

### 3 On Meaning, Meaning, and Meaning\*

---

#### *Introduction*

Various crisscrossing 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

\* © Ruth Garrett Millikan 2005.

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 proposal is that there are these three basic kinds of linguistic meaning:

- (1) conventional linguistic cooperative functions, to be called ‘stabilizing functions’;
- (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.<sup>1</sup> The meanings that characterize the public part of a language are fully extensional. I will have to say quite a lot about that before I am finished. Let us begin with simpler matters.

## 1. *Stabilizing Functions*

When speakers are conforming to the conventions of a language, speaker meanings coincide with conventional meanings. But speakers

<sup>1</sup> 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.

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: 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, my claim is that conventions do not generally require regularities of behavior, 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 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 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 thus 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 hearers' behaviors toward dogs. Such a speaker will eventually stop trying to use the word 'dog' for these purposes if they are never achieved. Also, 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—advice, instructions, directions, friendly requests, sanctioned directives, and so forth—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 example, 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: 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 (1969: 42, 58). Later, he modified this to allow that *almost* everyone conform *almost* all of the time (1969: 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.<sup>2</sup> 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 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 (Ch. 1; Millikan 1984: ch. 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

<sup>2</sup> 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.

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: ch. 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. At 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 (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.<sup>3</sup> Similarly, consider the various grammatical moods in a language. Sometimes you impart beliefs with the indicative mood but at other times you use it to give orders. Sometimes you ask questions with the interrogative mood but at 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: ch. 4; Ch. 10). Then, understanding the force of a particular utterance using the form will require disambiguation

<sup>3</sup> 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.

from context, just as understanding the meanings of individual words often does.

## 2. *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.<sup>4</sup> 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 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'

<sup>4</sup> In the case of most animal communication determined genetically.

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: ch. 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: ch. 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.<sup>5</sup>

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 my (1984: ch. 10) and (2004: ch. 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.

### 3. *Conceptions*

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.<sup>6</sup> 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

<sup>5</sup> See n. 4 above.

<sup>6</sup> See (Millikan 1984: ch. 4) on the most general stabilizing function of all referential terms.

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 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,<sup>7</sup> 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: chs. 4, 9; 2000: ch. 6), some practical capacity to

<sup>7</sup> The very fact that psychologists can make a study of how color constancy is perceived seems to cast doubt on the idea that colors are secondary qualities, at least in the Lockean sense of that term.

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 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: chs. 3–4) and the right story about perception through language (Millikan 1984: ch. 9; 2000: ch. 6; 2004: ch. 9), as well as a story about abilities that allows for their fallibility (Millikan 2000: ch. 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's 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 bat's 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 (Ch. 6; 1984: ch. 16–17; 2000: ch. 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 (Ch. 6; 2000: ch. 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 (Ch. 6; 1984: ch. 9; 2000: chs. 3, 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 Claus 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 ways 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: chs. 18–19; 2000: ch. 7; 2004: ch. 19).

#### 4. *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.<sup>8</sup>

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: ch. 12). Thus, for any hearer who 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

<sup>8</sup> 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: ch. 12).

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 1984: ch. 13; 2004: ch. 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.<sup>10</sup>

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

## REFERENCES

- Boër, S. E., and Lycan, W. G. (1986), *Knowing Who* (Cambridge, Mass.: MIT Press).
- Davidson, D. (1968–9), 'On Saying That', *Synthese*, 19: 130–46; repr. in A. P. Martinich (ed.), *The Philosophy of Language*, 4th edn. (Oxford: Oxford University Press).
- Donnellan, K. (1966), 'Reference and Definite Description', *Philosophical Review*, 75: 281–304.
- Grice, H. P. (1968), 'Utterer's Meaning, Sentence-Meaning and Word-Meaning', in *Studies in the Way of Words* (Cambridge, Mass.: Harvard University Press).
- Lewis, D. (1969), *Convention: A Philosophical Study* (Oxford: Blackwell).
- (1975), 'Languages and Language', in K. Gunderson (ed.), *Language, Mind and Knowledge* (Minneapolis, Minn.: University of Minnesota Press); repr. in H. Geirsson and M. Losonsky (eds.), *Readings in Language and Mind* (Oxford: Blackwell, 1996), 134–55.
- Millikan, R. G. (1984), *Language, Thought and Other Biological Categories* (Cambridge Mass.: MIT Press).
- Millikan, R. G. (2000), *On Clear and Confused Ideas* (Cambridge: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).

<sup>10</sup> To accord with this treatment, modal contexts need to be understood as contexts in which the truth conditions concern representations, rather than possible situations or worlds.

- Sellars, W. (1956), 'Empiricism and the Philosophy of Mind', in K. Gunderson (ed.), *Minnesota Studies in the Philosophy of Science, i* (Minneapolis, Minn.: University of Minnesota Press), 253–329; repr. in Sellars (1963).
- (1963), *Science, Perception and Reality* (New York: Humanities).
- Strawson, P. F. (1974), *Subject and Predicate in Logic and Grammar* (London: Methuen).

## 6 Why (Most) Kinds are not Classes\*

---

In the last century many hundreds of experiments were run by psychologists trying to discover how people ‘classify’ or ‘categorize’ items under kind words such as ‘dog’, ‘chair’, and ‘fruit’. The position I have taken on such words is that they do not designate classes but units of another kind entirely. There do exist some, a very few, uncompounded nouns that designate classes, but words like ‘dog’ and ‘chair’ and ‘fruit’ are not among them. Here I will introduce you to this negative position. I cannot attempt to defend it at any length in a short chapter, but I will present a portion of it in the most intuitively understandable terms I can muster. The details are spelled out in *On Clear and Confused Ideas* (Millikan 2000), which I will refer to as ‘OCCI’.

One place to begin is with the claims of biologist M.T. Ghiselin (1974, 1981) and philosopher David Hull (e.g. 1978) about what biological species really are. To be members of the same species, individual animals must belong to historical lineages that have a common origin. They do not have to be similar to one another in any specified way. For example, there are no genes that every dog has in common with every other dog. Every dog gene has alleles. Similarly, there are no properties that every dog has in common with every other dog. Nor is it mere overlap in properties or resemblance to some paradigm that makes a group of dogs be conspecifics. Highly similar species but that have different historical origins do not form one species but several. Species, according to Ghiselin and Hull, are not similarity classes but big,

\* © Ruth Garrett Millikan 2005.

scattered, historical *individuals* enduring through time. They are entities somewhat like the Kennedy family, which is held together, of course, not by 'family resemblance' in Wittgenstein's sense but by blood-relatedness.

On the other hand, in the case of species, blood-relatedness is bound to be accompanied by considerable overlap in properties. If the species reproduces asexually, the reason is that progeny are clones. If the species reproduces sexually, then each of the genes in the gene pool has to fit in with a random selection of other genes from the pool so as to help produce a viable individual frequently enough to get itself reproduced often enough not to be eliminated from the gene pool. No single gene that changes the animal in very extreme ways can survive. This results in what is called 'homeostasis' in the gene pool. Thus, the various individuals within a species *mostly* resemble one another in a great variety of ways, but do not *all* resemble one another in any particular ways. But what pulls them together as a group is not just that they have common or overlapping properties, but that they tend to have common and overlapping properties *for a good reason*. One individual is like the next *for a good reason*. There is a good explanation of why one is likely to be like the next. Various kinds of inductions drawn over the members of a species are likely to be sound owing to certain kinds of causal connections among these members.

On this analysis of what pulls the members of a species together, species are not classes. Classes are defined by the members having certain common properties. Fuzzy classes may be defined by the members having overlapping properties or by their having many properties in common with a paradigm or paradigms. But the members of a class do not need to be like one another for any reason. They may be like one another quite by accident. Categories are classes or fuzzy classes. Species names are not names of categories.

Now, I need, first, to explain why this point is important and I need, second, to generalize it.

The point is important because it explains why it is possible to study a species as such, to gather stable information about it. If there is a reason why one dog is likely to be like the next in a good number of respects, then there is a reason why studying one dog is likely to yield

a considerable amount of probable knowledge about the next dog. In fact, of course, dogs are something it is possible to learn a great deal about. Consider how much time may be spent on learning about dogs by a student at veterinary college. True, all this knowledge is merely probable knowledge. Whatever one learns of the properties of dogs, it won't be analytic or necessary that every individual dog has each of those properties. But mere classes are not things one can learn anything at all about by induction. If there is no reason, given one member of a class, why the next member is likely to be like it, then if any inductions over the class turn out true conclusions, it can only be by accident. For example, it seems likely that there is no reason why one red triangular object should tend to be like the next in any respect other than redness and triangularity, so it is not likely that discovering, say, that one red triangular object is sweet will be of any use in predicting the taste of the next.

The way in which dogs are cemented into a unit is important, then, because it is only when individuals are cemented into a unit in some analogous way, such that there is a reason why one individual should be like the next, that we can obtain knowledge about this unit—unless, of course, by examining every member separately. It is obvious, then, why this sort of unit is the sort that tends to acquire a name. Names for mere classes are in most contexts quite useless. Names for units of this kind are names for the seeds on which all empirical knowledge is built, for all empirical knowledge is inductive.

The point about dogs is generalized by noticing other kinds of relations that tend to cement a unit together such that there is a reason why one individual within the unit or one part of the unit is liable to be like another. I have a name for units of this general kind, taken from Aristotle. I call them 'substances', a word which nonphilosophers may need to read as a new technical term but which philosophers may recognize as fairly traditional. Aristotle spoke of 'secondary substances'—the unit *dog* would be an example—and of 'primary substances', which were individuals. (Recall that what Ghiselin and Hull said about dogs was that they were big, scattered, historical *individuals* enduring through time.) There is a chapter in *OCCI* detailing many kinds of substances. Here I will discuss only a few, but enough, perhaps, to give you the flavor.

Substances fall roughly into at least three basic sorts, which I call 'historical kinds', 'eternal kinds', and 'individuals'. Historical kinds are like dogs. They are collections of individuals scattered over a definite spatiotemporal area that are causally related to one another in such a manner that each is likely to be like the next in a variety of respects. The two most obvious sorts of things that cause members of a historical kind to be like one another are these. First, something akin to reproduction or copying has been going on, all the various individuals having been produced from one another or from the same models. Second, the various members have been produced by, in, or in response to the very same ongoing historical environment; for example, in response to the presence of members of *other* ongoing historical kinds. A third and ubiquitous causal factor often supporting the first is that some 'function' is served by members of the kind, where 'function' is understood roughly in the biological sense as an effect raising the probability that its cause will be reproduced. It is typical for several of these kinds of causes to be combined. Artifacts are often good examples of this.

Consider chairs. Chairs have been designed to fit the physical dimensions and practical and aesthetic preferences of humans, who are much alike in relevant respects for the same reasons dogs are. Moreover, the design of a chair is pretty invariably influenced by the design of previous chairs, typically because these previous chairs have functioned well and were aesthetically pleasing within a cultural setting, relevant aspects of the cultural setting being reproduced elements as well. For these reasons, chairs form a rough historical kind. There are reasons that have nothing to do with any arbitrary points of definition why one knows roughly what to expect when someone offers to bring a chair.

Renditions of 'The Irish Washer Woman' or of *The Rite of Spring* form a historical kind. They are copied from one another or from scores that are transcribed from earlier renditions or copied from earlier scores. McDonald's restaurants form a historical kind. There are causes of their being so much alike. Professors, doctors, and businessmen form historical kinds, especially well-integrated ones when these groups are studied as limited to particular historical cultural contexts. Members of these groups are likely to act similarly in certain ways and to have attitudes in common as a result of similar training handed down from

person to person (reproduction or copying), as a result of custom (more copying), as a result of natural human dispositions (compare dog dispositions) or social pressures to conform to role models (copying again), and/or as a result of legal practices handed down from univocal sources. There is a reason why it may be productive to investigate, say, 'the attitude of American doctors towards acupuncture'. These attitudes are contagious. They spread.

The members of eternal kinds are like one another for a different kind of reason. They are alike because they possess a common inner nature of some sort, such as an inner molecular structure, from which the more superficial or easily observable properties of the kind's instances flow. The inner structure results in a certain selection of surface properties, or results in given selections of properties under given conditions. Popular examples of this sort of kind are the various chemical elements and compounds, along with various forms of these such as ice, liquid water, and steam. Portions of water have an inner structure in common that produces the same surface properties given the same temperature conditions. Strictly speaking, I suppose, gold, nitrous oxide, ice, and so forth are not kinds but stuffs, but samples of them are members of eternal kinds. Also, water molecules, electrons, protons, and so forth are examples of eternal kinds. Stars, planets, comets, asteroids, and geodes are eternal kinds, not because their properties flow always from exactly the same inner nature, but because they were formed by the same natural forces in the same sort of circumstances out of materials similar in relevant ways.

Eternal kinds can be said to have 'essences' in a very traditional sense, essences that are not nominal but real, often discovered only through empirical investigation. The reason that the members of these kinds have many properties in common is that they have a few fundamental properties and/or causes in common that account, given laws of nature, for all the others. Eternal kinds do form classes, all of whose members are alike in a variety of respects. But they are also much more than mere classes, because they are alike in these respects not by accident but in accordance with a causal explanation.

The last kind of (Aristotelian) substances are individuals. Individuals have been taken in modern times to have a very different sort of unity

than the unity of a kind, but there is a way in which the cement that holds a single individual together as it endures through time is quite a lot like the cement that holds a historical kind together. Ghiselin and Hull claimed that species are actually individuals, because they are held together not by a traditional essence but through historical causal connections. The other side of this coin is that individuals are rather like species. A species is a 'homeostatic system . . . amazingly well-buffered to resist change and maintain stability in the face of disturbing influences' (Eldredge and Gould 1972: 114; quoted in Hull 1978: 199). Similarly, an individual animal is a 'homeostatic system . . . amazingly well-buffered to resist change and maintain stability in the face of disturbing influences'. If a person is tall, brown-haired, knowledgeable about electronics, and a good piano player today, it is likely, though not certain, that she will have each of these traits also tomorrow. The various members of a species are like one another in part because they are, as it were, copied from one another. An individual physical object tends to have the same physical properties the next day as it had the day before because of natural conservation laws which tend not to copy, of course, but to preserve its properties from day to day. The effect, however, is much the same. The inference that an individual animal or inanimate object will probably have these and those properties tomorrow because it has them today is likely to yield a true conclusion for the same general sort of reason that an inference that other members of a species probably have these and those properties because this member has them is likely to yield a true conclusion. Individual objects are things that inductive knowledge can be collected about over time for the same sort of reason that historical kinds and, more broadly, eternal kinds are things that knowledge can be collected about over time.

I have explained why historical kinds, eternal kinds, and individuals, three basic kinds of (Aristotelian) substances, are similar with respect to the question why it is possible to gain inductive knowledge about one part of the cemented-together unity they compose from other parts. The reason this is possible for each is that it is not merely a class, either focused or fuzzy. Because substances are not classes, not units cemented together merely by some set of common or overlapping

properties, to have a concept of a certain substance is not to have a certain set of properties in mind, whether derived from paradigm cases or from exemplars. My next job then is to explain what it is to have a concept of a certain substance if not to have in mind a set of central properties. I am going to do this by explaining, first, what it is to have a concept of an Aristotelian primary substance—of an individual. Then I will generalize to other kinds of substances.

The idea that there are such things as ‘concepts’ of individuals is foreign to many psychologists and to many philosophers too. This is for interesting historical reasons that need not detain us here. If this use of the term ‘concept’ bothers you, then interpret me as just talking about thoughts of individuals or ideas of individuals. What is involved in being able to think of an individual?

One traditional twentieth-century answer to this question is that to think of an individual is to capture that individual with a description that uniquely identifies it. Another twentieth-century answer is that to think of an individual requires that you know how to identify it one way or another, perhaps by description and perhaps just by being able to recognize it, to differentiate it from other individuals, in perception. These views are close enough to the answer I would give myself that they will serve my purposes here. Something that they have in common and that I am sure is correct is the assumption that there is more than one way to think of the same individual; indeed, that there are innumerable ways to think of the same individual. An indefinite number of individuating descriptions apply to every individual. Similarly, there are, in general, numerous ways that the same individual might be recognized by sight, by characteristic sounds that they make, by smell (dogs are good at this), and so forth. Contrast the ways Helen Keller recognized her friends with the ways they recognized one another. Twentieth-century tradition had it, then, and I believe correctly, that there is no single or definite set of properties that one must either think of or be able to discriminate in order to have a concept, a thought, of an individual. Nor is there some central set of properties, some or most of which one must think of or be able to recognize in order to think of a particular individual. Similarly, I will soon claim, there is no central set of properties, all or some of which one must be able to think of,

recognize, or discriminate in order to think of the (Aristotelian) substance *dog*, hence in order to learn about dogs, to understand things said about dogs, and so forth. But I will come back to that part a bit later.

First, we have to deal with fallibility. The ways we have of recognizing individuals are always fallible in principle. Even supposedly individuating descriptions always presuppose that there is indeed one and only one thing fitting the description, something not guaranteed, for example, merely by the description containing superlatives. It might always be, for example, that no one is tallest or first in line, or first on the moon. More important, if you are actually to use an individuating description for purposes of recognizing an individual, you will have to recognize exemplifications of the properties mentioned in the description. But one's capacities to recognize objective properties are always fallible, for they depend on external intervening or mediating conditions such as lighting conditions, atmospheric conditions, sound absorption and reflectance properties of surrounding objects, obscuring conditions such as intervening objects, masking sounds and odors, and so forth. Nor is there an independent way of ascertaining what these mediating conditions happen to be in a particular case. There are always possible conditions under which you would misidentify or fail to recognize even your own mother or spouse.

Having the ability to recognize an individual, then, cannot be the same thing as being infallible in recognizing it. I have the ability to walk. It is one of my very best abilities. It does not follow that it cannot happen that I trip and fall when trying to walk. These reflections suggest that what we need here is an analysis of what it is to have an ability to do something, such as walking or recognizing your mother, that does not equate an ability with any simple sort of disposition. That analysis has been given in *OCCI*. But that all abilities are fallible is common sense, and I propose just to assume it here.

Tradition plus common sense suggests, then, not only that different people can have different kinds of concepts of the same individual by using quite different methods of recognition, but also that the methods any one person uses to recognize, hence to be able to think of, an individual will be fallible. Nor do these methods constitute a definition

of the individual. Your mother is not defined by the way you recognize her; say, by the look of her face and the sound of her voice. She doesn't have a definition, a set of properties, that make her be who she is. She is not a class that happens to contain only one member.

Similarly, the species *dog* is a unity that different people can have quite different kinds of concepts of, by using quite different methods of recognition. Whatever method a person uses for recognizing dogs, this method may always be fallible. Nor does the method that a person uses for recognizing dogs constitute a definition of what dogs are, even for that individual. The species *dog* is not just a class that happens to contain so-and-so many members.

What makes substances interesting is that there is often a great deal that can be found out and known about them. Often they have a great many properties. And it is typically the case that numerous of these properties and numerous sets of these properties will each be diagnostic of the substance. That is, each of these properties or property sets will be found only or typically when the substance itself is encountered. At least this will often be so within the spatial and temporal area inhabited by the person needing to recognize a substance. Mistakes that people might have made had they lived in different places and times are not relevant to their actual abilities to recognize substances. This is why it is possible for different people to have concepts of the very same substance by very different means. Children and chemists have different ways of recognizing sugar. You and Helen Keller have different ways of recognizing nearly every secondary substance, nearly every ordinary stuff, and nearly every ordinary eternal and historical kind.

Further, none of the ways that a person knows to diagnose the presence of a substance needs to be infallible. No particular set of properties used to diagnose a substance is ever definitional of it, although in the case, especially, of eternal kinds, empirical investigation may reveal (with probability) that, in fact, some sets are always correctly diagnostic. It is always logically possible that there is some other substance that has parts of its cemented-together unity that share the very same properties as the properties one is using, with practical success, for diagnosis of a certain substance. I can put this for

philosophers by saying that the possibility of 'twinearth water', certainly of 'twinearth dogs', and, indeed, of 'twinearth Mama', indistinguishable from your mother, is never ruled out by logic alone. It takes more than a set of properties in your mind to determine a substance. It takes a certain sort of causal glue in the world, holding that substance together. But given that glue in the world, conceptual access to that glued-together unity may be had by reference to any of many of its different parts or properties.<sup>1</sup>

In talking about what is involved in having a concept of a substance, I have quietly been making an assumption that I must now bring into the foreground as a claim. I have spoken of ways of recognizing a substance, and I have said that your ability to recognize a certain substance can depend on your inhabiting a certain space-time locale, one where certain diagnostic properties do mostly signify encounters with that substance rather than with others. The assumption I am making is that thinking of a substance involves the ability to recognize it, as it were, *in the flesh*, not merely the ability passively to contemplate its properties. We have thoughts of substances in order to be able to collect information about substances, which information we pick up on some occasions and then apply on others. To pick up information about a substance you must be in a position to interact with the substance, or with other things that interact with the substance, other things that are influenced by the substance or that influence it. Natural information is transmitted in the causal order, and you have to be in the causal order with whatever the information is information about in order to receive it.<sup>2</sup>

<sup>1</sup> Philosophers may detect a missing link in this analysis. The link is needed to connect the ability one has to recognize a particular substance to prior encounters with that particular substance rather than with similar substances on twinearth or wherever. That link is supplied in the description of abilities given in *OCCI*. What an ability is an ability to do is determined not merely by current dispositions but by the histories of the mechanisms responsible for those dispositions.

<sup>2</sup> I am using the notion *natural information* in a way somewhat like the way Dretske uses it in *Knowledge and the Flow of Information* (1980), yet not quite that way. For our purposes here the difference probably does not matter, but a careful description of the kind of information I have in mind is in Millikan (2004: chs. 3–5) where I call it 'local information'.

Now, if you think about that claim for a moment, you will see that it is a fairly radical one. Surely you can have a concept of the last dinosaur species on earth to become completely extinct and of the first baby to be born next year, and of any other substance for which, although you have never encountered it, you do know an identifying description. And you have these concepts without having the slightest idea how to identify these things in the flesh. Surely you can have a concept of molybdenum—you can think about it and ask questions about it—without being able to identify it in the laboratory. Surely you can have a concept of Socrates without being able to identify him in the flesh, even if you were to be transported back to ancient Athens. Let me tackle the descriptions first, then come back to molybdenum and Socrates, for they will prove far more interesting.

The descriptions are handled this way. That your circumstances are such that you never get a chance to use an ability that you have does not take that ability away from you. You won't lose the ability to swim just because they chain you to a post in the middle of the Sahara desert for the rest of your life. If you understand the terms in any description and know how to apply them—that is, you know how to recognize the other objects and properties and relations mentioned in the description—and if you are right that the description is identifying, then you know a way to identify the substance that the description describes. You would do so by encountering something that you can recognize directly as fitting that description, or by coming across something else that you recognize as carrying information telling what fits that description. There are many cases in which you just aren't at all likely to come across any such information, but that is irrelevant to whether you have a capacity to recognize the substance. I am assuming here a fairly usual reading of the notion of natural information, according to which information about the past and about the future are entirely routine kinds of information (but see n. 2). And I am about to claim that language is a standard medium through which natural information is transmitted, hence a standard medium through which substances are recognized exactly as they are recognized 'in the flesh' through other media such as light and sound.

Now consider molybdenum and Socrates. It seems an obvious fact that many of our concepts of substances have been acquired without

encountering those substances directly but only by hearing about them. Moreover, as Kripke (1972), Putnam (1975), and Burge (1979) have observed, we often have no unique descriptions of these substances in mind either. How then can we be said to know how to recognize them? The answer, I claim, is that speech is just as direct a medium for the perception of objects and events and their properties as is the light reflected off objects, the smells emanating from objects, the sounds emanating from events in the environment, or the mechanical stimulations caused by objects in direct contact with one's body. This is a thesis that requires defense, and I have defended it at length both in *OCCI* (ch. 6) and in Millikan (2004: ch. 9). Here I can only throw out the rough idea, hoping that if it strikes you as dubious you will look to these longer versions and defenses before final judgment.

The claim is that hearing and immediately believing a sentence about a fact or occurrence is in relevant respects just like, for example, seeing that something is the case or seeing that something has occurred and immediately believing it. 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 in other ways, 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). Recognizing a linguistic reference to a substance is as much a way of recognizing the substance 'in the flesh' as any other way of recognizing it. It is identifying it and recognizing natural information concerning it through one more medium of manifestation. Think of this medium, the speech of another person, as like an instrument that aids perception. The lens of one's eye is, of course, an instrument that aids perception. If one wears corrective lenses, they are another such instrument. The speech of another person is analogous to somewhat more complicated instruments of this kind. Like a camera, a radio, a CAT scan, or a microscope, another person who talks to you picks up information-bearing patterns from his environment, focuses them, translates them into a new medium, and beams them at you. 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 you in ways that, barring certain interferences, you have become tuned to interpret. Becoming tuned to interpret the information-bearing patterns that are common in a certain language community is coming to understand the language of that community. Similarly, a radiologist must learn to interpret the information contained on X-ray images and the auto mechanic must learn to interpret the information contained in the sounds emanating from ailing automobile engines.

The notion that understanding and believing what is said to you is just one more level of natural-sign reading on the same level as ordinary perception is to many people quite unintuitive. One reason is that what is given to you in ordinary perception is always given as in some quite definite current relation to you. It is given as happening at the time you perceive it, as happening relatively nearby, and often as bearing quite an exact spatial relation to you. This kind of information is needed to guide action, for how one can currently act on a thing always depends on its current relation to oneself. Ordinary perception is for immediate action, whereas what one learns through language is not typically used that way. Usually you are not told what exact spatial and temporal relations the objects and events being presented to you through language have to you here and now. But there are intermediate cases; for example, video recordings. It is clear enough that you *perceive* things happening when you watch a video, but, as in the case of language understanding, you do not perceive the spatial and temporal location to yourself of what occurs on video.

A second reason that the comparison between ordinary perception and language comprehension is unintuitive is that ordinary perception is so much more reliable than what one hears said, at least under common circumstances. It is not easy to fool ordinary perception. To create strong perceptual illusions requires a good deal of knowledge about the perceptual mechanisms and often quite special equipment, of the kind, for example, that optometrists have in their examination rooms. This is a difference of degree, however, a mere difference in frequency, not a difference in kind. Recalling that film dubbing is currently the rule rather than the exception, what differences are there, for example, among (1) believing what you apparently see when

you look through the peephole into an Ames room, (2) believing what you see when a film has been dubbed, and (3) believing what you hear someone say when it's false? In the modern world, if you want to believe only what's true, you often have to apply heavy filters to other methods of perception as well as to perception through language.

The upshot of these reflections is that we can understand how it is possible to recognize a substance through the information that language bears; indeed, how it is possible to come to be able to recognize a substance pretty much merely by learning a word for it. This is how we manage to have concepts of Socrates and, for most of us, how we manage to have concepts of molybdenum. To have a word for a substance is to have an essential part of an ability to recognize manifestations of it that are generated in a particular language community. That, I have argued, is why it is possible for small children to learn, as Chomsky puts it (1995: 15), 'a word an hour' between 18 months and 6 years of age.

## REFERENCES

- Burge, T. (1979), 'Individualism and the Mental', in P. French, T. Uehling, and H. Wettstein (eds.), *Studies in Metaphysics: Midwest Studies in Philosophy, iv* (Minneapolis, Minn.: University of Minnesota), 73–121.
- Chomsky, N. (1995), 'Language and Nature', *Mind*, 104: 1–61.
- Dretske, F. (1980), *Knowledge and the Flow of Information* (Cambridge, Mass.: MIT Press).
- Eldredge, N., and Gould, S. J. (1972), 'Punctuated Equilibria: An Alternative to Phyletic Gradualism', in T. J. M. Schopf (ed.), *Models in Paleobiology* (San Francisco, Calif.: Freeman, Cooper), 82–155; quoted in Hull (1978).
- Ghiselin, M. (1974), 'A Radical Solution to the Species Problem', *Systematic Zoology*, 23: 536–44.
- (1981), 'Categories, Life and Thinking', *Behavioral and Brain Sciences*, 4: 269–83.
- Gilbert, D. (1993), 'The Assent of man: Mental Representation and the Control of Belief', in D. M. Wegner and J. W. Pennebaker (eds.), *Handbook of Mental Control* (Englewood Cliffs, NJ: Prentice-Hall).
- Hull, D. L. (1978), 'A Matter of Individuality', *Philosophy of Science*, 45: 335–60; repr. in Sober (1994: 193–216).

- Kripke, S. (1972), *Naming and Necessity* (Cambridge, Mass.: Harvard University Press).
- Millikan, R. (2000), *On Clear and Confused Ideas: An Essay on Substance Concepts* (Cambridge [OCCI]: Cambridge University Press).
- (2004), *Varieties of Meaning* (Cambridge, Mass.: MIT Press).
- Putnam, H. (1975), 'The Meaning of "meaning" ', in Keith Gunderson (ed.), *Language, Mind and Knowledge; Minnesota Studies in the Philosophy of Science, vii*. (Minneapolis, Minn.: University of Minnesota Press).
- Sober, E. (1994) (ed.), *Conceptual Issues in Evolutionary Biology* (Cambridge Mass.: MIT Press).