

On Meaning, Meaning and Meaning

Introduction

Various criss-crossing distinctions have been drawn in the philosophical tradition between kinds or dimensions of linguistic meaning or between meaning and other dimensions of linguistic function. In this chapter I'll try to collect together from various books and papers the results of my own investigations on different aspects of meaning. The underlying idea is that to understand how language works, one must look, first, to the cooperative functions that various language forms perform, understanding these on a biological model as what these forms accomplish that keeps them in circulation. To explain the cooperative function of a language form is to explain its survival value, the source of its proliferation, what it does that accounts at the same time for the fact that speakers continue to use it and that hearers continue to react to it often enough in standard ways. Next we should look at language mechanics, at how language forms perform their functions. For some language forms there are conditions in the world that are necessary to support their functions and that vary systematically with certain variations in the forms themselves. These are truth conditions, and they are determined by a kind of "meaning" that I will call "semantic mapping functions" -- "functions," this time, in the mathematical sense. (Semantic mapping functions determine truth conditions; truth conditions only delimit and do not determine semantic mapping functions. I will get to this.) Last we need to describe the psychological mechanisms that are involved in implementing the functions of various language forms, the ways that speakers and hearers manage to produce and understand these forms so as to promote performance of their cooperative linguistic functions.

There is also an important distinction, of course, between speaker meaning and linguistic or conventional meaning. This is the difference between the cooperative function, with its associated truth condition and so forth, of a public language form and functions that individual speakers may use or try to use the form to serve. I will discuss this distinction, but only to set speaker meaning aside. My basic proposal is that there are these three basic kinds of linguistic meaning:

(1) Conventional linguistic cooperative function, to be called "stabilizing function."

(2) Conventional semantic mapping functions ("functions" in the mathematical sense) which determine truth and other kinds of satisfaction conditions.

(3) Methods of identifying --to be called "conceptions" and "conceptual components"-- that govern individual speakers' grasps of referents and of truth or satisfaction conditions, hence help to determine their dispositions to use and understand various conventional language forms.

I will argue that the third of these, conceptions governing individual speakers' grasps of referents and satisfaction conditions, may exhibit little or no overlap among competent members of the same language community. Thus none of the aspects of meaning that I will define corresponds at all well either to any traditional notion of intension or to any Frege-related notion of sense.¹ The meanings that characterize the public part of a

¹ This despite the embarrassing fact that in my (1984), "(Fregean) sense" was the name I gave to what I now call "semantic mapping" and "intension" my name for what I now call "conception." I had my reasons, but they were not good.

language are fully extensional. I will have to say quite a lot about that before I am finished. Let us begin with simpler matters.

I. Stabilizing functions

When speakers are conforming to the conventions of a language, speaker meanings coincide with conventional meanings. But speakers often use language forms for purposes that diverge from conventional meanings. To distinguish speaker meanings from conventional linguistic meanings, we need to know what a public language convention is. Probably the best known theory of convention and its application to language is that of David Lewis (1969, 1975). I will clarify my position by comparing it with his.

Lewis describes a convention as a regularity in the behavior of a population such that within the population there is mutual knowledge (1) that everyone conforms to the regularity, (2) that everyone prefers to conform given that the others do and (3) that everyone expects everyone else to conform for the same reason he does. The reason each prefers to conform is that conforming solves a coordination problem. A coordination problem arises when people have a purpose in common which must be achieved by joint action, where the contribution that each must make will vary depending on what each of the others contributes, and where there is more than one acceptable way of combining contributions to produce a successful outcome. Then coordination is necessary. It is best for everyone if everyone makes his contribution according to the same solution plan. To each it doesn't matter as much which plan is chosen as it matters that the same plan is chosen by all. In many cases, Lewis says, the plan that is chosen will be the one for

which there is a precedent. It has been used before, which makes it a salient plan, one that comes to mind and that each participant assumes will come to the mind of the other participants. Each participant thus steps into his role in this plan on the assumption that the others will adopt their roles according to the same plan. When a precedent for solving a coordination problem spreads in this way a convention is born. Thus Lewis claimed that social conventions of all kinds, including linguistic conventions, are supported by rational beliefs and intentions concerning one another's thoughts. It is true that children and idiots may conform to the conventions of language without having reasons of this sort but, Lewis claims, "they are not parties to the convention and their linguistic competence is incomplete" (1969 p. 51).

Now I agree that the conventions of language arise and spread because they solve certain kinds of coordination problems. Not all conventions solve coordination problems, however. And for those that spread because they do, the "because" is almost never a reasoned because but some more mundane kind of causal because. The rest of us conform to linguistic conventions in exactly the same unreasoned way that the idiot and the child do. Further, despite apparent consensus among philosophers that conventions always involve regularities of behavior within a group (Searle 1969; Lewis 1969, 1975; Schiffer 1972; Bach & Harnish 1979; Gilbert 1983, 1989/1992; Recanati 1987²), my claim is that conventions do not generally require regularities of behavior,

² Recanati (1987) takes it that conventional language devices "indicate" or, using the linguists' term, "mark" uses of language, that is, conventions mandate that these

either de facto or de jure. In particular, conventional coordinations, including linguistic coordinations, do not, in general, require regularities of behavior. These claims were defended in Chapter 1. Here I will just review the high points.

A convention, in the sense that a natural language contains conventions, is merely a pattern of behavior that is (1) handed down from one person, pair, or group of persons to others --the pattern is reproduced-- and (2) is such that, if the pattern has a function, then it is not the only pattern that might have served that function about as well. Thus if a different precedent had been set instead, a different pattern of behavior would probably have been handed down instead. Putting a wreath on the door at Christmas time, dyeing eggs for Easter and drinking green beer on St Patrick's day are conventions in this sense. In Japan the convention is to eat with chopsticks, in America, with a knife and fork. Against Lewis, that these are conventions (1) does not necessarily mean that they solve coordination problems. Also (2) it does not necessarily mean that they are

devices shall be used only for those purposes. This is not explicitly stated, but see, for example, ' 22.

universally followed. Indeed, there are many conventions for which conformity is neither prescribed nor mandatory in any sense. Of course, some conventions, such as driving on the right in the United States, do solve coordination problems, are universally followed, and are mandatory. But that is not what makes them be conventions. Also, linguistic conventions do solve coordination problems, but they are neither universally followed nor mandatory.

When a conventional pattern of behavior is handed down because it is solving a coordination problem, the mechanism for this is usually quite simple. No matter how the precedent for the convention was originally set, if the coordination it effects is an obvious and important one, it will tend to proliferate without anyone's thinking about anyone else's thoughts. Like other higher animals, people repeat behaviors that have been successful in achieving wanted results in the past. Unlike most other animals, they tend also to copy behaviors of others that have been successful in producing wanted results. Behaviors that constitute solutions to coordination problems achieve results desired by all parties to the coordination, hence these behaviors will tend to be reproduced when similar results are desired. There is no need for the various parties in the coordination even to recognize the problem as a coordination problem let alone to think about one another's thoughts in order for the convention to proliferate. If other people are driving on the right, then I will drive safely only if I drive on the right. Thus I might learn to drive on the right without ever quite realizing that it is only a convention to do so. Exactly so, not only children but very smart primitives typically are unaware that the languages they speak are merely conventional.

Specific language forms continue to be reproduced by speakers within a language community merely because, often enough, they prompt hearer responses that contribute to the fulfillment of speaker purposes in speaking. Similarly, hearers continue to respond in conventional ways, for example, by believing or by doing what they are told, because, often enough, the result is rewarding for them. Often enough, believing or doing what one is told leads to believing or doing what is useful or what will keep one out of trouble. Speakers within a language community are, simply, adapted to an environment in which hearers are responding, sufficiently often, to the forms speakers produce in ways that reinforce these speaker productions. Correlatively, hearers in the community are, simply, adapted to conditions under which speakers, sufficiently often, produce these language forms in circumstances such that making conventional responses to them aids those hearers.

Consider, for example, a speaker whose purposes in using the word "dog" will be achieved only through calling attention to dogs or to facts that concern dogs or through changing behaviors toward dogs. Such a speaker will eventually stop trying to use the word "dog" for these purposes if they are never achieved. Similarly, a hearer whose language-understanding faculties turn his mind to dogs whenever speakers use the word "dog" will soon unlearn this response if speakers never use the word "dog" such that it carries information or expresses intentions that concern dogs. Similarly, consider those syntactic forms that get labeled "indicative" in various languages. These forms usually have a number of alternative functions, but no form will be so labeled unless one of its functions is to effect production of true beliefs having propositional contents carried by

other aspects of these sentences. These conventional forms are surviving in part because, often enough, this particular effect is of interest both to speakers and to hearers. Production of false hearer beliefs may occasionally interest speakers, but rarely serves the interest of hearers. A hearer unable to interpret the indicative sentences he hears so as sometimes to extract genuine information from them would soon cease to form beliefs on their basis. And if hearers ceased ever using indicative sentences as guides in forming beliefs, speakers would stop trying to use them for purposes that required imparting beliefs. Similarly, if it were not sometimes in the interest of hearers to comply with imperatives Cadvice, instructions, directions, friendly requests, sanctioned directives, and so forthC they would soon cease ever to comply. And if hearers never complied with imperatives, speakers would soon cease to issue them. Imperative syntactic forms would become obsolete.

A corollary is that the functions of public language forms are not on the same level as either speaker purposes or hearer purposes taken alone. The conventional functions of language forms are not, for examples, merely standard speaker purposes. Conventional language forms are selected for performing services satisfactory at once to both partners in communication. Their functions must balance speaker with hearer interests. Because the conventional function of a linguistic form will remain stable only if it continues to serve the interests of both speakers and hearers often enough, I call it a "stabilizing function." Linguistic "meaning" in the sense of stabilizing function is on an entirely different level from, for example, average speaker meaning.

Similarly, on this analysis a linguistic convention consists in a pattern that includes

both a conventional contribution by the speaker and a conventional contribution by the hearer. The hearer's contribution is as much a part of the convention as is the speaker's. Thus the linguistic convention includes important aspects of what Austin called the perlocutionary act. As such it effects a genuine coordination between speaker and hearer, each of whom must play his part if the coordination is to be successful. Contrary to this, Lewis claimed that "[a] member of the audience, as such, is not constrained by convention...Only when he takes his turn as communicator does he himself act in conformity to the convention of truthfulness in L" (1969, p. 179-80).

We can see why Lewis took this position. According to his analysis, a convention was a "regularity of behavior in a population" such that "everyone conforms to the regularity." Later in Convention he modified this to allow that almost everyone conform almost all of the time (1969 p. 78), but still, it is plainly false that almost every hearer of a directive complies with it and plainly false that almost every hearer of a description believes it. So on Lewis's account the hearer's response could not be part of a linguistic convention.³ Lewis also says, "forming a belief...is normally not a voluntary action and hence not an action in conformity with convention," and of directives, "[e]ven if the audience should act, the action may not answer to an interest common to the communicator and the audience" (p.180). But if, as I have claimed, a convention is

³ Lewis also claimed however that the conventions of a language involve speaking the truth, which may be just as implausible on his assumption that conventions require nearly universal compliance.

merely a reproduced pattern whose form is arbitrary with respect to its function, then there is no requirement on how voluntarily or how regularly the pattern is reproduced or on how often the pattern is broken, with either speaker or hearer failing to contribute his or her proper part (Chapter 1 above; Millikan 1984 chapter 4). Sometimes the speaker is not interested in genuine cooperation. Sometimes the hearer is not. Sometimes mistakes are made. Conventional coordination patterns need to succeed only often enough to avoid extinction.

Notice as well that alternative coordination conventions serving the same purpose often happily exist side by side in a community. Besides linguistic conventions, Lewis talked of "signaling conventions" which he did describe as involving the receiver's responses as well as the signaler's gestures. He illustrated with signals used by a man standing outside to help a truck driver back into a tight space. The helper and the driver both want to maneuver the truck safely into the space. How the signals are composed by the helper and how they are read by the driver have to be coordinated if this common end is to be achieved. Lewis apparently overlooked that both helper and driver might easily be acquainted with more than one signaling system commonly used for this purpose. So long as these different systems didn't happen to contain identical signals that meant different things, it wouldn't matter which system the helper chose to use, the driver would recognize the signal and follow it. What is necessary for success is only that the same precedent should be followed by both helper and driver on each individual occasion. What they do on other occasions doesn't matter. An initiating move by one party will immediately be recognized by the other as coming from a particular lineage of

precedent with which both are familiar. Thus it is that linguistic conventions are neither universally followed nor mandatory. Many alternative conventions can possess the same stabilizing function.

Language conventions can be considered as lineages of precedent. A public language is a huge web of crisscrossing lineages of reproduced patterns consisting of tokens of linguistic forms and responses to them. People listen to one another, then repeat words and idioms they have heard, syntactic arrangements they have heard and tonal inflections they have heard, arranging these into new combinations. Words, idioms, syntactic forms, tonal inflections and so forth are handed down from one person to others because these elements are helping to serve coordinating functions. These stabilizing functions are, in one of that term's various senses, their "meanings," the first of the three "meanings" listed in this essay's title. One thing to investigate then is exactly what kinds of stabilizing functions compose these meanings. What various jobs do linguistic forms do to keep themselves in circulation?

The thesis that linguistic conventions correspond to reproducing lineages of cooperatively-used tokens-with-responses has direct implications for the individuation of linguistic forms. For the purposes of semantics, what makes two tokens be tokens of the same linguistic type is not their sound or shape, or the phonemes or letters of which they are composed, or their surface syntactic arrangements. They are tokens of the same type only if they have been copied from the same pool of tokens reproducing in the same language community. They must be segments from the same historical lineage. Genuine words cannot be accidentally formed by the wind. Further, any genuine linguistic token is

automatically a token from one particular language or another. When discussing linguistic forms, reference to the form as being "in L_1 " or "in L_2 ," etc., may help the hearer to identify the form intended, but whether identified or not, if the form is a genuine natural language form, it already is essentially either in L_1 or L_2 or some other language without that. Otherwise it is not a linguistic form, but merely a describable shape or sound.

Unlike the lineages that make up animal species, linguistic lineages frequently acquire new functions without changing their physical forms. Similar to mutations in biological evolution are novel uses of conventional linguistic forms introduced by speakers through figures of speech or through Gricean implicature. If the hearer understands the figure or the implicature, the novel use will serve a new coordinating function. It may then be copied by other speakers and may in time be understood directly by hearers without having to go through the process of unpacking a figure or an implicature. Then a new lineage of tokens with a different stabilizing function has branched off from the original lineage but without any change in physical form. Suppose, for example, that a new metaphor is copied again and again. For a very long time, those who use it and those who understand it may continue to read it as originally derived from its original source. For most speakers and hearers, acquaintance with the old lineage and with the new lineage may both together and equally be responsible for its use and for its easy comprehension. Later, however, the new use may become as familiar as the old and may start to be proliferated quite independently. Then the metaphor becomes "dead." An entirely independent branch of the family has been formed from tokens with exactly the same physical form. The result is called "polysemy" -- one sound, many meanings.

Families of linguistic forms quite typically form wide-spreading bushes, many different branches having slowly formed over time, and more branches from those branches. Since branches often take a long time to separate off completely, the places where true branchings begin are not at all sharp. At a given time, exactly how many branches there are is not definite. Still, it is helpful to give a name to the branches of a given form that are currently fairly independent, each being well enough established that it would survive even if all the others should die out. These independent branches I call "least types." Least types correspond to the various relatively independent stabilizing functions (different "senses") of a polysemous language form (Millikan 1984 Chapter 4). However, typically the number of such "meanings" and the divisions between them are very far from precise.

Surface syntactic forms may also branch into independent least types. They too may be polysemous, having a number of branches that propagate more or less independently. For example, definite descriptions are sometimes used by speakers merely to identify a particular referent for a hearer, the description itself being of no interest at all. Other times definite descriptions are of interest in their own right, indeed, the speaker may not care whether the hearer identifies the referent or not so long as the description is remembered. Donnellan's distinction (Donnellan 1966), when understood this way, is a distinction between two stabilizing functions of definite descriptions that tend to divide these descriptions into two least types. Definite descriptions tend to be

polysemous in stabilizing function.⁴ Similarly, consider the various grammatical moods in a language. Sometimes you impart beliefs with the indicative mood but other times you use it to give orders. Sometimes you ask questions with the interrogative mood but other times you use it to make requests. A number of relatively separate conventions are helping to propagate the same surface syntactic forms. Differences in what Austin called "illocutionary" and also "perlocutionary" force may often be carried conventionally by syntactic forms that are polysemous in this way (Millikan 1984 chapter 4;Chapter 10

⁴ It doesn't follow that they are polysemous in semantic mapping function or that they correspond to more than one kind of truth condition. I will return to this soon.

Donnellan's claim that for one use of definite description it doesn't matter whether the description is empty so long as the speaker's intended referent is understood by the hearer is best interpreted as a claim about speaker meaning rather than linguistic meaning. Certainly it is not a stabilizing function of definite descriptions to bring things not correctly described by them to hearer's minds.

below). Then understanding the force of a particular utterance using the form will require disambiguation from context, just as understanding the meanings of individual words often does.

II. Semantic mapping functions (and satisfaction conditions)

What sets linguistic acts and, more generally, communicative acts apart from other acts with cooperative functions is that communicative forms work in part by mapping or, as Wittgenstein put it, "picturing." They correspond to states of affairs in accordance with semantic mapping functions that have been determined by convention.⁵ Directive communicative forms have as their stabilizing functions to yield states of affairs --completed actions-- that vary with variations in the sentences exhibiting these forms. For example, directive least types used in giving orders have as stabilizing functions to produce compliance, what constitutes compliance being determined along the lines Tarski proposed. The state of affairs that would result from compliance is the "satisfaction condition" of the directive sentence. Descriptive communicative forms have stabilizing functions that can be performed through normal mechanisms only if they correspond to states of affairs existing independently. For example, conventional fact-stating least types are designed to produce true beliefs in hearers, but a true belief will be formed by normal mechanisms only if the sentence corresponds to a world affair in accordance with its conventional semantic mapping function. False sentences do not

⁵ In the case of most animal communication, determined genetically.

cause true beliefs in hearers through normal mechanisms. The truth conditions of a descriptive sentence are also called "satisfaction conditions."

The semantic mapping function for a sentence determines the sentence's satisfaction condition, but the satisfaction condition does not determine the semantic mapping function. The semantic mapping function is given by rules according to which significant transformations of the sentence that conserve its syntactic form yield different truth or satisfaction conditions. Compare the sentence "It's raining" with the sentence "Rain is falling here now." "It's raining" contrasts with "It's snowing," "It's hailing," "It's sleeting," and so forth. All display the same syntactic form, the transformations that substitute "snow," "hail," and "sleet" for "rain" determining different satisfaction conditions in a systematic way. Similarly, "Rain is falling here now" contrasts with "Snow is falling here now," "Hail is falling here now," "Sleet is falling here now" and so forth, but it contrasts, further, with "Mist is rising here now," and with "Rain was falling in Rome yesterday." The truth conditions of "It's raining" and of "Rain is falling here now" are the same, but the semantic mapping is different. "Many drops of water are presently precipitating from the atmosphere and landing close to this place" also has the same truth condition but is articulated by yet another semantic mapping function. For vividness, compare the semantic mapping function for a bee dance with that of an English sentence having the same truth condition. Bee dances show by the angle of their axis where there is nectar relative to a line between the hive and the sun, but there are no transformations of the bee dance that would tell about nectar location relative to objects other than the hive and the sun, or about the location of anything other than nectar. Only reference to

the angle between the nectar and the line from the hive to the sun can be varied in the bee dance. Further, the bee dance is not subject to a negation transformation. No English sentence with the same truth conditions approaches this degree of inarticulateness. The semantic mapping of a sentence articulates it, placing it in a logical space of contrasting possibilities. Its truth condition is not, as such, articulated.

Stabilizing functions can vary while semantic mapping remains the same. Compare "Jane will close the door" to "Will Jane close the door?" And there are more interesting examples. Wilfrid Sellars claimed that the function of the sentence form "'X' means Y," as in "'Rot' means red" and "'und' means and," is to produce in the hearer a disposition to use "X" in the same way he already knows to use "Y." The "Y" in this rubric, Sellars said, is neither mentioned nor used in the usual way (Sellars 1956). It is used in a special way, held up, as it were, as a model (Millikan 2004 Chapter 7). Compare the function of the form "'X' and 'Y' are used the same way." Here "Y" is mentioned rather than used. This sentence has the same truth condition as "'X' means Y" but its function is different. Its function is to produce a belief about words whereas the function of "'X' means Y" can be performed even if the hearer lacks a concept of words (as very young children apparently do --Susan Carey, private correspondence).

Peter Strawson claimed that the function of the identity form "A is B," as in "Cicero is Tully," is to induce the hearer to merge all of the information he has accumulated under the concept he associates with the word "A" with the information he has accumulated under the concept he associates with "B," so that he no longer harbors this information under two separate concepts (Strawson 1974). More accurately, the stabilizing function

must be to induce the hearer to do this appropriately, such that the resulting concept is not confused or equivocal (Millikan 2000). If this is the function of the form "A is B," then its truth condition is the same as for the form "'A' has the same referent as 'B'," in which the "A" and the "B" are mentioned rather than used. But these two sentence forms do not have the same function. The hearer of "'A' has the same referent as 'B'" is to form a belief about words, hence needs concepts of words and also the concept of reference, whereas the hearer of "A is B" needs neither.

I have claimed (Millikan 1984, chapter 12) that the stabilizing function of the form "A does not exist" is correctly to induce the hearer to disengage his concept associated with "A" from ordinary referential uses, relegating it, for example, to pretend uses, or eliminating it entirely from his conceptual repertoire. Correlatively, the function of "A exists" is correctly to engage a previously disengaged concept associated with "A". But if these are their functions, the sentence forms "A does not exist" and "A exists" have the same truth conditions as do "'A' has no referent" and "'A' has a referent" though, again, the functions of these sentences are not the same. The latter have as stabilizing functions to cause beliefs about words.

Adding a different kind of example, the two uses of definite descriptions mentioned above in connection with Donnellan's distinction may correspond to two independent stabilizing functions of these, but these two uses require exactly the same conditions for truth. The world affair needed to make one least type of definite description serve its stabilizing function through normal mechanisms is exactly the same

as that needed for the other. In both cases the truth condition is Russellian.⁶

The study of semantic mapping functions should include a study of the peculiarities of indexicals and demonstratives. As I understand it, there is more than one way of describing their semantic mapping functions and I have suggested somewhat different though, I believe, compatible ways of thinking about the matter in (1984 chapter 10) and (2004 chapter 12). I omit discussion of these forms here, my purpose being only to make clear why we need to distinguish among the three broad aspects of meaning mentioned at the start.

III. Conceptions

⁶ See footnote 4 above.

I think that Frege made a mistake in positing something common beyond Bedeutung that is grasped by the mind of every competent speaker using the same unambiguous linguistic form. A related mistake suffuses the tradition of conceptual analysis in seeking shared "criteria" for the correct application of various terms, criteria taken to be learned, in some mysterious way, when one learns one's language. On the contrary, the public meaning of a simple referential term typically includes only its stabilizing function and its reference, and since the stabilizing function depends almost entirely on sentential context, the public meaning is essentially just reference.⁷ I intend this sweeping assertion to include terms for properties, kinds, stuffs, and so forth, which I will treat here as also being, in a broad sense, referential. The claim will need qualifications, but first I'll just try to explain it.

The idea to be opposed is that for different users to understand the same referential term as having the same meaning requires that their psychological processing be similar in certain ways. The idea to be defended is that "Meaning is not in the head." But I want to launch an attack on Frege's idea and on conceptual analysis that is more radical and exhaustive than the familiar offenses once launched by Putnam, Burge and Evans. My argument grows out of a view about thought structure, a view about what it is

⁷ See (Millikan 1984) chapter 4 on the most general stabilizing function of all referential terms.

to have a concept of a property, an individual, a kind or a stuff and so forth. It is fully articulated in (Millikan 1984, 2000). Here I can make only a small sketch.

Consider what is involved in being able to recognize, for example, shapes. Think of the variety of proximal visual stimulations to which a given shape may give rise when viewed from various angles, from different distances, under different lighting conditions, through various media such as water or fog, when colored different ways, when partially occluded and so forth. How shape constancy is achieved by the visual system, the capacity to recognize the same shape as the same under a range of conditions, is a problem of nearly unimaginable complexity that is still largely unsolved.

Similarly, how color constancy, texture constancy, size constancy and distance constancy are achieved are enormously complicated problems. We are also adept at perceiving sounds, especially speech sounds, as the same sound at origin whether near or far, through air or through water, muffled or distorted and so forth. What does seem clear in each of these cases is that no single rule is applied. Different clues are used by the perceptual systems in different circumstances, separately or together. For example, distance is perceived with the help at least of ocular disparity, tension in the focusing muscles, occlusion of one object by another, knowledge of the size of objects viewed, and atmospheric haze. We also recognize distances by touch and stretch using many different parts of the body, and we recognize distances of things that make noises fairly well by ear. And of course there are more readily noticed ways of recognizing distances as well, for example, by measuring with a ruler or a tape measure or just a string, or measuring as a surveyor does by triangulation, or measuring with an odometer or a

micrometer or by the time of the return of light... None of these ways of telling distances is infallible nor is any definitional of our concepts of distances. On the other hand, each adds something to our concepts of distances, nor could we have distance concepts at all were we not in command at least of some of these methods of recognition.

The situation is similar, if not always so extreme, with our grasp of other perceptual constancies. The perceptual systems do their work in flagrant violation of the ideal once set by champions of operational definitions. The more ways the better when it comes to methods of perceiving a property. After all, the ways in which empirical properties affect the various senses through intervening media is a thoroughly empirical matter, a question of natural law, not a matter of logic or definition. That is why neither phenomenalism nor verificationism could ultimately survive.

Now it is conceivable that all normal persons perceive some constancies, for example depth, in the same way, conceivable even that they are genetically programmed rather than perceptually tuned to perceive some of these constancies in standard ways. The issue is under debate. But surely whether one's perceptual capacities were entirely normal in this regard would not affect what one meant by the English words one used in designating depths or shapes or textures. Being blind in one eye so that one could not perceive depth using ocular disparity would not change what one meant by "near" and "far" nor, indeed, is it sensible to suppose that Helen Keller meant something different by "near" and "far" than you do. I don't want to debate about whether there are secondary

qualities,⁸ but surely we recognize in perception enough properties and relations that obviously are primary to make the point. There are many different ways of recognizing each of these properties, but none defines either the property or the words that stand for it.

Turning now to the opposite extreme, consider proper names. Besides having a referent does your name have a definition? What is involved in someone's understanding who's meant by your name --say, a child in your family, your child's teacher, a student of yours, the student's wife, a reader of your essays, the pharmacist who fills your prescriptions. Do these people all understand who's meant by your name in the same way? The reasonable answer is that there is no special thing common in the minds of all people who understand your name except, I have argued (Millikan 1984 chapters 4 and 9, 2000 chapter 6), some practical capacity to reidentify as such, in actual context, the least type that is your name (rather than the name of someone else with "the same name"), so as to recognize when information is being offered about the same person, you, again. Speaking more generally, what it is to have a concept of an individual is, in

⁸ The very fact that psychologist's can make a study of how color constancy is perceived seems to cast doubt on the idea that colors are secondary qualities, anyway in the Lockean sense of that term.

part, to have an ability to recognize, in one way or another, under at least some circumstances, when one is encountering information concerning that individual, and one recognizable way that one encounters information about a thing, besides through direct perception, is by encountering sentences that contain its name. To defend this position properly, the right characterization of information is needed (Millikan 2004, chapters 3-4) and the right story about perception through language (Millikan 1984 chapter 9, 2000 chapter 6, 2004 chapter 9) as well as a story about abilities that allows for their fallibility (Millikan 2000 chapter 4). But that names of individuals need not be associated either with independent publicly agreed on ways of recognizing these individuals in order to do their work or with agreed on descriptions associated with these individuals is generally accepted, I think.

Call the sum of the various ways that you have of recognizing a thing or, what amounts to the same, of recognizing when you are receiving information about a thing, your "conception" of that thing. Your conceptions of most common things have many components, for you have many ways of recognizing these things --no infallible ways, of course, but many fairly reliable ways. Whatever you know about a thing is part of your conception of it too, for whatever you know might help you to identify it, or help prevent you from misidentifying it, under some circumstances. Some components of conceptions are explicit, involving the use of descriptions hence of prior concepts in their application. Other conceptual components are implicit, moving one directly from perceptual experience to an identification of what is perceived. My claim so far is that neither the names of perceivable properties nor the names of individuals are associated with

conceptions or conceptual components, either explicit or implicit, that all users of their names must possess in order correctly to understand these names. No specific way of identifying their referents is required. True, in some cases there does exist considerable overlap in the conceptions that most people use in understanding the referent of such a name, for example, since "Mark Twain" was Samuel Clemens' pen name, a large proportion of people know that Mark Twain was a writer, perhaps even that he wrote Huckleberry Finn. You can usually count on someone having that knowledge if the name "Mark Twain" is in their vocabulary. And the implicit conceptual components by which we recognize many common properties may be shared among a large majority of adults. But if someone were born with bats ears and could only hear shapes, this would not prevent him from learning the English words "round" and "square." I take it that this much is not highly controversial. The principle can be extended, however, to less obvious cases.

Many terms for kinds name kinds that are objective natural units, discovered rather than created by thought and language (Millikan 1984 chapters 16-17, 2000 chapter 2). These "real" kinds are important subjects for knowledge because there is a reason why the various members of the kind mostly resemble one another in a good number of ways, hence there is a reason why one can learn from observation of one or a few examples of the kind much that is likely to be true of other members. Most single terms designating kinds designate real kinds of this sort (Millikan 2000 chapter 3). Typically, these kinds not only have many properties, there are also many ways to recognize them. Think how many ways there are of telling that something is copper, or that a dog is present. Do you have to look to tell it's a lemon? Or that it's raining? How

much of what portion of The First Noel or The Lord's Prayer do you have to hear to recognize it? To have a valid concept of a real kind one needn't know the reason for the resemblance of its members, what natural principles hold the kind together. One only needs some fairly reliable ways of reidentifying the kind, the more the better, of course, since most ways are only applicable on some occasions. Like concepts of individuals, concepts of real kinds can be supported by alternative conceptions, alternative methods of recognition, and there are no conceptual components that all users of a real kind's name must possess to understand it (Millikan 1984 chapter 9, 2000 chapters 3 and 5).

The third aspect of meaning, conception, is not then essentially public. It attaches in the first instance to idiolects rather than public languages. However, there usually is considerable overlap among people's conceptions corresponding to names of very common real kinds. Also, sometimes conceptual components are passed on explicitly from generation to generation, for example, the definitions of certain geometrical figures. One could, after all, "define" a circle, instead, as a closed plane figure with but one side of uniform curvature, but it is not conventional to do so. And in the case of fictional names, and empty names like "phlogiston" and "witch" when these are mistakenly thought to have referents, there is no public meaning beyond certain traditional explicit conceptual components, traditional descriptions, passed down from person to person. There is no more to public meaning in these cases than public conception, indeed, public conception that is highly subject to drift. Santa Clause acquired red and white attire and reindeer rather late in his career while phlogiston and witches took on different diagnostic properties over time in the eyes of different investigators. Water and dogs, by contrast,

are surely recognized by us, practically all of the time, in exactly the same way they were by the ancients.

Traditional descriptions associated with empty terms fail to reach anything real, hence do not correspond to real abilities to identify. But having empirical concepts, having thoughts of objects, properties and so forth, essentially involves abilities to identify. It seems to follow that empty terms do not express real concepts. This brings us to the externalist core of this chapter on meaning.

The claim is that the meaning of an empirical term is, in the first instance, its referring to something, and only in the second instance, ways one has of identifying this thing through various of its manifestations. Wittgenstein was right, after all, that the primary check on whether we mean the same by our words is agreement in judgments, but agreement in judgments proves nothing about agreement in the methods of identifying used in making those judgments. It is clear that nothing inside the head or mind can determine, in and of itself, whether one's dispositions to react to sensory stimulations with would-be thoughts of individuals, properties and kinds manifest real abilities to identify such things or not. Similarly, nothing inside the head determines whether those explicit inference dispositions whose job is to help one identify such things are actually doing their job. But if a would-be thought has conceptual components that are explicit, and if the prior concepts in the descriptions employed in these components are not themselves empty, then there is a legitimate, though secondary, sense in which even a term expressing an empty thought can have a meaning ---because components of its conception have meanings. Indeed, if the term is public, it will have conceptual

components that are both explicit and traditional, having been handed down from speaker to speaker, hence it will have a sort of public meaning. But suppose there were an empty concept that had only implicit conceptual components, that was not anchored by any prior nonempty terms. Such a term would have no more claim to membership in the realm of the intentional or the semantic than a sneeze. It would merely be a quirkish regular response to certain sensory stimulations, resulting, presumably, from the faulty operation of systems designed to design genuine concepts, genuine thoughts, through experience, but that had failed in that task.

A crucial task incumbent on any advocate of meaning externalism is to explain how we acquire evidence through experience that our concepts are not empty, that they are anchored externally in what is objectively real. The externalist is obliged to accompany claims about the ontology of meaning with a plausible epistemology of adequacy for empirical concepts. She must construct an epistemology of meaning to support her claims in the philosophy of mind. I consider this an urgent matter, though one sorely neglected in the current literature on externalism. The epistemology of concepts, or of meaning, is the subject of (Millikan 1984 chapters 18-19, 2000 Chapter 7, 2004 chapter 19).

IV Replacing intensions and Fregean senses

As said at the start, none of my trio of meanings corresponds at all well either to any traditional notion of intension or to any Frege-like notion of sense. Both these latter notions were introduced on the assumption that a grasp of certain ways of identifying or certain properties by which a thing may be identified must be shared by users of any

public term that refers to it, whereas I claim that any such grasp is, in the first instance, a private matter. But this claim needs to be defended with an alternative explanation of the phenomena that lead to the postulation of intensions and Fregean senses. I take it that there are three central classical arguments for something like intensions or senses, one from the informativeness of sentences asserting identity, a second from the need to analyze statements asserting existence, the third from the behavior of referential terms in intensional contexts. I will briefly discuss each in turn.⁹

First, identity statements. I have agreed with Strawson that the stabilizing function of an identity sentence, "A is B" is to encourage the hearer to merge under a single concept all of the information she has accumulated under the concept she associates with the word "A" with that under the concept for "B." More precisely, it serves its stabilizing function by joining the conception the hearer has associated with "A" to the conception associated with "B" so that these now (correctly) govern the same concept. Rather than inducing beliefs --compare beliefs to mental sentences-- it alters conceptions, ways of identifying (Millikan 2000 chapter 12). Thus for any hearer who

⁹ There is a fourth argument (that is not classical) from the need to understand how intentional attitudes are to be described so as to play their usual role in psychological explanation. This need is discussed in my (2000) chapter 12.

associates a different conception with "A" than with "B," the effect of a true identity statement "A is B" obviously is different from that of "A is A." This can be true and important even if no two hearers who react in the stabilizing way to "A is B" happen to share their conceptions associated with "A" or with "B."

Second, existence statements. Statements of the form "A doesn't exist," I have claimed, induce a hearer to disengage his concept associated with "A" from ordinary referential uses, eliminating it entirely or reserving it only for pretend uses. Statements asserting the existence of A reverse this effect. To engage or disengage a concept is the same as to engage or disengage the conception that governs that concept. The forms "A exists" and "A doesn't exist" can serve these functions regardless of how diverse listeners are in the conceptions they associate with the name "A." But as noted earlier, there is usually a good deal of overlap in conceptions for names that are very common, and names that have no referents can only be passed on by description, so they are especially likely to have conceptions that are largely public (though perhaps shifting).

Third, intensional contexts. A well known way of extensionalizing intensional contexts was suggested by Davidson in "On Saying That" (1968-9). His idea was that a sentence such as "Galileo said that the earth moves" is true just in case uttering the words inside the "that..." clause of this sentence makes the speaker and Galileo into "same sayers," people who have uttered words with the same import. I have adopted a similar view but generalized it, claiming that when one representation is held up or put on display in order to show what another representation is like, the kind of similarity intended may concern any aspect of meaning, or may even concern some aspect of the

vehicle of the displayed expression (Millikan 1884 chapter 13, 2004 chapter 7). For an example of the latter, consider "John kept insisting that it wasn't a woodchuck but rather a groundhog!" It is clear here that the similarity intended must concern the very words "groundhog" and "woodchuck" since these two are names for the same.

Intentional attitude contexts yield to a similar analysis. In "John firmly believed that it was not a woodchuck but a groundhog," again something about the words "woodchuck" and "groundhog" is surely at stake. One way to understand this is to assume that an embedded sentence displayed in an intentional attitude context refers to an intentional attitude that is relevantly like one it would be its own stabilizing function to produce. The sentence "It was not a woodchuck but a groundhog," if it were to serve its stabilizing function with John as a hearer, would produce just the mental state John is in, right down to the last conceptual component. That the message concerns not merely some proposition associated with John's mental state but also his conceptions, including the very words through which he would try to recognize information coming in about the subjects of his thought, is clear on the (pragmatic) assumption that John does not think a thought that shows, from the inside, that it is contradictory. This reading also nicely accommodates the fact that definite descriptions appearing inside intentional attitude contexts are sometimes read as attributing the description to the thinker as part of his conception and sometimes as attributing to him only a thought of the description's referent. Thus "Ralph thought that our venerable dean was a spy" might or might not imply that Ralph knew that the man he thought was a spy is our dean. This is entirely natural if the intentional attitude description works by displaying a sentence whose

function is to produce an attitude like the one being attributed and if definite descriptions have alternative stabilizing functions corresponding to Donnellan's distinction, as discussed above.¹⁰

The three aspects of meaning that I have discussed are thus sufficient to account for those properties of natural language traditionally associated with meaning.

¹⁰ To accord with this treatment, modal contexts need to be understood as contexts in which representations, rather than possible situations or worlds, are the basic subject matter.

References

- Boër, S.E. and W.G. Lycan 1986. Knowing Who (Cambridge MA: MIT Press)
- Bach, K. and R. M. Harnish 1979 Linguistic communication and speech acts (Cambridge MA: The MIT Press).
- Davidson D 1968-9 "On Saying That" Synthese 19:130-146. Reprinted in A.P. Martinich ed. The Philosophy of Language, fourth edition (Oxford: Oxford University Press)
- Donnellan K 1966 "Reference and definite description" Philosophical Review 75:281-304
- Gilbert, M. 1983 "Agreements, conventions, and language," Synthese 54:375-407.
- Gilbert, M. 1989/1992 On social facts (Princeton: Princeton University Press 1992, reprinted from Routledge 1989 edition).
- Grice, H. P. 1968, "Utterer's Meaning, Sentence-Meaning and Word-Meaning" In Studies in the Way of Words (Cambridge MA: Harvard University press).
- Lewis, D. 1969 Convention: A Philosophical Study (Oxford: Basil Blackwell).
- Lewis, D. 1975 "Languages and Language". In K. Gunderson, ed., Language, mind and knowledge (Minneapolis: University of Minnesota Press). Reprinted in H. Geirsson and M. Losonsky eds., Readings in Language and Mind (Oxford: Blackwell) 1996, 134-155.
- Millikan, R.G. 1984 Language, Thought and Other Biological Categories (Cambridge MA: The MIT Press).
- Millikan, R.G. 1998 "Proper Function and Convention in Speech Acts," in ed. L.E. Hahn, The Philosophy of Peter F. Strawson, The Library of Living Philosophers (LaSalle IL: Open Court) 25-43.

- Recanati, F. 1987 Meaning and force: the pragmatics of performative utterances
(Cambridge: Cambridge University Press).
- Searle, J. 1969. Speech acts: an essay in the philosophy of language. Cambridge:
Cambridge University Press.
- Sellars, W. 1956. "Empiricism and the philosophy of mind." In Minnesota studies in the
philosophy of science, vol. I, ed. K. Gunderson, 253-329. Minneapolis: University
of Minnesota Press. Reprinted in (Sellars 1963).
- Sellars, W. 1963 Science, Perception and Reality (New York: The Humanities Press).
- Schiffer, S. 1972. Meaning. Oxford: Clarendon.
- Strawson. P.F. 1974 Subject and Predicate in Logic and Grammar. London: Methuen
and Co.