that

(KH) To know what s means in *L* is to have the ability to know what a speaker of *L* would be saying were she to utter s.

For several reasons, (KH) is very much a first approximation, but among the problems it confronts there is one that makes it doubtful that there's enough to (KH) to make it worth trying to fix.

The problem to which I allude is simply this. We are able to know what a speaker is saying in uttering a sentence in part *because we know what the uttered sentence means*; but that couldn't be so if knowing the meaning of a sentence was being able to know what would be said in an utterance of it. A man confronts you in Central Park, and says to you, 'Madame, your daughter just bit my dog'. Even though you never before encountered that sentence, you knew straightway that in uttering it the man was telling you, and thus saying, that your daughter bit the man's dog. How were you able to move from your perception of the man's utterance to your knowledge of what he said? It's reasonable to suppose that part of the answer is that *you knew what the uttered sentence meant*. This couldn't be part of the explanation of how you arrived at your knowledge of what the man was saying if your knowing what the sentence meant just was your ability to know what would be said in an utterance of it. This problem with (KH) points the way to a better account of what it is to know what a sentence means.

Let's stay with the irate man in the park. He produces an acoustical blast—his utterance of 'Madame, your daughter just bit my dog'—and you immediately know that the man is stating that your daughter just bit his dog. Here we have a bit of information processing that begins with your auditory perception of his utterance and ends with your belief about what he said. Implicated in this process is an information-processing state that constitutes your knowing what the uttered sentence means, a state that plays the knowledge-of-meaning information-processing role crucial to your understanding of the utterance. This suggests that:

(K-IP) To know what s means in *L* is to be in (or disposed readily to be in) a token of a state-type capable of occupying the knowledge-of-meaning role in your understanding of a literal *L* utterance of s.

But what is the nature of the state that plays the knowledge-of-meaning information-processing role?

It's not a state that realizes any sort of propositional knowledge. If it were we should expect that one arrives at one's belief about what was said on the basis of an *inference* whose premises include the proposition that the uttered sentence means such-and-such, or has some other meaning-constituting property. But the information process that takes us from an auditory perception of the utterance of a sentence to a belief about what

was said in that utterance does not appear to be any sort of *inference*. It seems much more like the sort of non-inferential processing at work when one has a visual experience and on that basis believes that there is a red ball in front of one. So what, then, might be the nature of the state that plays the knowledge-of-meaning role?

The meaning of a sentence, I have proposed, is an ordered pair of the form $<A, P>$, where A is a kind of speech act and P is a kind of propositional content. If s means $<A, P>$ for S and S utters s on its own and is speaking literally, then, for some speech-act type Ψ of kind A and proposition q of kind P, S Ψs q in uttering in s, and your understanding S's utterance of s consists in your knowing that in uttering s S Ψ-ed q. It's therefore reasonable to suppose that the information process that takes you to your knowledge that S Ψ-ed q contains a state that directly or indirectly represents s as linked with its meaning, $<A, P>$. *Whatever state represents that linkage is the one that plays the knowledge-of-meaning role*.

No doubt more than one kind of state can play the knowledge-of-meaning role. For example, a monolingual speaker of Norwegian might have been told that 'Eating frogs is morally permissible' is equivalent in meaning to a certain Norwegian sentence, and that propositional knowledge will be the state that plays the knowledge-of-meaning role when he understands an utterance of the English sentence. But in the normal case of one who understands the utterance of a novel sentence of her own language, nothing like that is going on. In the first place, as already noted, the state that plays the knowledge-of-meaning role won't be a state that realizes propositional knowledge about the uttered sentence, and, in the second place, the state that represents a sentence as linked with its meaning will itself by determined by other states in the information-processing sequence that effect linkages of the sentence's words with their meanings.

Suppose, as many linguists and philosophers believe, that implicated in the information process that takes one from one's perception of the utterance of a sentence of language L to one's knowledge of what was said in that utterance is an internal representation of what the linguists call a generative grammar of L and what the philosophers call a compositional meaning theory for L. Suppose such a theory is internally represented. Then here's what might be going on when ordinary guy Al understands his wife Sue's utterance of 'He refuses to eat it'.

Al, like each of us, is a computational information-processing device, and his processing of English utterances involves both an internally represented generative syntax of English and an internally represented compositional meaning theory for English. When Sue uttered 'He refuses to eat it', her noise making caused Al's ear drums to vibrate in certain ways, thereby causing certain electrical signals to be sent via Al's hearing nerves to his auditory input processor. That processor delivers a representation of the sounds produced by Sue as input to Al's language processor. There the representation is

transformed in various ways and a representation is produced that is the input to the language sub-processor that contains the representation of a syntax for English. The output of that processor is then the input to the language sub-processor that contains the representation of a compositional meaning theory. That processor in turn computes (inter alia) representations of the words in Sue's utterance as paired with their meanings and uses those representations to compute a representation of the sentence Sue uttered as paired with its meaning (perhaps the content of that representation is an ordered pair whose first member is the sentence and whose second member is the sentence's meaning, an ordered pair of the form $<A, P>$). That representation is the output of the language processor and the input to certain central-processing mechanisms, and a few milliseconds later Al knows that Sue has told him that their dog Fido refuses to eat the $200 Kobe steak they bought for him.

The partial account, or caricature, of what might be going on in the information processing involved in Al's understanding of Sue's utterance is of course a partial gloss, or caricature, of the sort of account the Chomsky-inspired linguist expects to be her account of knowledge of language. If everything is OK so far, we are now well positioned to have the philosopher use the linguist's conception of knowledge of language to solve his, the philosopher's, knowledge-without-meaning problem, and thereby arrive at the philosopher's account of knowledge of language.

5. A Symbiotic Answer to the Philosopher's Knowledge-of-Language Question

The philosopher answers her knowledge-of-language question by saying what it is for any given language to be used by any given person. To do this the philosopher must solve her meaning-without-use problem. I suggested that (K-IP) is a plausible account of what it is to know what a sentence means. I then used the kind of account the linguist would give in answer to her knowledge-of-language question to suggest that the state that typically plays the knowledge-of-meaning role is a state that is generated in part by a generative grammar and which represents the uttered sentence as linked with its meaning. Now, if x uses L, then every sentence of L means for x what it means in L. It doesn't follow from a sentence's having meaning for a person that the person knows what the sentence means. The sentence might be too long or too convoluted for the person even to entertain. But we can use the Chomsky-inspired account of knowledge of meaning to solve the philosopher's meaning-without-use problem, and thereby to say in what knowledge of language, in the philosopher's sense, consists. Or at least we can solve it by way of stating a condition whose satisfaction not only entails that L is used by x but is also a condition (or close to one) that is actually

satisfied, and thus a condition that pertains to how we actually understand utterances in the languages we know.

My suggestion is this:

> *L* is used by *x*—and thus a language *x* knows—if *x* belongs to a group of communicators *G* who regularly communicate by uttering sentences of *L*, and when they do, the knowledge-of-meaning role is played by states that pair expressions with their meanings via an internally represented generative grammar for *L*.

So my solution to the meaning-without-use problem is this. What makes *L* a language that is used by *x* even though the used fragment of *L* is also a fragment of infinitely many languages other than *L* is that it's the language determined by the internally represented generative grammar implicated in *x*'s language-understanding processes. That's in effect the same response Chomsky made to Lewis's account of the language relation. In disagreeing with Lewis, he said that what makes *L* a language that is used by a population *P* is that "each person in *P* has a grammar determining *L* in his mind/brain" (ibid., p. 84). Since that tells us what it is to use a particular language, it also tells us what it is to know a language, in that sense of knowing a language in which you ipso facto know the languages you use.

But what about David Lewis's objection that if the grammar that determines the used language is the one written into the brain, then "whenever we resort to extrapolation to answer questions of syntax and semantics, we are engaged in risky speculation about the secret workings of the brain . . . [and] that seems wrong"? The sort of "extrapolation" to which Lewis alludes concerns hypotheses about the rules of syntax and semantics that actually govern the language we actually use. Suppose I believe that the correct generative grammar for my language is whatever grammar is written into my brain. I don't need to tell readers of this journal that it's no small feat to come up with a grammar that merely fits the fragment of my language that I'm likely to use. Suppose, however, that after considerable labor I believe that a certain grammar Γ gets that fragment right. Now, Γ determines not just the language fragment in question but an entire infinite language: it assigns an LF and meaning to each of infinitely many sentences, all but a small minority of which have no chance of ever being used by me or anyone else. Since Γ fits the pretty large fragment on which I can reliably pronounce, it's a pretty good bet that it's a straight grammar, probably the only straight grammar I can come up with that does fit that fragment. So now I must speculate: is Γ the grammar I actually use in my language processing, or is that grammar perhaps one that coincides with Γ on the relevant fragment but thereafter recognizes as its only sentence the string

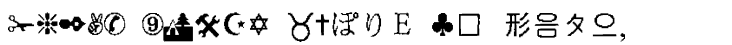

or is it that grammar which is like Γ but assigns to every sentence with more than 100 occurrences of 'cabbage' the proposition that God is great as its meaning, or is it still some other bent grammar? Call me a risky speculator, but I'll go with the straight grammar Γ. I'll go with it for the same reason that a scientist will infer a straight theory from the evidence that underdetermines it, rather than infer one of the infinitely many bent theories that fits the same evidence.

Notes

1. For present purposes, 'using a language' pertains only to using a language as a public language of communication, and not, say, as a system of mental representation, a *lingua mentis*.
2. This example is recycled from Schiffer 1993, pp. 236–237.

References

Chomsky, N. (1980). *Rules and Representations* (Columbia University Press).
Lewis, D. (1969). *Convention* (Harvard University Press).
———. (1992). "Meaning without Use: Reply to Hawthorne," *Australasian Journal of Philosophy*, 70: 106–10.
Schiffer, S. (1993). "Actual-Language Relations," *Philosophical Perspectives*, 7: 231–58.
———. (2003). *The Things We Mean* (Oxford University Press).