

The Illusion of Semantic Reference

Christopher Gauker

University of Cincinnati

Version of December 2013

When I look at a chair, I know that the word “chair” applies to it. When I look at anything that is not a chair, I know that the word “chair” does not apply to it. In general, I feel, I know what kind of thing the word “chair” applies to. I know the extension of the word “chair”. So it seems clear to me that the word “chair” has a definite extension. There may be some borderline cases, but from that I draw no radical conclusions; all that means is that the boundaries of the extension are a bit fuzzy. And yet, when we philosophers try to say in a general way what it takes for a set to be the extension of a word, all of our theories seem to fail. They fail so badly that some of us have given up on the very idea that there is such a thing as a term’s having an extension. What is going on?

Similarly, we usually think we know quite well what a singular term, such as a proper name or demonstrative expression, refers to. If I say, “Obama is making another compromise with Senate Republicans,” then I am quite sure that the name “Obama”, as I use it, refers to the man who happens to be the President of the United States. Or if I say, “That one looks better,” I may have no doubt that I am referring to the fish filet that the grocer has his finger on at the time of my speaking. And yet when we try to explain in a general way what relation holds between a singular term and its referent, we find ourselves thwarted at every turn. Why is this so hard?

The question can be reformulated in terms of concepts in place of words. Some will say that a word such as “chair” has an extension only insofar as it *expresses* a mental representation—the concept *chair* (see references below). So if there is a problem, it must pertain in the first instance to the concept, not the word. In terms of concepts, we may formulate our question thus: When I look at a chair, I know that it is a chair; I know that it is correct to *think* that it is a chair. I think of it *as* a chair. When I look at a nonchair, I know that it is *not* correct to think that it is a chair. I do *not* think of it as a chair. So it seems quite plain that my concept *chair* has a definite extension (give or take a few borderline cases). And yet we have been unable to define in a general way the relation between a concept and its extension. Why?

What is going on, I will contend, is that we are misled by our ordinary practices of attributing knowledge of meanings. When the intuition that we know the extensions of our terms is examined closely, it gives way to a deeper intuition that we know the meanings of our terms. But our intuition that we know the meanings of our terms can be accounted for in a way that does not entail that we are acquainted with a semantic reference relation or any property that we have tried and failed to explicate. The account will go by way of an account of the role that attributions of knowledge of meaning play in the conduct of conversation.

1. *Semantic reference*

We have an intuition to the effect that we know the extensions of our terms. Yet, when we go to explicate the nature of the relation between a term and its extension, we come up empty-handed. This is puzzling, because from the fact that we know the extensions of our terms, we are tempted to infer that the relation of reference must exist. Before I can

try to dispel this puzzlement, I need to say more about what the problematic relation is *supposed* to be.

Sometimes the term “reference” refers only to a relation that holds between each singular term and the individual object that it is said to “refer” to. But the term “reference” may be, and is, used more generally to refer to a relation that holds in addition between predicative expressions, such as nouns, adjectives, and most verbs, and their *extensions*. What this relation is supposed to be is a relation between words and (typically) non-words that plays various theoretical roles in the theory of language, which I will itemize presently. To distinguish this relation from others, we may call it *semantic* reference.

Above all, the concept of semantic reference has a place in a certain technique – call it *referential semantics* – for defining the truth of sentences. For example, we say that the sentence “Socrates is snub-nosed” is true if and only if the name “Socrates” *refers* to Socrates, the predicate “is snub-nosed” *refers* to the set of snub-nosed things, and Socrates is a member of the set of snub-nosed things. We say that the sentence “Everything has a cause” is true if and only if for every object x there exists an object y such that the pair $\langle y, x \rangle$ is a member of the set of pairs that “causes” refers to. More generally, we may formulate the conditions under which a sentence is true *relative to a context* in terms of, inter alia, the extensions of the predicative expressions that it contains. (See, for instance, Larson and Segal 1995.)

This approach to defining the conditions under which a sentence is true may play a role in further important projects. We may think of truth, defined in terms of reference, as a special case of truth in a possible world. Truth in a possible world is like *truth*

except that at a world other than the actual world some of the terms of the language may have different extensions than they have at the actual world. Then for each sentence, we can define the proposition expressed by that sentence as the set of worlds in which the sentence is true. This definition is useful in turn because it allows us to define the logical validity of arguments (consisting of a set of premises and a conclusion) as follows: An argument is valid if and only if for every context, the intersection of the propositions expressed by the premises in that context is a subset of the proposition expressed by the conclusion in that context. (Here I follow Barwise and Etchemendy 1999 in assuming that the usual technique of defining truth in a *model* and then defining logical validity in terms of that is just a formally tractable *Ersatz* for the definition of logical validity in terms of possible worlds. But I supplement their account by allowing that truth is relative to context as well as to world.)

The notion of a proposition that we define in terms of the concept of reference generalized to truth in a possible world may also be put to use in explaining the nature of successful communication. We might say that linguistic communication between a speaker and a hearer is successful if and only if the hearer correctly grasps the proposition that the sentence the speaker utters expresses in the context that pertains to the speaker's utterance. (For theories of communication that accept this basic picture more or less, see Lewis 1975, Loar 1981, Sperber and Wilson 1986, and Davis 2003.)

We normally think of words as elements of a spoken language such as English, but in recent decades many philosophers have conceived of semantic reference as first of all a property of words in a language of thought, distinct from any language we speak (Loar 1981, Fodor 2001, Davis 2003, Carston 2008, among many others). For present

purposes I can play along with this assumption (although in other writings I reject it; see my 2011). On this assumption, *concepts* may be identified with words in the language of thought. (Other modes of individuating concepts will be useful as well.) On this assumption, the reference of words might be explained in terms of the reference of concepts. What a spoken word refers to, it might be said, is just whatever is referred to by the word in the language of thought that the spoken word conventionally expresses.

Semantic reference contrasts with speaker reference. Speaker reference is something that speakers *do*. We are speaking of speaker reference when we ask, “To whom were you referring a moment ago when you said, ‘*She* won’t even let us have a goldfish’?” Speaker reference and semantic reference can come apart in that what a *speaker* refers to by means of some words may not be what those words *semantically* refer to (because the speaker may make a mistake or flout convention). But for many theorists the two notions come together when it comes to explicating the semantic reference of *demonstratives*, such as “this” and “that”, for it is often assumed that the semantic reference of a demonstrative can be defined in terms of (even if not just *identified with*) what the *speaker intended* to refer to (e.g. King, forthcoming). (For a critique of this approach, see Gauker 2008.)

The present puzzle concerns only semantic reference, not speaker reference. The intuition I want to account for is the intuition that we know what our terms refer to. The conclusion we may be mistakenly tempted to draw from this is that our terms semantically refer. My focus through most of the paper will be on predicative expressions (nouns, adjectives and verbs), not singular terms. But in the penultimate

section, I will take up the issue of singular terms and explain how they elicit a similar, but not entirely identical, dialectic.

2. *The prospects for a theory of semantic reference*

Theories of semantic reference are aimed at an explication of the concept of reference that would allow it to play the roles described in the previous section in explicating truth conditions, logical properties and communication. What is wanted is something nontrivial, informative of the form,

Scheme for a theory of reference: Term t refers to e relative to parameters $P_1, P_2, P_3 \dots$ if and only if $t \dots e \dots P_1 \dots P_2 \dots P_3 \dots$

The parameters $P_1, P_2, P_3 \dots$ may include such things as language, world, time, context of utterance, etc. Let us say that an informative explication of reference of that form would be a *robust* theory of reference.

In the past half century or so, the philosophical community has produced a whole series of different attempts to fill in this scheme. The theories of Kripke (1972), Fodor (1987), Millikan (1986, 1989), Cummins (1991), Lewis (1983), can all be construed as aiming at something along these lines, if we bear in mind that the primary target of a robust theory of reference may be either mental representations or spoken languages. All of these theories could be described as seeking, broadly speaking, “naturalistic” accounts of reference, by which I mean accounts in which the terms on the right hand side of a theory that satisfies the above scheme are recognizable as describing natural properties of the sort that might be studied by the sciences. But I do not assume that we could not be satisfied by a different sort of theory that did not qualify as naturalistic. For example, we

might be satisfied with a theory of reference in which normative terminology appeared in the right hand side. Theories of that kind may be found in Sellars (1954), Brandom (1994) and Gibbard (2012).

In this paper I will simply take for granted that none of these theories of reference has been successful and that this history of failure is sufficient reason to abandon hope that we can have a robust theory of reference. (Some of my own criticisms may be found in my 2003 and 2011. For one of many other critiques, see Hattiangadi 2007.)

Moreover, I will assume that if the reference relation is supposed to play the role I have described in a theory of truth, logic and communication, then reference has to be explicable robustly. It is not satisfactory to simply declare reference an irreducible primitive, for then one cannot answer simple questions, such as why the extension of “dog” is the set of dogs and not the set of dogs minus Rover. Given these assumptions, we can no longer expect to define truth, logical properties and successful communication in terms of reference in the manner I have described. The reference relation, conceived as serving such definitions, just does not exist. This paper is primarily aimed at those who are willing to entertain this conclusion, at least as a hypothesis.

Fortunately, in denying that we can expect to have a robust theory of reference that will play the requisite role in a theory of truth conditions, logic and communication, we need not deny that we will be able to understand the nature of language, communication and scientific progress. We can still hope to understand all of these things, but in ways that do not rest on a relation of reference that we do not understand. Currently there are several programs in the philosophical literature that aim to understand some of these things in ways that do not depend on the existence of an explicable

reference relation. These include the deflationism of Horwich (1990, 1998) and the inferentialism of Brandom (1994) (who, though he also proposes a theory of reference, does not rest his semantics on it). I would particularly like to recommend my own semantics of assertibility in a context (Gauker 2005) and my conception of language as a tool for achieving collective goals (Gauker 2011).

3. The content of the resilient intuition

However, my topic in this paper is not alternatives to referential semantics. My topic is the resilient intuition I began with: We feel that we know what our terms refer to, that we know the extensions of our terms. Those who believe that there is a relation of reference that we can put to work in the way I have described in defining truth, logical properties and successful communication will be able to explicate the content of this intuition as follows: What we intuitively feel we know is what our terms stand in this relation to. But if, as I will assume, there is no relation of reference that can play the roles that philosophical theories assign to it in explicating truth, logical properties and successful communication, then our intuition cannot be correctly explicated in this manner. What else, then, can be said on behalf of our intuition that we know what our terms refer to? That is my question. Proponents of the competing programs (such as deflationism and inferentialism) might have their own answers to this question, but I will not attempt to take issue in any way with these other theories.

Someone might claim that we have an equally resilient intuition that truth can be explained in terms of the reference relation and that our intuition that we know what our terms refer to is tightly bound up with this. I think this claim can be resisted. We cannot credibly claim to have an intuition that referential semantics is correct or that the theory

of communication that builds on it is true. Those are transparently theoretical claims that have no claim on our allegiance apart from their virtues as theories. Even leaving aside doubts about the nature of the semantic reference relation, it has proven very difficult to devise a referential semantics that explains how the truth conditions of natural language sentences rest on the relation of reference between terms and their extensions. The aspiration of constructing a comprehensive referential semantics for a language like English is one that a person acquires only through a lot of formal instruction. (I do not take this as a strike against referential semantics. The same will be true of any approach to formal semantics. The point is only that it cannot be claimed that the referential approach is directly supported by “intuition”.)

For this reason, I also resist characterizing our intuition as an intuition that *our terms do refer* or that *our terms do have extensions*. That could be innocent enough if taken only as another way of saying that we have an intuition that *we know what our terms refer to* or that *we know the extensions of our terms*. But we need to avoid the suggestion that we have an intuition that referential semantics is correct. I intend my formulation of the content of our intuition – that we have an intuition *that we know what our terms refer to* – to be neutral. If we were prepared to accept referential semantics, then we could appeal to it in explaining the content of our intuition, and we might also appeal to the facts of referential semantics in explaining the legitimacy or even the provenance of our intuition. If we reject referential semantics, we will be free to explicate the intuition in ways that do not presuppose the correctness of referential semantics.

4. *The critical counterfactual*

Before I try to say what is right in our intuition that we know the extensions of our terms, I first want to review several ways of trying to undermine it. Maybe on closer inspection it turns out that this intuition is just confused. I think it is not just confused, but it will be subject to significant modification.

Consider again my intuition that for any object that I can imagine I can tell whether the term “chair” applies to it. The intuition might be expressed as a counterfactual, thus: For any middle-sized object, if I were to get a good look at it in close proximity under good light, then I would be able to tell whether “That is a chair” is true of it. Call this the *critical counterfactual*.

Thus one could raise doubts about the intuition by raising doubts about the critical counterfactual. *Would* I always be able to tell whether “That is a chair” is true of a thing? What if I were suddenly struck by the mad insight that there is no essential difference between a chair and a stool, and so stools deserve to be called “chair” too? What if I decided that a thing is a chair only if *I* would like to sit in it (regardless of what others do)? Perhaps such possibilities can be ignored. In general, the evaluation of a counterfactual conditional is relative to a context in which certain possibilities are excluded or ignored (Gauker 2005). But what, in the present instance, is the general principle by which possibilities are ignored? If the ground on which I ignore those possibilities is my intuition that I know the extension of “chair”, then, it seems, I have no defense against the objection that both my intuition and the critical counterfactual that expresses it are mistaken.

The trouble with this objection is that it misunderstands the nature of counterfactual assertions. I say to myself, “If such and such were to happen, then I would assert, ‘That is a chair’,” and in accepting that, I may indeed be ignoring certain possibilities. And it may be that the reason why I am justified in ignoring those possibilities is ultimately only that I do know the extension of “chair”. Even so, I may be justified in asserting the counterfactual without presupposing that I know the extension of “chair”. So this counterfactual, leaving certain possibilities aside, may be a perfectly good way for me to express my intuition. Noting its truth, in a context in which I am in fact justified in ignoring certain possibilities, may be even be a way for me to discover for the first time, that I do know the extension of the word “chair”.

5. Context-relativity

I imagine a few chairs and on that basis persuade myself that I know the extension of “chair”. But isn’t that just because I only consider paradigm cases? I imagine sitting in a classroom, and I can easily decide whether each of the things in the classroom is a chair. But suppose I perform the experiment while sitting at a campfire. Would I then be so sure whether a certain log, which a person could sit on, was or was not a chair? What if I and my chair were shrunk to the size of speck of dust? Would the thing I was sitting on still be a “chair”? Perhaps I only fool myself into thinking I know the extension of “chair” by confining myself to certain paradigm cases. If I really set my imagination free and considered all kinds of cases, then I would realize that I cannot so confidently apply the term “chair” after all, and in this way I might undermine the intuition I began with.

This is not persuasive. At most we are going to dig up some more borderline cases, some context-relativity, or some unsuspected relationality. Maybe the term “chair”

is context-dependent; its extension varies from one context to another. Maybe a log is in the extension of “chair” relative to campfire contexts but is not in the extension of “chair” relative to indoor contexts. Maybe speck-sized chairs are in the extension of “chair” relative to heretofore unrealized speck-sized contexts. Or maybe “chair” is implicitly, at the level of semantics, a relation word, meaning always *chair-for*. Similarly, to say that something is “tall” is always to say that it is *taller than* some contextually given reference object. So the things that we normally call “chairs” are chairs *for* normal humans, and the imagined speck-sized chairs are chairs *for* speck-sized humans. In that case, the reference of “chair” might be a set of ordered pairs, such that for each pair in the extension, the first member is a chair for something, and the second member a thing that it is a chair for.

There is something for us to figure out here, namely, how we want to write the semantics for “chair”, and this is not decided by imagining a few chairs. But once we have figured out the basic form of our semantic theory, intuition tells us that we will not have any trouble in knowing what the extension of “chair” is for any given context or what objects stand in the *chair-for* relation to any given reference object. So the context relativity of semantic reference is not ultimately a challenge to our intuition that we know the extensions of our terms.

6. *The possibility of error*

Why does it strike me as so clear that I know the extension of “chair”? I can form mental pictures of all kinds of things and in each case I can easily decide whether the term “chair” correctly applies. So we might infer that in general, for any object x , if in imagining a thing that looks like x I would imagine something to which the word “chair”

applies, then x is in the extension of “chair”; and otherwise it is not. No questions are begged in speaking of whether a term “applies”. To say that a term *applies* to something that I imagine means simply that if I were to form a mental image of that thing (whether real or not) and ask myself “Is that an N ?” (or “Does that N ?”), then I would answer, “Yes”. In this way, I assure myself that I know the extensions of my terms.

This kind of reasoning is persuasive at most for words for kinds of thing that I can recognize just by looking. The reasoning can perhaps be extended somewhat by means of additional steps. I might imagine various animals and be unsure whether some of them should be called “mouse” or “rat”. But I can recognize a few clear cases to which “mouse” applies and can conceive of the rest of the extension of “mouse” as comprising those things that stand in a certain relation of sameness-of-species to those clear cases. Or again, I cannot tell just by looking whether a piece of yellow metal belongs to the extension of “gold”, but I can provide a definition of the extension using terms that it seems to me I understand perfectly well, such as “stuff composed of atoms having a nucleus containing exactly 79 protons”.

In all of these cases, however, – even in the case of “chair” – there is the possibility that I will make mistakes. I may apply a term incorrectly on some occasions, and I may provide mistaken definitions. Or I might not know the extensions of the terms of the definition (“proton”). Of course, these mistakes can often be corrected. Other people more knowledgeable than me about the world or about words can correct me when I make mistakes. Sometimes the whole community may make mistakes. Sometimes the whole community, due to general ignorance, may be *disposed* to make a

mistake. But even then, further observation and experiment, or at least a major revolution in the sciences, might enable us to correct our mistakes.

Certainly we cannot obtain in this way any general definition of the relation of being in the extension of a term. We might try: An object belongs to the extension of a term if and only if the term can be *correctly* applied to the object. But if to say that an application is “correct” is just to say that a sentence containing the term is *true*, then the definition is question-begging, since we expect to appeal to semantic reference in explaining truth. If “correct” just means that our applications are the product of a methodologically sound procedure, then the definition will be extensionally incorrect in many cases, since nothing guarantees that every application that results from a methodologically sound procedure will be true. In many cases, moreover, sound procedure will yield no verdict on the truth of a sentence at all because we may be dealing with questions that simply cannot be answered on the basis of humanly available evidence (such as questions about small-scale events in the remote past).

Accordingly, I am forced to concede that I may not exactly *know* the extensions of my terms. Still, I can maintain that I am at least *competent* in the application of my terms. And from the fact of that competence it is still tempting to infer that a word possesses a definite (though perhaps vaguely bounded) extension. A word has to have a definite extension, one might reason, because only that extension provides the standard by which my competence might be measured. My competence in the application of “chair” is demonstrated in my ability to apply it, in assertions of the form, “That is a chair”, only to objects that belong to the actual extension. I am competent in the use of a word only to the extent that the class of things to which I am disposed to apply it

approximates to the actual extension. Thus I might still take my competence with “chair” as evidence that the term “chair” has a definite extension.

So the concession that I do not exactly *know* the extensions of my terms does not unravel the puzzle that we started with. Our competence in the application of terms still gives us reason to suppose that our terms have a definite extension. And so it is puzzling that philosophers have not been able to construct a viable account of what it means for a term to have a certain extension. In view of this, we may continue, speaking loosely, to characterize the starting intuition as the intuition that we *know* the extensions of our terms.

7. *Knowing the extension versus knowing the meaning*

I imagine a variety of objects and confidently label some of them “chair” and confidently deny that label to others. On that basis I conclude that I know the extension of the word “chair”, which means of course that the word “chair” does have an extension. But what does this mean, that I “know the extension”? When we scrutinize this question, it becomes evident that the intuition we started with is somewhat ill-expressed when formulated in terms of extensions rather than meanings.

Though I “know the extension” of the word “chair”, that does not mean that I stand in a relation of acquaintance to each member of the extension. Some chairs are hidden in basements in Uzbekistan. Some chairs will not be made until after I am dead. Moreover, the thought experiment by which I persuade myself that I know the extension of “chair” has no more to do with *actual* chairs than it has to do with merely possible chairs that I *imagine*. So the first conclusion we should draw is that I know the *meaning* of the word “chair”. Only the meaning, and not the extension, “includes” in some sense

the merely possible chairs that I think I would apply “chair” to. If in addition we may say that I know the *extension* of “chair”, then that is only because the meaning, or intension, that I know determines the extension of the word “chair” in the actual world.

In any case, it is not very easy to see what “knowing the extension of a term” could amount to other than knowing what *kind* of thing the term applies to. If I know the extension of “chair”, I of course do not have the entire membership spread out before my mind, each member clearly distinguished from every other. My knowledge could be acquaintance with the kind *chair* or the property *chairhood* and my knowledge that the word “chair” *expresses* that kind. Or my knowledge might be an ability to correctly apply the word “chair”. In either case, my knowledge could equally well be described as knowledge of the meaning of “chair”. So my intuition that I know the extension of “chair” seems to reduce to, or depend on, my knowledge of the *meaning* of “chair”.

So our question, “What accounts for my intuition that I know the reference of my terms”, has been usurped. It has now become, “What accounts for my intuition that I know the meanings of my terms?” This transformation of our question from a question about reference, or extensions, to a question about meaning is helpful. It shows where we have to focus our diagnostic effort. Our premise is that I confidently label certain objects, or imagined objects, “chair” and confidently withhold that label from certain other things. On that basis, we infer that I know the meaning of the word “chair”. What we have to account for is precisely the inference from that premise to that conclusion. What we have to account for is our intuition that we know the meaning of our words, where a knowledge of the meaning of a word in the pertinent sense is something that can be inferred from a number of episodes of labeling.

8. *Skeptical solutions*

So our question turns out to be: What do we know in knowing the meanings of our terms? But in answering this question, we do not need to preserve all of the theories that lead from our intuition that we know the meaning of our terms to the conclusion that our terms have definite extensions. In particular, we do not have to preserve the assumption that if a predicative expression has a definite meaning, then it has a definite extension. As I explained in section 2, those patently theoretical assumptions can be addressed by putting forward alternative theories of truth, logic and communication, and, as I explained in section 3, they cannot rightly be described as objects of intuition. I will begin by explaining what we are doing when we attribute knowledge of meaning to another, and then I will complete my answer by applying that explanation to my own case.

My account of knowledge of meaning will be a variety of “skeptical solution” in the sense of this term that Kripke introduced in his discussion of Wittgenstein’s puzzle about rule-following (1982). I take it that a *straight* solution to a question about the nature of a thing is a solution that provides a naturalistic reduction of that kind of thing to things of a kind not in question or at least defines that kind of thing in a general, but non-question-begging way. A *skeptical* solution to a question about the nature of a kind of thing is a solution that simply tells us something about the use of a word that, as we say, denotes that kind of thing. (For a defense of this interpretation of Kripke, see Gauker 1995.)

Thus, a skeptical theory of a kind of thing may address the social utility of our talk of that kind of thing and deny that there is anything more to say about the essential

nature of that kind of thing. For example, we might give a skeptical theory of *knowing how to add*. This might take the form of an account of the social status we grant to a person in saying that he or she “knows how to add”. In declaring that someone “knows how to add”, we mark a person as capable of playing certain roles in society. (See Kripke’s remarks on this at pp. 92-93.) It may be that we do that only after we have put that person through a variety of tests, but that does not mean that “knows how to add” just *means* that the person has passed the tests. We would not be able to say in any general way what tests are necessary and sufficient to justify us in labeling a person as “knowing how to add”. Since I have not proposed to present a theory of knowing how to add in this paper, I will not attempt to say what the social function of labeling a person as knowing how to add might be.

In any case, a consequence of a skeptical theory of knowing how to add may be that if any two people whom we have properly labeled as “knowing how to add” both receive the same addition problem, then they will *probably* give the same answer. The basis for our expectation that they will give the same answer will not be that they both know how to add, however. Rather, as the skeptical theory will affirm, our observations of their adding behavior will license us in granting to both of them the social status *knows how to add*, and those same observations, quite apart from these people’s having this status, will license an inductive inference to the conclusion that each will give the same answer as the other to questions about addition in the future.

9. *Meaning in ordinary language*

In ordinary, nonphilosophical discourse, the term “meaning” has many uses that are completely irrelevant for present purposes. For example, we might speak of the historical

“meaning” of some political events, or we might speak of what someone “meant” to achieve in performing some action. Of interest here is only the kind of meaning we might call *linguistic meaning*.

My hypothesis is that, in ordinary discourse, talk of linguistic meaning occurs in the context of *conversational repair*. Person *A* does not understand something person *B* has said and asks, “What do you mean?”, whereupon *B* defines the terms in question or just reformulates the point in other words. For example, person *A* might ask person *B*, “What do you mean by *officious*?”, and *B* might answer by supplying some other words that *B* might have used in place of “officious”, in the expectation that *A* will not likewise need to ask about the meaning of those other words. In such a case, we speak of linguistic meaning as something that a person *does* by means of some words. In other cases, we may ask about the meaning of some *words* or a symbol (“What does ‘officious’ mean?”), but even in that case we expect the answer to take the form of some other words that we understand that might be used in its place. In a foreign land we might ask about the meaning of some traffic sign, and the answer might be, “That sign means ‘Do not park here!’.”

We can also speak of linguistic meaning in the course of instructing a person in how to converse with other people. For example, we might ask a child, “Did you do that *intentionally*?” and his mother might intercede, saying, “He doesn’t know what ‘intentionally’ means.” So then we might try, “Did you do that *on purpose*?” or “Did you *want* to do that?” Or in trying to figure out why a Chinese guest in a hotel asked the housekeeper for a “blanket” on a warm summer day, we might speculate that by

“blanket” she meant *towel*. In that case, we are suggesting to the housekeeper that he, the housekeeper, might act as though the guest had said “towel” instead of “blanket”.

In normal cases, where communication goes smoothly, we know how to respond to people’s words without thinking about meanings. We might say that understanding another’s words always involves “interpretation”, because our response depends on the speaker’s choice of words, but understanding another’s words is not usually a matter of ascribing meanings. Rather, responding to the words of a member of our linguistic community is normally the flip side of the same skill that we exercise in speaking the language of that community ourselves, which certainly does not require us to explicitly ascribe meanings to our own words. Usually meanings are ascribed only when the process of communication does not proceed smoothly, and in that case the role played by attributions of meaning is to supply other words that we can utilize without having to address the sorts of doubts that we express by asking about meanings. (In saying this I am rejecting Davidson’s claim (Davidson 1984) that linguistic communication does not rest on shared practice. For more on this issue, see Gauker 1994, chapter 2.)

In some cases, there might be some doubt about whether in choosing our response to a person we need to interpret him or her by substituting some other words for those that he or she actually used. If we decide that some such interpretation is called for, we may introduce our interpretation by saying something like, “He does not know what the word ‘officious’ means,” and we may state our interpretation by saying, “What he thinks it means is . . .” Alternatively, we may decide that no interpretation is needed, and in that case, we may answer our original doubt by saying, no, the speaker knows the meaning of his or her words. In so doing, we grant the speaker a status, that of being someone whose

use of those words we may respond to without substituting other words. That we answer such doubts in this way, by granting the speaker the status *knowing the meanings* of his or her words, is my skeptical theory of knowledge of meaning.

10. Testing knowledge of meaning

In granting someone the status of knowing the meaning of a word, we wager our own status as knowers of the meaning of that word. We lose that wager if, as a result of our granting that status to someone else, others are unwilling to grant it to us. So if there is any question whether we should grant that status to someone, we will want some assurance. For instance, if there is reason to doubt it, we might question whether someone knows the meaning of “chair”. That is, we might want to know whether we should respond to him as if he had used some other words in place of “chair”—perhaps “step ladder”. In that case, we could put him to the test by asking him to tell us, for each of a few objects, whether it is a “chair”. Or, more likely, we would pay attention to his spontaneous uses of “chair”. If we agree with all of his answers, or accept his uses as correct, then that settles the matter well enough. We conclude that he *knows what the word “chair” means*, and in consequence of that, we do not, in conversing with him, go through an extra step of imagining what he might have said in other words instead of using the word “chair”.

More generally, if the question arises whether we should grant to someone the status of knowing the meaning of some word, we can settle the question by testing for what we might call *communicative compatibility*. Communicative compatibility in the use of a word can be demonstrated by testing for and finding agreement in questions formulated by means of that word. Two people may enjoy communicative compatibility

in the use of a word even though they do not agree on every question formulated in terms of it; but if two people enjoy communicative compatibility in the use of a word, then they should at least be able to understand their disagreements. Each should be able to locate the basis for their disagreement, and this basis should not be either one's thinking that the other does not know the meaning of some words. You may grant that I know the meaning of "gold" while expecting that I will sometimes disagree with you over the labeling of something as "gold"; but even then you expect that my mistakes will be explicable without reference to meanings (for instance, as the result of misleading appearances) and that I will either yield to persuasive arguments or offer some objection that you can understand.

Saying that someone knows the meaning of a word does not *mean* that the person has passed certain tests for communicative compatibility with respect to the use of that word. Nor does saying that provide an *explanation* for the communicative compatibility we have observed. No, my hypothesis is that to say that someone knows the meaning of a word is to grant a certain status, as I have explained. However, in granting that status to a person we do expect that the person will not, so to speak, betray our trust by failing tests of communicative compatibility. The basis for that expectation may be whatever basis we might cite, if doubts were raised, for granting the status. That basis may be tests of communicative compatibility that the person has passed, or it may be simply that the word in question is a common word in the community in which the person was raised.

Whatever the basis for one's granting the status of knowing the meaning of a word may be, that basis will be one's own basis for expecting that the person will go on passing tests of communicative compatibility with respect to that word. If the grant of

the status originates with ourselves, then that basis will be the attributee's having passed some tests for communicative compatibility. But in granting to someone the status of knowing the meaning of the word, we may license others in expecting that the person will generally pass tests of communicative compatibility with respect to that word, even if they themselves have not tested the attributee for communicative compatibility with respect to the word in question. If we give that license to others, then their expectation of communicative compatibility will be grounded in the license we have given them.

11. Attributing knowledge of meaning to oneself

Suppose I ask myself whether I know what the word "chair" means. I may answer this question by putting myself, in imagination, to the same sorts of tests that I would put others to if I wondered whether they knew what "chair" meant. This need not be a pointless charade. In the case of some words, I might find that I am not so sure about some applications, not because those are borderline cases but just because I really do not know the meaning of the word. If I pass the tests, then, just as I would if I were administering the test to someone else, I will grant to myself the status of knowing the meaning of the word "chair". Others, of course, recognizing the biases I may be subject to, will not rely very heavily on the grants of this sort that I make to myself.

Thus I have explained my sense that I know the meaning of the word "chair". In performing the thought experiments, I give myself in imagination the sorts of tests I would give to others in deciding whether to grant them the status of knowing the meaning of "chair". If I pass those tests by my own lights (I might not), I grant myself the status of knowing the meaning of the word "chair". As a dividend, I have explained why

knowing the meaning of “chair” implies competence in the use of “chair”. It is because we do not grant the status unless we have a basis, either in our own inductive evidence or in the license extended to us by others, for expecting communicative compatibility in the future.

This way of explaining our confidence that we know the meaning of a word gives us no reason to think that we should be able to explain in a general way the relation of reference between a term and an extension. We find no reason to think that in advance of any application there is a discrete or fuzzy range of things comprising all and only the things to which the term is truly applicable. This way of explaining our confidence leaves us free to accept the lesson that our repeated failures to define the reference relation seem to force on us: There is no such thing as semantic reference.

12. The attribution of status

What is the difference between *describing* a person and *granting a status* to him or her? I have said that in attributing to a person knowledge how to add or knowledge of the meaning of a word, we are granting a status, not describing, as if that account of what we are doing somehow allowed us to evade the question, “What are the truth conditions for the attribution?” The correctness of what I have said so far may be judged independently of any account of the distinction between describing and granting a status, so long as we have an intuitive grasp of the distinction. But to try to undergird this intuitive grasp, let me briefly take a stab at drawing the distinction. For starters, the difference between describing a person and ascribing a status lies in the network of human relations to which

the two kinds of act belong. If *B* relies on *A*'s descriptions, by accepting them, then *B* presumes that *A* is *credible*. But if *B* relies on *A*'s attributions of status, then *B trusts A*.

In granting someone a status, we do not describe a state of mind from which communicative compatibility with respect to the use of that word will in the future flow. Much less do we describe the history of communicative compatibility on which our expectation of future communicative compatibility is grounded. Rather, we license others in treating in a certain way the person to whom the status is granted. This "licensing" is unlike the licensing by a government office, in that in this case the power to license is sustained only by the willingness of others to be licensed by it. Hence if what is licensed leads to failure, the power to license is diminished or taken away. In that respect, the act of granting a status is at the same time a kind of wager. What is wagered is that granting the status will not result in one's being deprived of the power to grant the status.

In attributing knowledge of the meaning of a word to a speaker, we license our hearer in expecting communicative compatibility from the attributee with respect to the use of that word. As a consequence of this licensing, stake our own reputations as attributors of the status on the attributee's continuing to be communicatively compatible with respect to the use of the word. In presuming, on the basis of our attribution, that the attributee will be communicatively compatible in the future, our hearer wagers with us, which is to say that he or she places his or her trust in us.

One difficulty we face in distinguishing between describing a thing and granting a status is that the act of describing too may be viewed as a kind of wagering of our own reputation. If I describe something as having certain properties, and I prove wrong, then

to some extent I harm my reputation as a reliable describer of things. Nonetheless, there is a difference between describing and granting a status in how the speaker expects to influence others. Although a speaker does put his or her reputation on the line in putting forward a description, it is not the trade in the power to grant the status —the trade in reputation — that drives the exchange. The practice of describing things would have a value quite apart from its power to enhance or diminish the reputation of speakers, so long as speakers were generally reliable. Whereas if a speaker grants someone a status, it is precisely the trade in power to grant the status that drives the exchange. The attribution of status serves a useful purpose only insofar as speakers risk their reputations in attributing a status.

13. Singular terms

I began with a quandary. On the one hand, I have an intuition that I know the extensions of my terms. On the other hand, the prospects for a general account of the relation between terms and their extensions that we might hope to use in a theory of truth will now strike many philosophers of language as quite poor. So far I have focused on the case of predicative expressions, but the quandary is just as deep in the case of singular terms. This is what we should expect, since the hoped-for explanation of truth in terms of reference requires a theory of reference that works as well for singular terms as for predicates.

Just as we have a problem saying in a general way what the relation is between a predicate and its extension, so too we have a problem saying in a general way what the relation is between a singular term and its referent. Here we might divide the literature

into attempts to account for the reference of proper names (e.g. Devitt, this volume; Bianchi, this volume), attempts to account for the reference of spoken demonstratives (e.g. Gauker 2008; King forthcoming) and attempts to define the relation between a singular thought and its objects (e.g. Bach 1988, Jeshion 2010). This debate, more than the debate over the general reference relation, still seems to have some life in it. Still, I think it will be fair to say that at this stage in history one would be justified in doubting the possibility of any general account of the relation between a singular term and its referent.

And yet, I have the intuition that I generally know what my singular terms refer to. As in the case of predicates, our intuition that we know what our singular terms refer to encourages the belief that our singular terms stand in a relation of semantic reference to objects. The reason is that if in fact our singular terms did stand in a relation of semantic reference to objects, then that fact would provide an account of our intuition, and would suggest an account of the legitimacy and provenance of that intuition. So those of us who have reached the point of doubting that we will ever have a plausible account of the reference relation for singular terms have reason to account for that intuition in some other way.

I am aware that some of the things I assert by means of a name may not be true. But I have no other conception of the bearer of the name than the collection of judgments I might express by means of it. So I can readily see that it is possible that a name that I thought refers to something does not refer to anything or only refers to something very different from what I had supposed. Likewise, I recognize that if I utter the demonstrative “that” and accompany my act with misleading gestures, one might deem

my utterance to have no reference at all. We might even say that I can refer with “that” to something quite different from what I intended to talk about. But these are unusual cases. So despite them, we can stand by our intuition that by and large we are “in command” of the reference of our singular terms.

A difference between the case of singular terms and the case of predicates is that in the case of singular terms not much would be accomplished by substituting an intuition that we know the meanings of our words for the intuition that there is a fact of the matter about the reference. In the case of predicates, the substitution seemed acceptable, because the acts of imagination by which we persuade ourselves that we know the extensions of our terms do not distinguish between actually existing members of the extension and merely possible members of the intension. The substitution of the intuition about meaning for the intuition about extension was useful, because it left us with an intuition amenable to skeptical diagnosis in terms of our practices of attributing knowledge of meaning. Of course, semantic theories may distinguish as well between the meanings of singular terms and their referents, but the substitution of the intuition that I know the meanings of my singular terms for the intuition that I know the extensions of my singular terms does not seem to be an equally useful reformulation of the question.

However, in the case of the reference of singular terms, I can offer a different sort of skeptical diagnosis. My diagnosis in this case will be consonant with my diagnosis of our sense that we “know the meaning” of our predicates in assuming that in ordinary discourse semantic terminology serves the purpose of conversational repair. Our conviction that we know what our singular terms refer to, I suggest, may be traced to our practice of asking, “Who are referring to?” and expecting an answer. If we use a singular

term in a way that does not permit our hearers to understand, we are liable to be asked “Who are you referring to?” or “What were you referring to?” Someone might say, “That customer from Ohio called again today and asked to speak to the manager, but the manager wouldn’t speak to him, because he was still angry”, and we might ask, “Who were you referring to, when you said ‘he was still angry’?”. In such cases, we typically hold speakers responsible for answering the question satisfactorily.

So in asserting that someone knows what a given singular term refers to, we are saying that, when it comes to responding to that person’s speech with respect to his or her use of that singular term, no conversational repair is necessary. To assert in this way that no conversational repair in this regard is necessary is not to describe the speaker’s state of mind but to grant a status. The status granted is that which one possesses when others are licensed to respond to the person’s words without reinterpretation. We do not generally grant such statuses without reason, however, but on the basis of tests of communicative compatibility. In particular, the speaker’s answer to the question, “Who were you referring to” must be one that *makes sense*. When we test our own communicative compatibility in the use of singular terms we generally find, not surprisingly, that we pass the test, and so we readily grant ourselves the status of knowing what our singular terms refer to.

14. Summary

I began by posing a puzzle: Why does it seem so obvious that I know the extensions of my terms, though all attempts to explicate the semantic reference relation have failed so badly (sections 2 and 3)? The initial intuition that I do know the extensions of my terms

survived a couple of challenges (sections 4 and 5). It was bruised by the observation that the possibility of error in applying our terms is never fully removed (section 6), but that observation was not sufficient to shake our confidence that we possess a certain competence in the application of terms, which is sufficient to sustain the puzzle. But as it turned out (section 7), our confidence that our terms have definite extensions reduces to our confidence that we know the meanings of our terms.

What remained of our initial puzzle, then, was to explain our intuition that we know the meanings of our terms. I proposed to answer this in a way that left behind the theories of meaning that seem to demand a robust theory of reference. I then explained the concept of a skeptical solution and proposed to give a skeptical solution to the question concerning knowledge of meaning (section 8). Knowledge of meaning, I suggested, is a status we grant to a person when we have no need to interpret him or her by substituting other words for those that he or she spoke (section 9). The status is granted when the agent in question passes tests of communicative compatibility (section 10). In light of this I explained my intuition that I know the meanings of my words: On the basis of a few thought experiments, I can grant myself the status of knowing the meaning of my words (section 11). This left us with a question about the difference between describing a person and attributing a status, which I took a stab at answering (section 12). Finally, I reiterated the basic dialectic, with modifications, to show that it pertains to the case of singular terms as well (section 13).

In closing, I would like to emphasize that what I have done here is at most assuage a recalcitrant intuition. The hard work of dispensing with semantic reference lies

in constructing alternative accounts of those linguistic and mental phenomena that we might have thought we should explain in terms of semantic reference.¹

References

- Bach, Kent. 1988. *Thought and Reference*. New York: Oxford University Press.
- Barwise, J. & Etchemendy, J. 1999. *Language, Proof and Logic*. Stanford: CSLI Publications.
- Bianchi, Andrea. This volume. “Repetition and reference.”
- Brandom, Robert. 1994. *Making It Explicit: Reasoning, Representing and Discursive Commitment*. Cambridge, MA: Harvard University Press.
- Carston, Robyn. 2008. “Linguistic Communication and the Semantics/Pragmatics Distinction.” *Synthese* 165: 321–345.
- Cummins, Robert. 1991. *Meaning and Mental Representation*. Cambridge, MA: MIT Press.
- Davidson, Donald. 1984. “Communication and Convention.” In his *Inquiries on Truth and Interpretation*, pp. 265-280. Oxford: Oxford University Press.
- Davis, Wayne A. 2003. *Meaning, Expression and Thought*. Cambridge University Press.
- Devitt, Michael. this volume. “Should Proper Names Still Seem So Problematic?”
- Fodor, Jerry A. 1987. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.

¹ I am grateful to an anonymous referee for perceptive comments on an earlier draft, and I thank Andrea Bianchi for his comments and very sound advice.

- Fodor, Jerry A. 2001. "Language, Thought and Compositionality." *Mind and Language* 16: 1–15.
- Gauker, Christopher. 1994. *Thinking Out Loud: An Essay on the Relation between Thought and Language*. Princeton, NJ: Princeton University Press.
- Gauker, Christopher. 1995. "A New Skeptical Solution." *Acta Analytica* 14: 113–129.
- Gauker, Christopher. 2003. *Words without Meaning*. Cambridge, MA: MIT Press.
- Gauker, Christopher. 2005. *Conditionals in Context*. Cambridge, MA: MIT Press.
- Gauker, Christopher. 2008. "Zero Tolerance for Pragmatics." *Synthese* 165: 359–371.
- Gauker, Christopher. 2011. *Words and Images: An Essay on the Origin of Ideas*. New York: Oxford University Press.
- Gibbard, Allan. 2012. *Meaning and Normativity*. New York: Oxford University Press.
- Hattiangadi, Anandi. 2007. *Oughts and Thoughts: Rule-Following and the Normativity of Content*. New York: Oxford University Press.
- Horwich, Paul. 1990. *Truth*. New York: Oxford University Press.
- Horwich, Paul. 1998. *Meaning*. New York: Oxford University Press.
- Jeshion, Robin. 2010. "Singular Thought: Acquaintance, Instrumentalism and Cognitivism." In Robin Jeshion, ed., *New Essays on Singular Thought*, pp.105–140. New York: Oxford University Press.
- King, Jeffrey. Forthcoming. "Speaker Intentions in Context." *Noûs*.
- Kripke, Saul. 1972. "Naming and Necessity." In Donald Davidson and Gilbert Harman, eds., *Semantics in Natural Language*, pp. 253–355. Dordrecht: D. Reidel.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.

Larson, Greg and Gabriel Segal. 1995. *Knowledge of Meaning: An Introduction to Semantic Theory*. Cambridge, MA: MIT Press.

Lewis, David. 1975. "Languages and Language." In Keith Gunderson, ed., *Language, Mind and Knowledge*, pp. 3–35. Minneapolis: University of Minnesota Press.

Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61: 343–377.

Loar, Brian. 1981. *Mind and Meaning*. Cambridge: Cambridge University Press.

Millikan, Ruth. 1986. "Thoughts Without Laws: Cognitive Science with Content." *Philosophical Review* 95: 47-80.

Millikan, Ruth. 1989. "Biosemantics." *Journal of Philosophy* 86: 281–97.

Sellars, Wilfrid. 1954. "Some Reflections on Language Games." *Philosophy of Science* 21: 204–28.

Sperber, Daniel and Deirdre Wilson. 1986. *Relevance: Communication and Cognition*. Cambridge, MA: Harvard University Press.