

Chapter Six
Substance concepts through language:
Knowing the meanings of words

§6.1 Perceiving the World Through Language

As I have described substance concepts, having these need not depend on knowing words. Preverbal humans, indeed, any animal that collects practical knowledge over time of how to relate to specific substances needs to have concepts of these. On the other hand, it is clear that language interacts with substance concepts in vigorous ways, completely transforming the conceptual repertoire. Putnam argued for what he called "the division of linguistic labor," according to which laymen can borrow on the concepts of experts. Though offering an entirely different analysis, I will conclude similarly, that the public language plays a crucial role both in the acquisition of substance concepts and also in their completed structure.

The story so far about substance concepts seems to collide with the obvious fact that many of these concepts, both for children and adults, have been acquired without encountering the substances "themselves" but only by hearing about them. With regard to these very same substances, moreover, we are often in the position that Kripke (1972) and Putnam (1975) observed, having no unique descriptions of them in mind either, so that descriptionist theories of how extensions are determined also do not fit these cases. I will argue that this entire problem falls away if we view speech as a direct medium for the perception of objects.

It is traditional to assume that gathering information by being told things is a radically different sort of process from gathering information directly through perception. But there is reason to think that this difference has been greatly exaggerated.¹ In fact, uncritically believing what you hear said may be surprisingly like uncritically believing what you see. For example, there is experimental evidence that what one is told goes directly into belief unless cognitive work is done to prevent this, just as what one perceives through other media does. Loading the cognitive systems with other tasks, such as having simultaneously to count backwards by threes, has the effect of facilitating belief fixation regarding whatever one hears or reads (Gilbert 1993).

There are two things that distinguish direct perception quite sharply from the acquisition of information through language, but neither implies a difference in immediacy. In direct perception, the spatial and temporal relation of the perceiver to the object perceived is given, whereas it is not normally given through language. If you see the cat, you normally see also its spatial relation to you, and whatever you perceive it doing is done at the time you perceive it. But if you hear John talking about Xavier, you do not usually hear about Xavier's spatial relation to you, nor is it automatic to know what temporal relation the Xavier-doings that John relates bear to you. There are intermediate cases, however, between ordinary perception and gathering information through

¹Gareth Evans may have had a view similar to the one I present below, though his notion of "information" was different. See Appendix A and Appendix B.

language. For example, when watching television, the spatial relation of perceiver to perceived is not given either, nor, unless the program is live, is the temporal relation, yet one perceives that the newscaster frowns or smiles just as immediately as one would in his presence. So this alone is not a reason to distinguish perception sharply from learning about the world through language.

The second feature that distinguishes perception is its near infallibility. It is remarkably difficult to deceive people about what they are actually seeing or hearing. This is why "seeing is believing." On the other hand, given a modern understanding of the mechanisms of perception and a substantial technology it is possible to manage materially to fool the human ear and eye. False appearances can be arranged in the laboratory. And false appearances are now easily arranged using modern communications media. Though generally overlooked in this connection, the latter offer much the most common current illustration of the persistence of perceptual illusion. After seeing her daddy on television, the small daughter of a friend asked him, "Daddy, how did you get in there?" But I mainly have in mind more radical cases such as dubbed films and cartoons.

In a similar way, persistent illusions are easily arranged through language and they are abundant. That is, sentences are often false, and even when you know they are false, they continue to present the same false appearances. They do not shift and appear to say something different. In water, oars look bent and the reflections of the trees show them moving in ripples. We are not, however, tempted to believe that our oars are bent or that the trees are moving in ripples. Similarly, "I'm dying" uttered by a laughing 8 year old does not tempt us to believe that someone is dying (Gendler1998). But the appearance is as much of someone dying as the appearance in the rippling water is of rippling trees. Perhaps you will say, but it doesn't sound as if the laughing eight year old is dying. True, those aren't the sounds an eight year old would make if she were really dying. It also isn't the look a tree would have if it were really rippling. Trees don't flex that way. Still, the persistent illusion is that the trees are flexing that way and that the eight year old is dying. In sum, hearing sentences may be quite a lot like watching the media, or like watching reflections, which in turn is quite a lot like watching the original.

Think of the matter this way. There are many ways to recognize, for example, rain. There is a way that rain feels when it falls on you, and a way that it looks out the window. There is a way that it sounds falling on the rooftop, "retetetetetet," and a way that it sounds falling on the ground, "shshshshsh." And falling on English speakers, here is another way it can sound: "Hey, guys, it's raining!"² Nor should you object that it is not rain you hear in the last case but rather "a sentence." Or a sound? Is it then a sound that you hear rather than rain on the roof? Is it a TV screen that you see rather than Bill Clinton? A pattern of ambient light rather than the TV screen? Best of all, perhaps all you see is a visual impression? Which ones of these things are the real or direct objects of perception?

² Thank you, Crawford Elder

You can, if you like, hear or see any of these things. What you see when you look depends, first, on where you focus your eyes; it depends, second, on where you focus your mind, your attention. True, philosophical tradition, and the psychological tradition following after, has resolutely held that for each of the physical senses there is just one layer of the world that it perceives directly; all other layers are known only through inference. This premise I am denying. There is no single "given" layer of perception (again, see MacLennan 1998, quoted in §5.5 above). This, of course, was argued strongly in the broad tradition that includes both Wilfrid Sellars and Willard Quine, so it should not be an unfamiliar idea. Perception, many have held, is "theory laden." Perceptual judgments do not arise through inference, but neither is the content of a perceptual judgment an epistemological given. I would argue that the Sellars/Quine tradition was mistaken in the reasons they gave for this conclusion. Their view was that applying a concept is making a transition from stimulation or sensation into thought, and that the content of the thought produced is defined by the network of inference dispositions (nowadays, the "cognitive role") in which the concept is enmeshed. I am proposing a different theory about the content of a concept and about how its extension is determined. But the conclusion is the same. The substances referred to in perceptual judgments are not epistemological givens but are discovered through a process of fallible construction, fallible learning. They are distal objects, and there is no necessary restriction on their level of distality.

According to the contemporary motor theory of speech perception (Liberman & Mattingly 1985, 1989; Mattingly & Studdart-Kennedy 1991), phonemes are not sounds, not acoustic phenomena, but gestures made by the vocal tract. That is, what you hear as the same phoneme again is not acoustically the same, but is the same movement or same posture aimed at by the vocal tract. Furthermore, the processing of speech sounds, when these are perceived as speech, is through different channels than the processing of other sounds. Thus the end organs of perception, the ears, determine more than one mode of perception. My further suggestion is that when one is listening to speech in order to gain information, the ears hear not just through the acoustics to the speech gestures, but through the speech gestures to the world. What the young child perceives in the presence of speech is not sounds, nor phoneme strings, nor words but, in the first instance, the world. (If a bat can hear that something is square, so can I. It's only fair.)

The child comes into the world without any knowledge of how minds work, without any knowledge of what goes on inside people when they speak (indeed, we ourselves seem to be a bit short on such knowledge). The child does not develop concepts, for example, of beliefs and desires, until several years after the onset of speech. It is clear then that children cannot possibly understand language in the way Grice (1957) described, by understanding that the speaker intends them to believe...and so forth.³ But it is also true that the young child has very little phonological awareness. Indeed, contrary to much public opinion, the

³ For a critique of Grice, see (Millikan 1984), Chapter 3.

difficulty of becoming aware of structure at the phonological level is probably the most common first cause of dyslexia (Morais et al 1979; I.Y. Liberman et al 1984; Lundberg et al 1988). Clearly the child has no concepts of phonemes and cannot understand speech by drawing inferences from the patterns of phonemes it hears.

It is also true that children learn very few words by ostentation. They learn them by hearing complete sentences containing them (e.g., Gleitman 1990, Grimshaw 1994, Pinker 1994a). For the young child, language serves simply as another medium of perception, a medium through which to perceive the world—exactly as the child perceives the world through its eyes without knowing anything about light, and through its sense of touch and smell without knowing anything about physical forces or chemicals. How can that be, you may say, since Mamma's words are right here while the dog she talks about is way over there? Well, how do you perceive yourself in the mirror? What's funny about language, I have said, is that it does not show your own relation to the things you perceive through it.

§6.2 Tracking Through Words: Concepts Entirely Through Language

But there really is no need to exhaust this point here. In the present context, the part that really matters is that believing what one hears said is a way of picking up information about substances, and that it is by learning a language that a child becomes able to pick up information in this way.⁴ It sounds a bit queer to speak of learning a word for a substance as learning a way to identify that substance. But just as the relation of one part of the pattern on the TV screen to another part can manifest the relation of one part of Bill Clinton to another, the relation of a word to other words in a sentence can manifest the configuration of a substance in relation to other substances and properties in the world. The semantics of natural languages is productive; alterations performed upon sentences correspond systematically to alterations in what the sentences represent, just as in the case of pictures, though the mapping functions involved are, of course, far more abstract. So if learning what a substance looks like can be learning how to identify it, similarly, learning a word for a substance can be learning a way to identify it. In both cases, what one learns is to recognize or understand manifestations of the substance as manifestations of it; one learns how to translate information arriving in one more kind of package at one's sensory surfaces into beliefs.

Learning a language is, in part, just learning more ways to pick up information through the senses and put it away in the right boxes. A difference, of course, is that this way of picking up information is much more fallible than in the case of ordinary perception. But no human ability is infallible. Furthermore, just as substances are sometimes look-alikes in the flesh (twin brothers), many substances are sound-alikes in words (John_(Doe) and John_(Roe)). But substances are tracked through the medium of words not merely by means of the same words manifesting the same substances. Like more direct manifestations of substances, words and sentences occur in context, allowing methods of tracking

⁴ "Information," in this passage, means informationC as defined in Appendix B.

to be used that are analogous to more ordinary tracking, in that they rely in large part on expected spatial, temporal and causal relations (c.f., trajectory) rather than the persistence of properties. How do I recognize that as John's elbow poking out over there behind the lamp? Well, I saw John head that way with a book just a moment ago. Some of these relations are natural, as the natural relation between a speaker's experience plus the context of his speech to the subject of the information he is trying to convey. One will usually know which "John" a speaker is talking about in a way analogous to the way one knows whose elbow that is. Other relations between word and referent are governed more closely by convention, as in the interpretation of certain anaphoric pronouns and certain indexicals.

Recognizing a linguistic reference to a substance is just another way of reidentifying the substance itself. It is identifying it through one more medium of manifestation. Think of this medium as like an instrument that aids perception. Like a camera, a radio, a cat scan, or a microscope, another person who talks to me picks up information-bearing patterns from his environment, focuses them, translates them into a new medium and beams them at me. Or think of living in a language community as like being inundated in one more sea of ambient energy. Like the surrounding light, surrounding people transmit the structure of the environment to me in ways that, barring certain interferences, I can become tuned to interpret.

It is even possible, indeed it is common, to have a substance concept entirely through the medium of language. It is possible to have it, that is, while lacking any ability to recognize the substance in the flesh. For most of us, that is how we have a concept of Aristotle, of molybdenum and, say, of African dormice. —There, I just handed you a concept of African dormice, in case you had none before. Now you can think of them nights if you like, wondering what they are like—on the assumption, of course, that you gathered from their name what sorts of questions you might reasonably ask about them (animal questions, not vegetable or mineral or social artifact questions). In many cases there is not much more to having a substance concept than having a word. To have a word is to have a handle on tracking a substance via manifestations of it produced in a particular language community. For the person who remembers faces easily, one look at a new person may be enough to implant the ability to recognize that person again, thus enabling a concept of them. For the person who remembers words easily, one hearing of a new substance through a word for it may be enough to implant the ability to recognize the substance again through that word, thus enabling a concept of it. Simply grasping the phonemic structure of a language and the rudiments of how to parse it enables one to help oneself to an embryo concept of every substance named in that language. It enables one conceptually to track these substances and easily to discover under what sorts of substance templates they fall. That, I suppose, is why it is possible for small children to learn a new word every hour (§5.1).

The basic phenomenon here is the same as that underlying Putnam's theory of the "Division of Linguistic Labor" (1975) and Burge's claim that constitution of the very content of ones thought sometimes passes through the

word usages of a surrounding language community (1979, 1982, 1986). But the explanation I am proposing of this phenomenon is quite different. The image created by both Putnam and Burge is that when I have a concept through language I take out a loan that the experts are prepared to pay up. There are experts out there who "really" have the concept while the rest of us really don't. But even if we soften this just to the claim that some people out there have (or had—consider our concept of Socrates) the concept in a way that was different from ours because focused without reliance on public language, still the image is wrong. My claim is that having a concept grounded only through language is no different than having a concept grounded only through, say, vision. Such a concept is in no way secondary. True, others must help me to have such a concept, just as a television may have to help me if I am to see Bill Clinton. But just as I really do see Bill Clinton on television, having a concept through language is really having a concept. It is really thinking of something.

§6.3 Focusing Reference and Knowing the Meanings of Words

Words serve in huge numbers as seed crystals around which fuller conceptions of substances are then quickly formed. That is why there can be such differences between the concepts available in cultures not historically related, and why poor Helen Keller was, as she later described it, pretty much unable to think until Sullivan taught her some language. Gelman and Coley (1991) are surely right that "a word can serve to stake out a new category, which then must be explored in more depth" (p. 184; see also Gopnik & Meltzoff 1993). Words also are handles to hang onto, helping to stabilize concepts so as gradually to eliminate equivocation in thought, granted that those who speak to us have unequivocal concepts themselves.

Acquiring adequate substance concepts involves learning to focus one's thought, such that all of the incoming information scattered over time about each substance is put into one slot, and the right constancies projected for it. Learning to do this is what Perner called "focusing reference" (§4.8) Learning words for substances is in part a matter of focusing reference. Substances are tracked through words and also in other ways. If information about a substance arrives through language and the substance is tracked through a word, but the information culled in this way is put under a concept used also to track a different substance using other means, there will be equivocation in the resulting concept. We say in such cases that the person "does not know the meaning of the word," or thinks that the word means something different than it really does. A perfectly parallel case would be mixing a person known to you only through phone calls with a different person known from glimpses at the beach. One could just as well say, using the same sense of "meaning," that they did not know the meaning of the voice over the phone.

In Chapter One I suggested that preschoolers who take tigers to be "kitties" may be confused, not about the meaning of the word "kitty," but about how to identify housecats. From our present perspective, however, thinking tigers are "kitties," that is, putting tiger information away in the same slot as information gotten from observing housecats and from hearing about "kitties," is being confused about tigers as well as about housecats. The child has not yet

managed to focus on only one substance. Perhaps the child calls the whole genus Felis by the name "kitty." It does not follow that the child means Felis by "kitty." Rather, the child's word "kitty" may hover between referring to felines generally and housecats specifically. The child may be putting all information gleaned through language and specific to housecats in the same bin as information gleaned about tigers and lions at the zoo. The child's conception of "kitty" will then be equivocal, part of it tracking felines generally while another part is channeled through the word "kitty," hence is bringing in information much of which is wrong for felines generally. Then the child's referent is not the same as the referent of the English word "kitty," so it is certainly true that child does not yet know the meaning of "kitty." But the public word "Felis" is not equivocal, so the child does not mean Felis by "kitty." (Suppose, on the other hand, that a foreigner uses the word "kitchen" to refer to chickens. It is very unlikely that she will have gathered in any information through language about kitchens and actually put it in her mental chicken bin, or that she has any disposition to do so. This is an entirely different case from that of the child who calls tigers "kitty.")

Because it is possible for a conception to be channeled completely through language, it is possible to have a substance concept through nothing but a word plus a grasp of its substance template and enough relevant grammar. Many people find this completely unintuitive, and I sympathize. But my point is that filling out the concept into a more and more adequate one happens in degrees. There is no special thing that gets added at some later point that suddenly makes it into a "real concept." It can be filled out more; it can get better and better. But there is no magic moment when it has attained some essence required for true concepthood. There is no magic place to draw the line between merely knowing a word and also knowing what the word means.

Traditionally it is supposed that learning what a word means is coming to exercise the "same concept" in connection with the word that adults do. But, I have argued, a concept is an ability, and there is an ambiguity in the notion "same ability" that shows up also in the notion "same concept." Sometimes what counts as the same ability is what accomplishes the same; other times it is what accomplishes the same by the same means. In the terminology I am using, the organic chemist and the child both have the concept of sugar but they have quite different conceptions of it (§1.9). Having the same substance concept as someone else involves being able to reidentify the same substance they can. Identifying a substance the same way that someone else does, having the same "conception" of it, is an added frill.⁵ Assuming that knowing English requires having the concepts that correspond to English words, an advantage of talking this way is that then Helen Keller gets to know English.

What do we mean, then, when we speak of someone as coming to understand "the meaning of a word"? If the word denotes a substance, there is

⁵ There is a further complication, for it is possible for a single person to have two separate concepts of the same substance if they have failed to coidentify these, for example, if they do not understand that Samuel Clemens is the same man as Mark Twain. See §12.8.

a sense in which its meaning is, just, its referring to that substance. To know what the word means is just to have a concept of the substance that includes knowing to reidentify it by means of hearing that word. But of course the child may not be very good at identifying the substance. The child may make gross mistakes that an adult would not make. Is there then a richer sense in which a child can come to understand "what adults mean" by the word? Is there such a thing as "THE adult conception," of a substance? Given the numerous and diverse methods by which it is possible to learn to identify almost any substance, it seems that there could not possibly be.

On the other hand, for some (how many?) substances, it may be that there are core methods by which nearly every adult (the "nearly" is for Helen Keller) knows to reidentify them. Or there may be certain conditions under which any adult would recognize the substance, or examples of the substance that any adult would recognize given a chance to examine them. There also are occasional words that almost everyone who knows them learns to associate with certain facts about their referents, such as that Hesperus is seen low on the horizon in the evening, Phosphorus in the morning, or that Mark Twain was an author. ("Mark Twain," after all, was a pen name.) Also, occasionally words are passed down from generation to generation along with explicit conventional definitions. For example, every child who is taught the rudiments of geometry in school is taught that all points on a circle are equidistant from its center. Then there may be a sense in which the child does not fully understand "the meaning" of the word for that substance until her competence at identifying the substance has been filled out to match adult standards. The child must have an adult's conception of the substance. If this is what is to be meant by "knowing the meaning" then knowing how to track a substance only by recognizing its name would not be nearly enough for "knowing the meaning."

The difficulty is that there seems to be no way to draw clear lines around "the adult conception" of a substance. For example, do you, in this sense of "knowing the meaning," know the meaning of the word "molybdenum"? —or "brisket"? — or "African dormouse"? Perhaps your intuition is to say that you don't know what any of these words mean? I see no way to avoid a merely verbal dispute at this point if we persist with the question: who gets to count as really knowing the meaning?