

MEANING WITHOUT USE: REPLY TO HAWTHORNE

David Lewis

Surely it is our use of language that somehow determines meaning. But if we try to say how, we must face the fact that only a tiny part of our language, or any human language, is ever used. There are many reasons why a meaningful sentence might never be suited to serve anyone's conversational purposes, and so might go unused. For instance, take length. Even the most abominable stylist will never write a sentence more than, say, a hundred words long. (Never? — Well, hardly ever.) But almost all of the infinitely many meaningful sentences of English, all but a finite minority, are longer than a hundred words. Almost all are longer than a thousand words, almost all are longer than a million words . . . So almost all sentences have meaning without use.

Years ago, Stephen Schiffer raised meaning-without-use problem against my own account of use and meaning, which ran as follows. A *language* L is a function that assigns truth conditions to certain verbal expressions, called the *sentences* of L; one of them is *true in L* or *false in L* according as the truth condition assigned to it by L is satisfied or not. To be *truthful in L* is to avoid uttering any sentence of L unless it is true in L; to be *trusting in L* is to expect others to be truthful in L. We (or any population) *use* L iff, by convention, we are truthful and trusting in L. (A convention is a regularity in the population that is sustained, in a certain way, by expectation of others' conformity. The details need not concern us here.) The *meaning* of a sentence, for us, is the truth condition assigned to that sentence by the language we use. That, I said, is how use determines meaning.¹

Now for the problem. If a sentence of L is never uttered at all, *a fortiori* it is never uttered falsely in L; if it is expected never to be uttered at all, *a fortiori* it is expected never to be uttered falsely in L. So for the unused part of L — which is almost all of L — truthfulness and trust in L go trivial. And this is so whether or not L gets the truth conditions right. Let L₁ be our actual language, the one that does get the truth conditions right. Let L₂ be a different language that agrees with L₁ on all the sentences we ever use: all we might ever utter, and all we do not expect others never to utter. On the unused sentences, however, L₁ and L₂ may

¹ David Lewis, 'Languages and Language', *Minnesota Studies in the Philosophy of Science* 7 (1975) pp. 3-35; reprinted in David Lewis, *Philosophical Papers*, Vol. I (Oxford: Oxford University Press, 1983). A generally similar account appeared in David Lewis, *Convention* (Cambridge, MA: Harvard University Press, 1969; Oxford: Blackwell, 1986); however, that account lacked the appeal to 'trust in L' that figures in the present discussion. Here I simplify the story by ignoring indexicality, ambiguity, non-indicative moods, sub-sentential linguistic expressions, and hyperintensional differences of meaning.

differ wildly. They may assign conflicting truth conditions to the same unused sentences, or they may disagree about which unused expressions are sentences at all. Yet we are truthful and trusting in L_1 and L_2 alike: trivially so in the differing unused parts, non-trivially and conventionally so in the used part that L_1 and L_2 have in common. It is not perfectly clear how my analysis classifies a case where truthfulness and trust are partly trivial and only partly conventional. Maybe we use both L_1 and L_2 , maybe we use neither. But, either way, L_1 and L_2 are on a par. Our truthfulness and trust do not distinguish them. So I have not explained why the correct meanings of unused sentences, and the correct ascription of sentencehood to unused expressions, are those given by L_1 rather than L_2 .

In reply to Schiffer I proposed this solution. Let us say that trust in L is more than just the expectation that sentences of L will not be uttered falsely in L . Say rather that one is *trusting in* L with respect to sentence S to the extent that one's conditional subjective probability that S will be uttered, given that S is true in L , exceeds one's unconditional probability that S will be uttered. Equivalently, to the extent that the 'likelihood ratio'

$$\frac{\text{Probability (S is uttered / S is true in L)}}{\text{Probability (S is uttered / S is false in L)}}$$

which measures the extent to which the truth of S in L is confirmed when S is uttered, exceeds one-to-one. Even if all these probabilities are minute — even if they are literally infinitesimal — still comparisons between them make sense. Trust thus defined does not go trivial just because one expects S never to be uttered, so long as the probability of utterance is not quite zero. Even for the unused sentences, we have one remaining vestige of use: non-trivial, conventional, trust in L , consisting of a difference between minute probabilities. It is this vestigial trust, so I said, that distinguishes our actual language L_1 from the impostor L_2 .

John Hawthorne now revives the issue.² He says that my solution fails because the subjective probability of utterance of the unused sentences is exactly zero. Then even the redefined version of trust in L goes trivial. The compared probabilities are all zero, and the likelihood ratio goes undefined, regardless of whether L gets the truth conditions right or wrong. (He needn't say we have zero probabilities always, for every unused sentence and every member of the population. But if we have them sometimes, that is bad enough.) I do not agree. I suppose I am typical among language-users; and although I find it improbable that a hundred-word sentence will be uttered, I find it more improbable that a thousand-word sentence will be uttered, and still more improbable that a million-word sentence will be uttered, and so on *ad infinitum*. This would make no sense if, after a point, the probabilities were zero exactly. But it does make sense. So they are not zero; just very, very small.

Why think otherwise? One bad reason — not given by Hawthorne, I hasten to add — is based on an operational definition of subjective probability. What would you pay for a bet whereby you win a rich reward if a certain very long

² John Hawthorne, 'A Note on Languages and Language', *Australasian Journal of Philosophy* 68 (1990) pp. 116-118.

sentence is ever uttered? (And uttered not for the sake of winning the bet.) Exactly nothing, if the sentence is long enough. But that needn't mean that the probability is zero. It just means that there is no coin so worthless as not to exceed the minute expected value of the bet. Or perhaps it means that the bet is not worth the trifling effort it would take to buy it. Other operational tests are more revealing. If forced to choose, would you not take the bet which you win if a certain thousand-word sentence is uttered instead of the bet which you win if a certain million-word sentence (without countervailing advantages) is uttered? But would you not take the second bet instead of a bet on some still worse sentence?

Anyhow, subjective probability should not be tied by definition to particular operational tests. Rather, it is a theoretical concept, defined implicitly by its role in decision theory. This theory contains constitutive constraints of rationality, and one aspect of rationality is a modicum of open-mindedness whereby genuine possibilities do not get probability zero. (Infinitesimal, maybe, when infinitely many alternatives seem more or less equiprobable. But infinitesimal isn't zero.) Utterance of a very long sentence is a genuine possibility, in the sense that matters. If the sentence is long enough, its utterance may be a physical impossibility, to be sure; but mere physical impossibilities are not entitled to probability zero, since we cannot know the laws of nature with absolute certainty.

So I disagree with Hawthorne: trust in L, taken in the sense of likelihood ratios, does not go trivial. But I disagree still more with my earlier self. It's all very well to define trust in terms of likelihood ratios, but doing so does not solve the meaning-without-use problem. To that extent, Hawthorne is right.

Consider some very long sentence. Let it be not only long but complicated: clauses within clauses within phrases within clauses . . . , and abundantly interlaced with cross references to 'the latter', 'the former', 'the aforementioned', 'condition (b*)', and so on *ad nauseam*. Of course you don't expect to hear this sentence uttered. The subjective probability is minute. But what if you did hear it? Would you think this was a successful job of truth-in-L-telling? Not likely! You'd think the speaker was trying to win a bet or set a record, or feigning madness or raving for real, or doing it to annoy, or filibustering, or making an experiment to test the limits of what it is humanly possible to say and mean. You wouldn't think he was even trying to be truthful in L. Still less would you think he was trying effectively, armed with skill enough to overcome the complexities of the sentence. In short, the lion's share of your subjective probability would go to hypotheses under which the utterance of the sentence had little to do with whether it was true in L. And likewise now, when you haven't heard the sentence uttered, the lion's share of your minute subjective probability that it will be uttered goes to these same hypotheses. Whether the sentence is uttered is probabilistically independent, near enough, of whether it is true in L. The likelihood ratio is one-to-one, or near enough. It does not go undefined. We can tell whether you are trusting in L. But the likely answer, for a great many unused sentences, is that you are not.³

³ It won't help much to restrict attention to utterances in 'serious communication situations', a

So if, as before, L_1 gets the truth conditions of unused sentences right and L_2 gets them wrong, we cannot hope to distinguish L_1 from L_2 as the language in which we are truthful and trusting. They will be on a par. And the same will be true, for similar reasons, if L_2 differs from L_1 about which expressions are sentences at all. The meaning-without-use problem is with us still.

But it has an obvious solution: extrapolation. First, use somehow determines meaning for the fragment of the language that is actually used. There 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.⁴

Hawthorne has not overlooked the obvious solution, of course, and neither did my earlier self. But we were both scared off it by Kripkenstein's challenge (formerly Goodman's challenge).⁵ 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.

We 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.

³ *continued . . .*

move that I consider at one point in 'Languages and Language'. That sets aside the cases of filibustering, etc., but leaves the case of an ineffective attempt at truthful communication.

⁴ Once we have the rules, we may use them to correct or refine our assignment of truth conditions within the used fragment, if we do so with enough restraint not to overthrow the basis from which we extrapolated the rules. *Correction* (here I am indebted to John Hawthorne): suppose the extrapolated rules assign truth condition T to sentence S; and suppose it is common knowledge that nobody has any way to tell whether T obtains; or that everybody already knows that T obtains; or that everybody already knows that T does not obtain. If the rules are right, the truth condition for S cannot be manifest in a pattern of truthfulness and trust, even if S is in the used fragment. If there is any such pattern it should be disregarded. S should be assigned its truth condition indirectly, in virtue of the extrapolated rules. *Refinement* (here I am indebted to M.J. Cresswell): suppose two languages L_1 and L_2 have the same sentences and differ only slightly, as follows, in the truth conditions they assign. Wherever T_1 and T_2 are truth conditions assigned to the same sentence by the two languages, T_1 and T_2 will coincide except in some case that all those in the population take to be quite improbable. (Example: the population consists of devout theists; S means in L_1 that Max likes cookies; S means in L_2 that God knows that Max likes cookies; and similarly for other sentences.) Then it may happen that the population is, by convention, truthful and trusting both in L_1 and L_2 to (near enough) equal degrees. But use of the extrapolated rules might favour L_1 over L_2 . (Parsing S in accordance with these rules, we find words that mean 'Max', 'likes', and 'cookies', but no word that means 'God' or 'knows'.

⁵ Ludwig Wittgenstein, *Philosophical Investigations* (Oxford: Blackwell, 1958), especially sections 185-242; Saul Kripke, *Wittgenstein on Rules and Private Language* (Cambridge, MA: Harvard University Press, 1982); Nelson Goodman, *Fact, Fiction, and Forecast* (Cambridge, MA: Harvard University Press, 1955), ch. II.

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. We have no difficulty in telling the difference. (Except insofar as bentness admits of degree, and some grammars are only a little bit bent. But this is a complication we can ignore.) We can reasonably hope that all straight grammars that agree on the used fragment will agree everywhere. We have no ironclad guarantee of this, but also no cause for alarm. After all, the used fragment is large and varied. The wild differences are between straight and bent grammars. The notion of extrapolation presupposes the distinction between straight and bent. It means going on according to the same straight rules. It is not radically underdetermined. We can speak of extrapolation with a clear conscience.⁶

The lesson of Kripkenstein is not that extrapolation is an illegitimate notion. What is illegitimate, rather, is a simple-minded analysis of extrapolation, one that does without the bent-straight distinction by overlooking bent rules altogether. Something else that is illegitimate, at least for those of us who have not embraced circularity, is an overall plan of analysis that postpones the bent-straight distinction to a later chapter and yet uses it in an early chapter. If we must rely on the bent-straight distinction to reach an analysis of meaning, we may not afterward analyse straightness in terms of straightforward (short, simple, non-disjunctive) expressibility in our language. Likewise, *mutatis mutandis*, if we must rely on the bent-straight distinction still earlier to reach an analysis of the content of thought. If that means carrying more baggage of primitive distinctions or ontological commitments than some of us might have hoped, so be it.⁷ But thinking that Kripkenstein proscribes talk of extrapolation is like thinking that Zeno stops us going from place to place.

Princeton University,

Received March 1991

⁶ Maybe there is a grammar somehow 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.

⁷ For catalogues of baggage available to be carried, see my 'New Work for a Theory of Universals', *Australasian Journal of Philosophy* 61 (1983) pp. 343-377; and D.M. Armstrong, *Universals: An Opinionated Introduction* (Boulder: Westview Press, 1989).