

TWO-DIMENSIONALISM AND THE FOUNDATION OF LINGUISTIC ANALYSIS

Andrew Melnyk

Can two-dimensional semantics provide a foundation for linguistic analysis, i.e., the conceptual analysis of words in a natural language? In what follows, I make a case for skepticism.

I understand linguistic analysis to be the use of a certain method—the method of hypothetical cases—to gain a priori knowledge of certain necessary truths. The method consists of asking oneself, for a variety of hypothetical situations, whether one would apply a given word to something in that hypothetical situation. The necessary truths that this method (together with further reflection) supposedly enables one to know a priori are necessary truths that can be formulated by using the given word, e.g., necessary truths that express necessary and sufficient conditions for the correct application of the word.

To illustrate, consider how (according to advocates of linguistic analysis) one might come to know that, necessarily, being an unmarried man is a necessary and sufficient condition of being a bachelor. Suppose that, for all of the (sufficiently varied) hypothetical situations one considers, one would say that ‘bachelor’ applies to *all* the unmarried men in that situation. That suggests that being an unmarried man is a *sufficient* condition for being a bachelor. And suppose that, for all of the (sufficiently varied) hypothetical situations one considers, one would say that ‘bachelor’ *doesn’t* apply to all people who *fail* to be unmarried men in that situation. That suggests that being an unmarried man is also a *necessary* condition of being a bachelor.

Now if one comes to know that bachelors are unmarried men through linguistic analysis, one must have understood the term, ‘bachelor’. And whatever exactly is the process by which one comes to understand words, it surely requires perceptual experience of the world. So understanding ‘bachelor’ requires that one have had perceptual experience of the world. But it does *not* follow that the knowledge that, necessarily, being an unmarried man is a necessary and sufficient condition of being a bachelor would be a posteriori if acquired through linguistic analysis. Such knowledge would still count as a priori if, to gain it, one would require no *further* perceptual experience than is required for coming to understand ‘bachelor’ in the first place. And advocates of linguistic analysis claim, of course, that this requirement is met. The method of hypothetical cases can only be used by someone who understands ‘bachelor’; but if one *does* understand ‘bachelor’, then one needs no *further* perceptual experience either to use the method of hypothetical cases or to draw a conclusion from the results of doing so. And ‘further’ perceptual experience includes two things: not only perceptual experience had *after* one comes to understand ‘bachelor’, but also perceptual experience, had *before* one comes

to understand ‘bachelor’, that goes beyond what is required for coming to understand ‘bachelor’. Intuitively, the knowledge that, necessarily, being an unmarried man is a necessary and sufficient condition for being a bachelor would count as a priori if acquired through linguistic analysis because merely understanding ‘bachelor’ is sufficient (together with certain capacities for a priori reasoning) for one to acquire the knowledge.¹

Linguistic analysis is the use of a certain method to achieve a certain goal. But it is appropriate to use a method to achieve a goal only if the method is *capable* of achieving the goal. So linguistic analysis is appropriate only if the method of hypothetical cases is capable of generating a priori knowledge of necessary truths. But how it does so cannot be magic. So linguistic analysis is appropriate only if there is some explanation of *how* the method of hypothetical cases is capable of yielding a priori knowledge of necessary truths, an explanation in terms of what goes on in our minds when we use the method. It seems inevitable that any such explanation will construe the necessary truths discovered by the method of hypothetical cases as arising somehow from the meanings of words.² The explanation may then be able to explain how one acquires knowledge of necessary truths naturalistically, but still preserve the a priori character of this knowledge, by not requiring knowers to have had any *more* perceptual experience than they needed to come to understand certain words. The appropriateness of linguistic analysis therefore turns on whether such an explanation of how the method of hypothetical cases can yield a priori knowledge of necessary truths—what I will call a ‘foundation’ for linguistic analysis—exists.

In section 1, I present what I call the ‘classical foundation’ for linguistic analysis. In section 2, for a reason that will emerge in section 4, I present a semantically externalist view of the meaning and reference of general terms, according to which no foundation for linguistic analysis is possible.³ In section 3, I present a two-dimensional account of the same terrain, focusing exclusively on Frank Jackson’s version of two-dimensionalism, and sketch how it can, it seems, provide a foundation for linguistic analysis.⁴ Finally, in section 4, I suggest a reason why the two-dimensional account might not in fact be able to provide such a foundation.

1 The Classical Foundation for Linguistic Analysis

What I am calling the ‘classical foundation’ for linguistic analysis is intended to be a reconstruction of what practitioners of linguistic analysis during its heyday would have offered as a rationale for their practice if pressed to give one.⁵ It presupposes a cluster of inter-related theses about meaning, reference, and linguistic understanding that are nowadays very widely rejected, though initially attractive, and it will provide a useful foil for views discussed later.

The central thesis that the classical foundation presupposes is that the *meaning* of a *general* term (e.g., ‘gold’, ‘bachelor’) is a particular *set of properties* that is conventionally associated with the term. In one sense of the word ‘concept’, such a set of properties can be said to be a concept, because the set of properties specifies *what it is to be F* (or maybe what it *would* be to be F), where ‘F’ is the general term conventionally associated with the set of properties; and what it is (or would be) to be F is the conceptually necessary and sufficient condition for being F.⁶ It is concepts in *this* sense that linguistic analysis analyzes; hence the traditional name for linguistic analysis, ‘conceptual analysis’.⁷ For example, ‘gold’ might be conventionally associated with the properties of being shiny, yellow, and malleable, so that ‘Gold is highly prized’ would mean that the shiny, yellow, and malleable stuff is highly prized; and being shiny, yellow, and malleable is what it is (or would be) to be gold, the conceptually necessary and sufficient condition for being gold. Likewise, the empty but still meaningful term, ‘unicorn’, might be conventionally associated with the properties of being a horse and having a single horn in the center of the forehead.⁸

A corollary of the central thesis that is crucial to grounding linguistic analysis is that *understanding* a general term, ‘F’, requires knowing, in some fashion, what it is to be F—requires knowing, in some fashion, that ‘F’ applies to something iff it has the properties that are conventionally associated with

'F'. Thus, understanding 'gold' requires knowing, in some fashion, that 'gold' applies to something iff it is shiny, yellow, and malleable. But understanding 'F' does not require the knower to verbalize the knowledge; a perfectly competent speaker-hearer of 'F' could go a lifetime without doing so. Understanding 'F' is merely *implicit* knowledge of the meaning of 'F', and manifests itself chiefly in one's competence as a speaker-hearer.

The central thesis about meaning and its corollary about understanding naturally suggests an account of how the method of hypothetical cases manages to yield a priori knowledge of necessary truths. The necessary truths that can be known a priori are, of course, the conceptually necessary truths that specify the necessary and sufficient condition for a general term to apply to something—for something to be F if the term is 'F'. And, according to the corollary, understanding a general term gives one implicit knowledge of the conceptually necessary and sufficient condition for the term to apply to something. Let us now add the plausible claim that one is drawing upon this implicit knowledge not only when, as a competent speaker-hearer, one considers *actual* circumstances and judges whether or not a general term applies to something in those circumstances but also when one considers *hypothetical* circumstances and judges whether or not a general term *would* apply to something in those circumstances. Because one is drawing upon knowledge, one's judgments are correct. And on the basis of sufficiently many such judgments about a general term, one can formulate an explicit statement of the conceptually necessary and sufficient condition for the term to apply to something. Knowledge of this statement counts as a priori because gaining it requires no perceptual experience of the world beyond what is required for coming to understand the term.

2 The Externalist Challenge to Linguistic Analysis

If the theses about meaning and linguistic understanding presupposed by the classical foundation for linguistic analysis are true, then linguistic analysis seems to have a solid foundation. On the other hand, linguistic analysis lacks a foundation if a semantically externalist view of the same terrain is true instead. Let me sketch a generic form of such an externalist view, without attributing it to any particular author(s), and then spell out its consequences for linguistic analysis.

According to this form of externalism, a general term, 'F', does not have a meaning of the sort envisaged by the classical view: there is no set of properties conventionally associated with 'F' such that to be F is to have the properties in the set.⁹ Two important consequences follow. First, the reference of 'F' cannot be determined by its meaning; its reference cannot be the thing, if any, that has the properties in the set of properties conventionally associated with 'F'. Second, understanding a general term cannot consist in implicitly knowing that 'F' applies to something iff it has the properties that are conventionally associated with 'F'. Given these consequences, this form of externalism must provide alternative accounts of how the reference of a general term is determined, and of what understanding a general term consists in. For present purposes, however, detailed accounts are not needed, and so I will merely sketch.

According to an externalist account of reference-determination, a general term refers to the real thing to which it stands in a certain non-semantic and natural relation, perhaps, but not necessarily, a causal-historical relation;¹⁰ standing in this relation to a real thing is what referring *is*.¹¹ The account doesn't claim, however, that the term is *conventionally associated with* (the property of) standing in this relation to a certain real thing. It also doesn't claim that understanding the term requires *knowing* that the term applies to whatever stands in this relation to a certain real thing. For example, on Saul Kripke's externalist 'picture' of reference, the term 'gold' on my lips refers to gold because it stands at the end of a lengthy historical chain of reference-borrowing, passing through countless users of the term, that got started when someone successfully introduced the term 'gold' as a term for a certain stuff which was in fact gold (Kripke 1980, 93–97).¹² But Kripke's picture does not claim that there is a conventional association among English-speakers between 'gold' and (the property of) standing at the end of this chain of reference-borrowing; and it does not claim that speaker-hearers who

understand ‘gold’ need to know, even implicitly, that ‘gold’ applies to whatever stands at the end of this chain.¹³ There is no standard externalist account of understanding a general term. But one externalist proposal is that understanding ‘F’ merely requires the ability to translate back and forth between occurrences of ‘F’ in English sentences and tokenings of a co-referential concept in thoughts with the same truth-conditional contents as the sentences (see Millikan 1984, 147–148; Devitt and Sterelny 1999, 187–190). Such an ability need not require propositional knowledge of meaning. Because inaccurate, or less accurate, habits of translation would tend to produce less success in interacting with other people and with the physical world, the ability might be acquired by exercising comparatively unsophisticated psychological capacities to abandon less successful behavior in favor of more successful behavior (e.g., the capacity to be conditioned).

If the externalist view of meaning and understanding that I have sketched is true, then the method of hypothetical cases seems incapable of yielding a priori knowledge of necessary truths. For on this view there *are* no conceptually necessary truths, which would require meaning of the sort envisaged by the classical view; so obviously the method of hypothetical cases cannot yield a priori knowledge of *them*.¹⁴ Consistently with this view there may, of course, be *metaphysically* necessary truths to the effect that something is F iff it is the G (e.g., that something is gold iff it is the element with atomic number 79). But understanding ‘F’, on the externalist view, does not require knowledge, not even implicit knowledge, of metaphysically necessary truths about being F; such knowledge would have to be acquired a posteriori (as in the case of gold). And there appears to be no further kind of necessary truth that might lead to a vindication of the claim that the method of hypothetical cases yields a priori knowledge of necessary truths.

3 A Two-Dimensional Foundation for Linguistic Analysis?

Let me now explain how two-dimensional semantics—or at least my reconstruction of Frank Jackson’s well-known version of it—might provide a foundation for linguistic analysis, starting with what it has to say about meaning, reference, and linguistic understanding. I limit myself to Jackson’s well-known version of two-dimensionalism, and only insofar as it applies to general terms, because its theoretical goals are reasonably clear and explicitly include providing a foundation for linguistic analysis (Jackson 1998, 2001, and in this volume; Chalmers and Jackson 2001). By ‘two-dimensionalism’ I will henceforth mean my reconstruction of Jackson’s two-dimensionalism as it applies to general terms.

Two-dimensionalism agrees with the classical account in identifying the meaning of a term, ‘F’, with a particular set of properties that is conventionally associated with ‘F’. But there is a major difference between how the two accounts view the set of properties. On the classical account (section 1), the set of properties conventionally associated with ‘F’ plays *two* roles: first, it specifies what it is to be F; second, it determines the reference of ‘F’—the reference of ‘F’ is the thing, if any, that has the properties in the set. On two-dimensionalism, the conventional association between ‘F’ and certain properties is along these lines: (i) ‘F’ is the term that people conventionally use when they want to talk about the thing that *actually* has the associated properties and (ii) when people hear other people utter ‘F’, they conventionally take them to be talking about the thing that *actually* has those properties.¹⁵ Because the unique bearer of the properties in the *actual* world may not bear those properties in *other* possible worlds, the set of properties conventionally associated with ‘F’ does not play the first role. But it still plays the second, reference-determining role: the reference of ‘F’ is the *actual* bearer of the properties in the set.¹⁶ Thus, the conventionally associated properties for ‘water’ might include falling from the sky as rain and filling the lakes and rivers, in which case the reference of ‘water’ would be the *actual* stuff that falls from the sky as rain, fills the lakes and rivers, and so forth; in some possible worlds, however, that actual stuff—H₂O, as it turns out—does *not* fall from the sky as rain, fill the lakes and rivers, and so forth. Even though, on two-dimensionalism, the conventionally associated properties do not specify what it is to be F, so that there is no *conceptually* necessary and sufficient condition

of being F, there can still be a *metaphysically* necessary and sufficient condition of being F. However, this metaphysically necessary and sufficient condition is not entailed by the meaning of 'F', i.e., by the properties conventionally associated with 'F'.

Understanding 'F', on the two-dimensionalist view, requires knowing the meaning of 'F'. But because this meaning does not specify a conceptually (or metaphysically) necessary and sufficient condition of being F, understanding 'F' does not require knowledge of such a condition; here, of course, two-dimensionalism breaks with the classical view and agrees with externalism. Understanding 'F' is also claimed not to require the ability to *state* which properties fix the reference of 'F'. Understanding 'F' is claimed to be merely *implicit* knowledge of which properties fix the reference of 'F', taking the form of an ability correctly to judge as true or false (or perhaps as indeterminate) every 'F'-involving 'application conditional' (Chalmers and Jackson 2001), that is, every conditional of the form, 'If the actual world turns out to be thus and so, then "F" applies to such-and-such' (e.g., 'If the actual world turns out to be one in which D₂O [deuterium oxide] is the stuff that falls from the sky as rain, fills the lakes and rivers, and so forth, then "water" refers to D₂O'). Because, on the two-dimensionalist view, understanding 'F' does not require implicit knowledge of a conceptually (or metaphysically) necessary and sufficient condition of being F, two-dimensionalism cannot explain how the method of hypothetical cases manages to yield a priori knowledge of conceptually (or metaphysically) necessary truths. But it need not try to do that, for it claims that there is another kind of necessary truth, and it seeks to explain how the method of hypothetical cases manages to yield a priori knowledge of necessary truths of *that* kind.

Necessary truth is traditionally understood as truth in all possible worlds. But two-dimensionalism draws an important distinction between possible worlds 'considered as counterfactual'—ways the world might have been—and possible worlds 'considered as actual'—ways the actual world might turn out to be. Possible worlds 'considered as counterfactual' are possible worlds as traditionally construed. Intuitively, 'What if your spouse had been a space alien?' is a very different question from 'What if your spouse turns out to be a space alien?' Both questions concern the same possibility; but the first question considers this possibility as a way the world might have been while the second question considers it as a way the world might turn out to be. To envisage the possibility that one's spouse should *turn out* to be a space alien is to envisage the possibility that one is dramatically mistaken as to how the world has actually been all along.

The meaning of 'F', on two-dimensionalism, can be represented formally as a function from each possible world *considered as actual* to the referent of 'F' (if any) in that world; this function Jackson calls the 'A-intension' of 'F'. The A-intension therefore takes each possible world considered as actual onto the thing (if any) in that world that has the properties conventionally associated with 'F', and it will typically take different possible worlds considered as actual onto different referents.¹⁷ Perhaps the A-intension of 'water' takes the possibility that D₂O turns out to be the stuff that falls from the sky as rain and so forth onto D₂O, but takes the possibility that alcohol turns out to be the stuff that falls from the sky as rain and so forth onto alcohol. The reference of 'F' can also be represented formally, as a constant function from each possible world *considered as counterfactual* to the thing that, in the *actual* world, has all the properties in the set conventionally associated with 'F'; this function is what Jackson calls the 'C-intension' of 'F'.¹⁸ Thus, on the assumption that H₂O is the actual stuff that falls from the sky as rain, fills the lakes and rivers, and so forth, the reference of 'water' can be represented formally as a constant function from each possible world considered as counterfactual to H₂O.

The distinction between possible worlds considered as counterfactual and possible worlds considered as actual enables a distinction to be drawn between two kinds of necessary truth. The first kind of necessary truth, often called 'metaphysically' or 'broadly logically' necessary, is truth in all possible worlds considered as counterfactual, and may be called 'C-necessary'. The second kind of necessary truth, which has no traditional name, is truth in all possible worlds considered as actual, and may be called 'A-necessary'. For example, the sentence, 'Water falls from the sky as rain', is not C-necessary, because had the world been too cold for water ever to evaporate then water would not

have fallen from the sky as rain. But if the A-intension of ‘water’ can be paraphrased as ‘the actual stuff that falls from the sky as rain and fills the lakes and rivers’, then the sentence *is* A-necessary—true in all worlds considered as actual—because *whatever* turns out to fall from the sky as rain and fill the lakes and rivers, whether H₂O, D₂O, alcohol, or something even better, it falls from the sky as rain.

It is now easy to see how two-dimensionalism might explain how the method of hypothetical cases can yield a priori knowledge of necessary truths. The first move is to claim that the necessary truths in question are *A-necessary truths*.¹⁹ Admittedly, A-necessary truths are not conceptually necessary truths. So linguistic analysis of ‘F’ as vindicated by two-dimensionalism would not produce a priori knowledge of what it is to be F; it would not achieve the longstanding philosophical goal of discovering the nature or essence of things. Jackson claims that it would nonetheless play an indispensable role in solving what he calls ‘location problems’ in philosophy (Jackson 1998). The second move is to appeal to the two-dimensionalist view that understanding ‘F’ requires implicit knowledge of the A-intension of ‘F’: the ability to evaluate (judge as true or false or indeterminate) every ‘F’-involving application conditional correctly, i.e., in conformity with the actual A-intension of ‘F’. This view promises to secure the a priori character of our knowledge of A-necessary truths. For suppose that one has the ability correctly to evaluate every ‘F’-involving application conditional. Then whenever, using the method of hypothetical cases, one considers whether ‘F’ applies to something in a hypothetical situation, one can simply treat the hypothetical situation as a way the actual world might turn out to be, and then exercise the ability. And having answered the question about the applicability of ‘F’ with regard to sufficiently many hypothetical situations considered as actual, and having reflected on the situations’ commonalities, one can rightly conclude (without the help of further perceptual experience) that something is F iff it is the actual bearer of certain properties. If this conclusion is true, it seems to constitute knowledge, and indeed a priori knowledge, because reaching it requires no perceptual experience beyond what was required to acquire understanding of ‘F’ in the first place.

4 A Problem for the Two-Dimensional Foundation

The two-dimensional account of how the method of hypothetical cases can yield a priori knowledge faces a problem. It takes no official stand on the nature of knowledge itself. But it would presumably be served best by some version of a reliabilist theory of knowledge, so that the products of linguistic analysis count as knowledge because they are produced by what is in fact a reliable belief-forming process, even if the subject has no evidence that it is reliable. But a reliabilist theory of knowledge must, in my view, include a ‘no defeater’ condition in order to avoid certain counterexamples (see Lycan 1988, 109–110). A true belief that arises from the operation of a reliable belief-forming process is not knowledge if the subject has a defeater, either a rebutting defeater (evidence that the belief is false) or an undermining defeater (evidence that the belief-forming process is not reliable); subjects know only if they have no defeater. I will argue that, even if the two-dimensionalist account of linguistic analysis is true, practitioners of linguistic analysis *who reflect on the account* have an undermining defeater for the belief-forming process that is claimed to operate in linguistic analysis. The defeater is the fact that, given the available evidence, the two-dimensionalist account of linguistic understanding is not clearly more probable than an externalist account on which, as we saw in section 2, linguistic analysis does *not* constitute a reliable belief-forming process. One can, of course, practice linguistic analysis without reflecting on the two-dimensionalist (or any) account of it; but philosophers who practice linguistic analysis are likely to reflect on it.

An undermining defeater for a belief-forming process need not show that the belief-forming process is *unlikely* to produce true beliefs. For example, suppose my teenage son tells his mother that he has done all his homework. Testimony, of course, is in general a reliable belief-forming process. But my son has often told his mother that he has done all his homework when in fact he hasn’t, and

she knows this. Also, the present episode seems no different from earlier episodes. She therefore has an undermining defeater, so that, even if in an excess of charity she believes him and he has told the truth, she does not *know* that he has done all his homework. But the prior occasions on which his testimony on homework has been inaccurate need not have formed a *majority* of all prior occasions on which he provided testimony on homework; all his mother knows is that his testimony on homework has *often* been inaccurate. I will not attempt further exploration of what it takes to be an undermining defeater. I will argue that philosophers who reflect on the two-dimensionalist account of linguistic analysis have an undermining defeater at least as strong, intuitively, as the one my wife has in knowing that my son's testimony has often been inaccurate. These philosophers should therefore judge that the products of their linguistic analysis do not constitute knowledge; at the very least they should withdraw any claim that they do constitute knowledge. In principle, the defeater that these philosophers have could itself be defeated, but only by the discovery of further empirical evidence. So, even if linguistic analysis can in the end produce knowledge, it would still not be *a priori* knowledge.

Suppose, then, that one practices linguistic analysis, and that the two-dimensionalist account of linguistic analysis is true. On reflection, one will notice that, according to the two-dimensionalist account, engaging in linguistic analysis constitutes a reliable belief-forming process only if understanding a general term, 'F', requires that one can evaluate all 'F'-relevant application conditionals correctly, i.e., in conformity with the actual A-intension of 'F'. Next, one will naturally wonder whether understanding 'F' really *does* require this particular ability. If it does, that must be because (i) understanding 'F' is implicitly knowing that 'F' means so-and-so, and (ii) what 'F' means is a certain particular A-intension, so that (iii) understanding 'F' is implicitly knowing that 'F' means this particular A-intension, i.e., is being able to evaluate all 'F'-relevant application conditionals in conformity with this particular A-intension.²⁰ According to two-dimensionalism, however, for the meaning of 'F' to be a certain particular A-intension (and not another one) just is for 'F' to be conventionally associated, in the right reference-fixing sort of way, with a certain set of properties (and not another set).²¹ But it is a posteriori whether 'F' is conventionally associated, in the right reference-fixing sort of way, with one set of properties rather than another; and so it is *also* a posteriori whether understanding 'F' really does require that one can evaluate all 'F'-relevant application conditionals in conformity with the *actual* A-intension of 'F' (and not with some other A-intension).

All that being so, what would actual empirical evidence for the view that understanding 'F' requires this particular ability look like? However the view is arrived at, it trivially implies the universal psychological generalization that everyone who understands 'F' has the ability to evaluate all 'F'-relevant application conditionals in conformity with the actual A-intension of 'F'. The view therefore predicts that numerous universal psychological generalizations hold among people who understand general terms in English. For instance, the view predicts that everyone who understands 'water' has the ability to evaluate all 'water'-relevant application conditionals in conformity with the actual A-intension of 'water', and hence (since there is only one A-intension of 'water') that everyone who understands 'water' would give *the same answers as everyone else* to every question of the form, 'If the actual world turns out to be thus and so, does "water" apply to such-and-such?'

Does the available empirical evidence suggest that such universal psychological generalizations actually hold? Presumably no one has ever discovered evidence for them as a social scientist would, by identifying a representative sample of people who understand (say) 'water', quizzing them with a representative sample of the right sort of questions, and then observing whether their answers exhibit uniformity. A different approach would exploit the fact that two-dimensionalists grant that the A-intension of a general term can often be paraphrased by a rigidified definite description, at least roughly. Suppose 'water' is such a term. Then the hypothesis that everyone who understands 'water' can evaluate all 'water'-relevant application conditionals in conformity with the A-intension of 'water' predicts that everyone who understands 'water' will think, more or less explicitly, that water is the actual G, where 'G' roughly paraphrases the A-intension of 'water'. And in principle there could be evidence that this prediction is correct. Thinking that water is the actual G does not, of course,

entail being disposed to use ‘the actual G’, rather than any other rigidified definite description true of water, to evaluate ‘water’-relevant application conditionals;²² but evidence that *any* prediction of a hypothesis is correct is some evidence for the hypothesis. Presumably, however, no one has ever identified a representative sample of people who understand ‘water’, asked them whether they think water is the actual G, and found that they all do. We must therefore fall back on casual empiricism. But when we do, I will now argue, such evidence as we find does not clearly favor, and may not favor at all, the conclusion that those who understand a general term, ‘F’, all think that F is the actual G, for some predicate, ‘G’.

First, one must have done more than hear someone use a term a few times to determine with reasonable confidence that the speaker thinks that a certain definite description is true of the term’s referent. How confident do you feel about which rigidified definite descriptions your friends and family think are true of the referents of ‘butter’ or ‘soil’ or ‘ingenuity’, terms that you must have heard them use scores of times? Either one must have examined a fair bit of speakers’ term-using speech or one must actually question speakers; and one is not likely to have met either condition with regard to more than a few speakers. One’s casual observation, then, has probably only revealed small numbers of people who understand ‘F’ and all think that the term’s referent is the actual G, for some ‘G’. Such small samples constitute weak evidence.

Second, however many confirming instances one’s casual observation reveals, they are unlikely to form a representative sample. One is most unlikely to have selected speakers so that they form a representative sample: they are just people one has run into through one’s job or reading habits or hobbies. And probably one has positive reason to fear that the sample is *unrepresentative* (e.g., if the speakers are highly educated, reflective, and articulate people from the distinctive sub-culture of academic analytic philosophy). One should hesitate to extrapolate far beyond it.

Third, positive instances of a generalization confirm the generalization only if there are no counterexamples. On the face of it, however, there are many counterexamples to the generalizations in question: people who understand a general term but who appear not to think that *any* definite description is true of the term’s referent.²³ The common philosophical example of ‘water’ may give a misleading impression, because all humans directly experience water every day in similar ways. But consider the better-known terms for chemical elements. Many of us count ourselves as understanding, say, ‘caesium’, even though all we know about caesium is that it is a highly reactive metallic element, and we know that caesium is not the *only* highly reactive metallic element. Many of us count ourselves as understanding ‘argon’ and ‘krypton’, even though all we know about argon and krypton is that they are noble gases, though not the only ones.²⁴ Such speakers, and others like them, are *prima facie* counterexamples to the hypothesis that, for every general term ‘F’, those who understand ‘F’ all think that F is the actual G, for some predicate, ‘G’.

Such speakers, however, are no threat to an *externalist* account of understanding a general term, the chief rival to the two-dimensionalist account. Suppose the externalist account mooted above is true: that understanding a general term merely requires the ability to translate back and forth between occurrences of the term in English sentences and tokenings of a co-referential concept in thoughts with the same truth-conditional contents as the sentences. Even if what explains *how* speakers can translate back and forth in this way is sometimes their thinking that a rigidified definite description is true of the term’s referent, it needn’t be the *same* rigidified definite description for *everyone* who understands the term; it needn’t even be the same rigidified definite description for a single speaker at different times in their life. Such an externalist account is therefore not committed to such psychological generalizations as that everyone who understands ‘water’ thinks, more or less explicitly, that water is the actual G, where ‘G’ roughly paraphrases the A-intension of ‘water’. At the same time, an externalist account need not deny that sometimes there *is* a single definite description that all people who understand a term think is true of the term’s referent. The account is therefore consistent with the evidence for the two-dimensionalist account that we have reviewed.

Let us now take stock. Those who reflect on the two-dimensionalist account of linguistic analysis have no systematic evidence for the entailed account of understanding a general term. They may have some unsystematic but weak evidence. There are, however, *prima facie* counterexamples to the account, cases that a rival externalist account of understanding a general term naturally accommodates. What to conclude? One might reasonably conclude that the externalist account wins by a whisker. Perhaps one might reasonably conclude that the two-dimensionalist account wins by a whisker. What one may not reasonably conclude is that on balance the evidence *clearly* favors the two-dimensionalist account; there is a good chance that the account is false, and that its externalist rival is true. Therefore, even if the two-dimensionalist account of linguistic analysis is in fact true, practitioners of linguistic analysis who reflect on the account along the preceding lines have an undermining defeater for the belief-forming process that, according to the account, operates in linguistic analysis. Further empirical evidence might defeat this defeater, and then the products of linguistic analysis would constitute knowledge (if the two-dimensionalist account of linguistic analysis is true); but they would not constitute a priori knowledge.

Even if the two-dimensionalist account of understanding a general term is not true, we still have dispositions to return confident answers to questions of the form, 'If the actual world turns out to be thus and so, does "F" apply to such-and-such?' The *extent* to which we have such dispositions—for how many general terms and for how many possibilities considered as actual—is not obvious from casual observation. But most of my students, for example, can be persuaded that, if what falls from the sky as rain and fills the lakes and rivers turns out to be D_2O , rather than H_2O , then 'water' refers to D_2O . Something must explain their dispositions to judge in this way, but if it is not their ability to evaluate all 'water'-relevant application conditionals in conformity with the A-intension of 'water', then what is it? The question would be a good one for psychological research. I speculate that their dispositions arise, not from their first-order semantic competence, but from their tacit endorsement of a *proto-theory of reference-determination*—which may be no more than the vague idea that words refer to those things in the world that we have dealings with when we use the words. Their (tacit) reasoning would then be: if what falls from the sky as rain and fills the lakes and rivers turns out to be D_2O , then our dealings with what we have been calling 'water' have in fact all along been dealings with D_2O , and so 'water' refers to D_2O . But this vague idea cannot *require* two-dimensionalism, because one way of developing it would eventuate in Millikan's radically externalist account of reference-determination according to which, very roughly, tokens of 'water' refer to the stuff that these tokens must correspond with for them to be able to perform their proper function of helping to produce true beliefs (Millikan 1984, ch. 6).

Notes

- 1 Does acquisition of knowledge count as a priori if it relies on a belief that is *innate*? On the one hand, though acquiring the belief does not require the *knower* to have had perceptual experience, it may require that the knower's *ancestors* had perceptual experience. On the other hand, a priori knowledge may not have a definition sharp enough to permit a definitive answer to the question.
- 2 The meanings of words, despite being partly constituted by mental phenomena, are mind-independent, just as minds themselves are mind-independent: their existence is at most *causally* dependent on *thought about them*. If the necessary truths allegedly discovered by the method of hypothetical cases arise from the meanings of words, then they are mind-independent too.
- 3 I treat only *general terms*; no one thinks we should do linguistic analysis on names. Only a portion of Jackson's exposition of two-dimensionalism in the present volume concerns general terms.
- 4 I focus exclusively on Jacksonian two-dimensionalism because the view is clearly intended as an account of public language—which cannot be said of, for example, David Chalmers' two-dimensionalism (Chalmers 2006; Schroeter 2017).
- 5 A sophisticated version of it has been defended more recently, for example, by Christopher Peacocke (1993).
- 6 The necessity is just truth in all possible worlds; I call it 'conceptual' to suggest that it is owed (somehow) to the nature of concepts.
- 7 In what I take to be a very *different* sense of 'concept', a concept is a kind of mental representation—perhaps a word in the language of thought, as in Jerry Fodor's theory (Fodor 1998).

- 8 A corollary of this central thesis (though it plays no part in grounding linguistic analysis) is that the *reference* of a general term, 'F', is the thing, if any, that has all the properties in the set of properties conventionally associated with 'F'—the thing, if any, that meets the conceptually necessary and sufficient condition for being F.
- 9 Or such that to be F merely *requires* having the properties in the set. A weaker form of externalism could allow that there is a set of properties conventionally associated with 'F' such that, if something is F, then it has the properties in the set, so that, for example, 'tiger' applies to something only if it is an animal. But the considerations invoked to support externalism as against a classical view tend to militate also against the weaker externalist view. For example, I once heard Berent Enç say that, early in their acquaintance, Fred Dretske nearly succeeded in persuading him that a buckeye was an animal. Enç was semantically competent with 'buckeye', for he had heard of buckeyes, knew, for instance, that they are common in Ohio (which is The Buckeye State!), and was weakly inclined to think that they are plants. He would not have lost that competence had Dretske succeeded. So we do not take semantic competence with 'buckeye' to require believing that buckeyes are plants. (Buckeyes are a kind of tree, though the term 'buckeye' is also used to refer to the tree's seeds, which resemble the eyes of deer.)
- 10 Ruth Millikan's externalist view doesn't take this relation to be causal-historical (Millikan 1984, 2005, 2017).
- 11 So empty terms do not refer. They may, however, be viewed as the products of processes that are *supposed* to produce non-empty terms, but that sometimes fail to do so; and they may play many of the same functional roles as non-empty terms (Millikan 1984, 136).
- 12 The theory originated by Kripke has been developed, for example, by Philip Kitcher (1993) and by Michael Devitt and Kim Sterelny (1999).
- 13 Proponents of causal descriptivism, of course, make both these claims.
- 14 The scope of the generic form of externalism that I have been presenting includes all general terms, even 'bachelor'. A version restricted to natural-kind terms would entail the weaker claim that there are no conceptually necessary truths about natural kinds. Stephen Biggs and Ranpal Dosanjh argue that externalism is true of many general terms, not just natural-kind terms (this volume).
- 15 Compare Jackson, in this volume, section 5: "the obvious [answer] is that 'water' is a word for telling about the watery stuff—that's its informational role".
- 16 In Kripke's terminology, the properties in the set are said merely to *fix the reference* (rather than, more ambitiously, to *give the meaning*) of 'F' (Kripke 1980).
- 17 The exceptions are when the properties are necessary properties; for example, $2=\sqrt{4}$ in all possible worlds, so 'the square root of 4' refers to 2 in all worlds considered as actual.
- 18 The reference is a constant function because 'F' is a rigid designator, referring, in every possible world in which it refers to anything, to whatever it refers to in the actual world. 'F' is rigid, because the definite description that fixes its reference is a *rigidified* definite description, 'the actual G'.
- 19 And that they are in principle expressible in a public language—because, on Jackson's two-dimensionalism, the A-intensions that ground A-necessary truths are the A-intensions of public terms. Versions of two-dimensionalism on which A-intensions (or primary intensions) belong to, say, terms in individual idiolects rather than to public terms will not have this consequence.
- 20 The meaning in question is what I have been calling 'meaning', which is distinct from reference.
- 21 The fact that a term is conventionally associated, in the right reference-fixing sort of way, with a certain set of properties *makes it the case* that the term has the A-intension that it has; but it does so only because the term's having the A-intension that it has *just is* the term's being conventionally associated, in the right reference-fixing sort of way, with a certain set of properties. One could conceivably claim instead that the semantic fact merely *emerges* from the social-cum-psychological fact, so that the two facts are distinct; but that claim is clearly not Jackson's.
- 22 To believe that water is the actual G is not yet to treat 'the actual G' as fixing the reference of 'water'.
- 23 Putnam famously made this point with his example of 'elm' and 'beech' (Putnam 1975). Kripke made the analogous point for proper names with his example of 'Richard Feynman' and 'Murray Gell-Mann'; both Feynman and Gell-Mann are known to some people who understand these names only as 'a famous American physicist' (Kripke 1980). The point is elaborated by Ruth Millikan (1984, ch. 9).
- 24 I know that argon is something called 'argon', but not that it is the only thing called 'argon'. (In fact, I know a person called 'Argon'.)

Further Reading

- Byrne, Alex, and James Pryor. 'Bad Intensions'. In *Two-Dimensional Semantics*, M. García-Carpintero and J. Macià (eds.), New York: Oxford University Press, 2006.
- Chalmers, David, and Frank Jackson. 'Conceptual Analysis and Reductive Explanation', *The Philosophical Review*, 110, 2001: 315–360.

- Jackson, Frank. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*, New York: Oxford University Press, 1998.
- Laurence, Stephen, and Eric Margolis. 'Concepts and Conceptual Analysis', *Philosophy and Phenomenological Research*, 67, 2003: 253–282.
- Millikan, Ruth Garrett. *Language, Thought, and Other Biological Categories: New Foundations for Realism*, Cambridge, MA: MIT Press, 1984.
- Millikan, Ruth Garrett. *Language: A Biological Model*, New York: Oxford University Press, 2005.

References

- Chalmers, David. 2006. 'Two-Dimensional Semantics'. In *The Oxford Handbook of Philosophy of Language*, E. Lepore and B. Smith (eds.). Oxford: Oxford University Press, pp. 575–606.
- Chalmers, David, and Frank Jackson. 2001. 'Conceptual Analysis and Reductive Explanation', *The Philosophical Review*, 110, 315–360.
- Devitt, Michael, and Kim Sterelny. 1999. *Language and Reality: An Introduction to the Philosophy of Language*. Cambridge, MA: MIT Press.
- Fodor, Jerry. 1998. *Concepts: Where Cognitive Science Went Wrong*. New York: Oxford University Press.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. New York: Oxford University Press.
- Jackson, Frank. 2001. 'Précis of *From Metaphysics to Ethics*', *Philosophy and Phenomenological Research*, 62, 617–624.
- Kitcher, Philip. 1993. *The Advancement of Science*. New York: Oxford University Press.
- Kripke, Saul. 1980. *Naming and Necessity*. Cambridge, MA: Harvard University Press.
- Lycan, William. 1988. *Judgement and Justification*. Cambridge: Cambridge University Press.
- Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Millikan, Ruth. 2005. *Language: A Biological Model*. New York: Oxford University Press.
- Millikan, Ruth. 2017. *Beyond Concepts: Unicepts, Language, and Natural Information*. New York: Oxford University Press.
- Peacocke, Christopher. 1993 'How Are a Priori Truths Possible?', *European Journal of Philosophy*, 1, 175–199.
- Putnam, Hilary. 1975. 'The Meaning of "Meaning"'. In *Mind, Language, and Reality: Philosophical Papers*, Vol. 2. Cambridge: Cambridge University Press.
- Schroeter, Laura. 2017. 'Two-Dimensional Semantics', *The Stanford Encyclopedia of Philosophy* (Summer 2017 Edition), Edward N. Zalta (ed.), <https://plato.stanford.edu/archives/sum2017/entries/two-dimensional-semantics/>.