

Chapter Four
The Nature of Abilities:
(How is Extension Determined?)

' 4.1 Abilities Are Not Dispositions of the Most Common Sort

The conception you have of a substance does not determine the extension of your concept. The extension is the extent of a certain substance in nature, not whatever you would identify as part of the extension. But the extent of which substance? That question is crucial. What determines, in the particular case, what particular substance one's perhaps stumbling, sketchy, and inadequate conception is aiming at? This chapter will make some progress toward answering that question. Further pieces of the puzzle will be added in Chapter Five and the last pieces will finally settle into place at the end of Chapter 14.

Substance concepts are abilities of a certain kind. They are, in part, abilities to reidentify their assigned substances. How are these substances assigned? It is not a function of the cognitive systems as handed down by natural selection to identify any particular substance. Natural selection did not endow me with the ability to identify either 1969 Plymouth Valiants, or gasoline, or my husband. What I was endowed with was the capacity to acquire these abilities. Thus the general form of the question what determines the reference of a certain substance concept is: What determines what a learned ability is an ability to do? It will help to tackle the matter in this entirely general form.

The question what abilities are deserves a lot of attention that it hasn't gotten. The modern philosophical tradition has unreflectively assimilated abilities to capacities and capacities to dispositions. This affords a slippery slope. A "capacity" can be either a living thing's abilities (the capacity of a camel to go without water for weeks) or a nonliving thing's dispositions (the capacity of gold to resist corrosion) but abilities and dispositions, in these contexts, are not the same. An ability is very much more than a disposition in one way, I will argue, and less in another.

In the usual philosophical sense, for something to have a "disposition" to behave in a certain way, say to do A, is for there to be circumstances under which it will do A in accordance with natural law. But to say merely that there are circumstances under which a thing will do A is nearly vacuous. If put in the right circumstances—for example, if hooked up to the right gadgets and so forth—any kind of thing can probably be made to supply a contributing cause to any kind of outcome you please. So usually one has in mind a disposition to do A in some definite circumstances C. To speak without restriction of a disposition to perform in a certain way must be either (1) to imply that the actual conditions, or the most likely conditions, are such as to realize the disposition, or such that it is often realized or (2) to assume or implicitly to refer to some specific sort of circumstances in which it will be realized. It is in the first way that people are said, for example, to have bad dispositions or sunny dispositions, or to be well disposed toward one another, and so forth. It is in the second way that salt is said to have a disposition to dissolve.

Sense (1) certainly is not the same as the notion of having an ability. Sunny dispositions are not, as such abilities. Moreover, we all have dispositions, in sense (1),

to depress the carpet on which we walk, to attract nearby mosquitoes and frighten nearby mice, and also to slip on ice. None of these are abilities. Moreover, we have many abilities that we have no dispositions at all, in sense (1), to realize. I have, for example, the ability to kill cats, to stand on my head on the commuter train, and to play bebop on the violin, but I have no sense (1) disposition to do any of these things. At the very least, something else must be added to a sense (1) disposition to make it into an ability. And it has to be subtracted that one is or is likely to be in conditions that realize it.

Perhaps abilities are dispositions in sense (2) if we fill in the conditions in the implicit antecedent of the conditionals correctly. An ability to do A is a disposition to do A if C_{what?} The obvious suggestion is, "if one tries."

The first thing to notice about this answer is that it will require us to unpack the notion trying, and that this is not easy. We cannot do it, for example, by making reference to an intention to do A, if intending to do x requires having a concept of A. That would probably leave many animals with no abilities at all, and it would certainly leave us without most of our cognitive abilities. Especially, it would leave us without the ability to identify substances. We cannot suppose that in order to have a substance concept one must first have a prior concept of the substance one is trying to identify so that one can intend to identify it. We would have to unpack "trying" by reference instead to some kind of purpose more primitive than that involved in explicitly intending, in terms, for example, of biological purpose. I will come back to this later.

But no matter how we unpack trying, an ability is not the same as a disposition to succeed when one tries. My abilities often fail me. I have the ability to walk but also, under certain circumstances, the disposition to slip or trip, and I do this exactly when trying to walk. I know how to cook, but I may still burn the dinner tonight. Many people with the ability to swim have drowned, presumably when they were trying to swim. We got into the question what abilities are by noting that the ability to reidentify a substance is fallible. Sometimes I misidentify things, but I would not do that if it were not, in some sense, my purpose to identify them. It seems that something else has to go into the antecedent besides "if I try," if abilities are to be understood as dispositions. If I try under what conditions?

There is a standard reflex answer to this sort of question: "under normal conditions." This combines disposition in sense (1) with disposition in sense (2). We saw that abilities can't just be what one has a disposition to do under the circumstances one is in or is likely to be in. Nor are they just what one has a disposition to do when one tries. Are they, then, what one has a disposition to do if one tries under the circumstances one is in or is likely to be in under circumstances that are "normal"? That is, are abilities dispositions for the most part to succeed if one tries? The reference to likelihood would explain why, for example, although lots of people try and succeed in catching fish on hook and line, we do not say that a person knows how to catch fish, but only that they know how to fish. Reasonably, this is because they do not regularly succeed when they try to catch fish. It might also help to explain how different people can have abilities to do the same thing, but some can be better at it than others. Some

are more likely to succeed when they try than others. So let us explore this answer.

I have an ability to swim granted I am likely to succeed if I try. I am now sitting at my desk in a dry sunny room. I now have the ability to swim, indeed I have had that ability ever since I was six. But if I try to swim sitting at my desk in a dry sunny room will I be likely to succeed? Not likely. I know how to swim, but I am not in a position at the moment to swim. Indeed, most of the time I am not in a position to swim. I go down to the lake only once in a while. So it can't be that knowing how to swim is the same as being likely to succeed if I try. CWhat has gone wrong?

' 4.2 Having an Ability to vs Being Able to

Focus on the antecedent, "if I try." Is the idea that if I just tried, I could do it? Cthat all that is needed for me to accomplish swimming is to add in a trying? That, we have just seen, cannot be the right idea. Certainly I could not immediately swim just by trying. It might be, of course, that if I tried I could bring about a situation in which I would then be in a position to swim. But this is not true for all abilities either. There is a difference, a surprisingly wide one, between being able to and having the ability to. I may know exactly how to invest a thousand to make a million, except that I haven't got any money at all. I have the required ability, but I am not able to exercise it, nor am I necessarily in a position to get myself into a position to invest a thousand to make a million. I may know how to make a marvelous gourmet dish for which the ingredients are completely unavailable. If there was a time when people knew how to make tasty dodo stew, they didn't suddenly stop knowing how on the expiration date of the last dodo. They still had the ability, but they were no longer able to apply it. True, in ordinary speech the word "ability" may be a little fuzzy around this edge. But let us settle on using it in this unwavering manner. Abilities don't disappear just because the world is uncooperative in supplying the necessary conditions for their exercise.

You may have noticed that I have been slipping back and forth between "knowing how to" and "having the ability to," whereas there are some differences in how these terms are typically used. One difference is exactly that "having the ability to" tends to slide a bit toward actually being able, whereas "knowing how to" does not. Whether I am in a position to do a thing virtually never bears on whether I "know how" to do it. Also, "knowing how" is more likely to be used for learned skills, "having the ability" for innate skills. I know how to do sums and ride a bicycle. I don't know how to see. Instead, I have the ability to see. I will continue to ignore this distinction. It won't be pertinent for our purposes.

' 4.3 Ways to Improve Abilities

Here is another way to interpret "if I tried." Usually I will not try something unless I believe I have a fair chance of succeeding. For example, you could not possibly induce me to try to swim while sitting at my desk in a dry sunlit room. You might induce me to pretend to try, but not really to try. Perhaps what we need then is this. One's abilities are what one has a disposition to do if one tries under the circumstances one would likely be in were one to try. Let us see how far this definition will take us.

Abilities can be more or less well developed. You can know very well how to read or swim or ride a bicycle, or you can know these things but not very well. I know a small

child who replied when asked "Are you learning how to play the violin?" with a denial: "No, I already know how to play the violin; I'm learning how to play it better." What does learning how to do something better consist in? In the case of violin playing, of course, it consists in learning how to turn out a more polished product. But there are two other dimensions in which abilities can improve.

You can learn to do the same thing under more circumstances. For example, you might learn to drive a car safely even on ice, or how to ride a bike without slipping even over muddy paths, or how to row a straight line even in a crosswind, or how to peel potatoes neatly with only an ordinary knife available rather than a potato peeler. If abilities are dispositions to do something when you try in circumstances you are then likely to be in, then it would be wrong to think of learning these various things as acquiring new abilities. It is the same thing you are aiming at, whether the wind blows crosswise or from the stern, whether the bike path turns out muddy or dry. You are trying for the same thing; it is just that the circumstances are different. You are not aiming at the circumstances. They were just the ones that happened to come along. (Of course it also is possible to aim, specifically, at rowing-straight-in-a-crosswind. You purposefully go out in a wind, and then purposefully row across it while trying to keep a straight course. That is a different thing to do. For example, in that case you don't succeed if you don't find a crosswind.)

Notice that learning to do the same thing under a wider variety of conditions need not produce a more reliable disposition to do that thing when you try. So long as throughout your learning you are pretty good at distinguishing conditions, given your current ability level, under which you would succeed from conditions under which you would fail, and so long as you don't try unless reasonably sure you will succeed, no change in the frequency of your failures necessarily occurs. So this might seem to be a way to get better at something, to improve an ability, without becoming more likely to succeed when you try. On the other hand, the harder it is to see the ground ahead with a project, the more difficult it is to see at the start whether conditions necessary for success lie ahead, hence whether to begin trying. Then learning how under more conditions will increase the probability of succeeding when you try. So far so good for our proposed definition.

This brings us to the second way in which you can improve an ability. You can learn to recognize better the circumstances in which you will succeed if you try. Adults fall down less often than children, and have to be rescued from drowning less often. This is not just because they are better coordinated. They are better at recognizing risky situations, better at knowing when not to try, or not to try this way but rather to try some other way. It is not just that adults are more cautious. They know better when to be cautious. This way of improving an ability necessarily makes it more reliable. It improves the chances of succeeding when you try, because it improves the chances of trying only at times when you will succeed.

Granted these ways that an ability can improve, notice that we can't interpret "normal conditions" for exercise of an ability as conditions people generally are in when they try that activity. Different people may know how to do the same thing in quite

different ways, under quite different conditions, so that the conditions under which they would try are quite different. Perhaps neither would be able to do it under the conditions the other finds most suitable. This sort of effect is especially evident and especially important in the area we are most interested in. Different people can have entirely different conceptions of the same substance, entirely different ways of recognizing it, so that neither would recognize it under the conditions the other would. People also can have skills that no one else has, skills developed under conditions, perhaps, in which no other human has been. Consider, for example, skills developed by space walkers engaged in specialized tasks. Similarly, a person can have a concept of a substance of which no other person has a concept. Any reference to "normal conditions," if it were to help in defining abilities would have to be made not just relative to the particular ability but also to the particular person who is able.

' 4.4 an Ability Is Not Just Succeeding Whenever One Would Try

Is it true that my ability to do A is a disposition that I have to do A under conditions I would probably be in if I tried? How is a probability concerning only one person to be defined? In fact it is a very tough question how to interpret a probability of this kind, and merely waving one's hands toward something called "individual propensities," as some now do in the philosophy of biology, is certainly a step in the wrong direction. But I will let these issues lie. Whatever sense might be made of the notion of an individual's own personalized probability of being in certain conditions when they try, still an ability can't be defined in terms of an individual's probability of succeeding should they try. Here's why.

At the start of this discussion, I argued that an ability can't be merely a disposition one has under conditions one is likely to be in, citing as counterinstances those dispositions that make up one's temperament, one's disposition to depress the carpet on which one stands, and to slip on ice. A reason to reject these dispositions as abilities is that so far as I, the organism, am concerned, these dispositions are entirely accidental, having nothing to do with purposes, either of nature's design or of mine. I then added a reference to tryings, thus importing a dimension of purpose of what happens not by accident but by design and this helped us quite a lot. But defining an ability as an individual's likeliness to succeed should they try in conditions they would likely be in if they tried snags on exactly the same problem one level up. The conditions I am in may not be conditions I was designed by nature to be in, or that my abilities were designed through learning processes to operate in. If (1) I have a disposition to succeed if I try, or if (2) I fail to have such a disposition, these could be purely coincidental matters, hence might not indicate to an ability.

Consider the second case first, the case where I have an ability but fail to have any such disposition. Suppose, for example, that the last dodo is dead but the dodo stew chef doesn't know it. And suppose there are people about still posing as dodo pedlars, and that it is quite likely the dodo stewmaker will be taken in by them. Then it will not be true that if the dodo stewmaker tries to make dodo stew he will most likely succeed. If he tries to make dodo stew it will be because he is mistaken in thinking he has bought dodo meat. Recalling the distinction noted earlier between knowing how to do A and

being in a position to do A, it is clear that something has gone wrong. The dodo stew chef surely does have the ability to make dodo stew, though he is in no position to exercise it. But he lacks a disposition to make dodo stew when under conditions he would probably be in if induced to try.

Now consider a case in which a disposition to succeed if one tries is present but does not correspond to an ability. Macaffee the cat has learned how to get himself let in the door by stretching up tall and scratching by its side window, thus making himself both heard and seen from inside. (Our cat does this.) Now the house changes owners, but Macaffee won't stay in his new home and returns to the old. The new owners, however, dislike cats and have no disposition at all to let him in. However, soon after his return they have an automatic door installed for a wheelchair occupant with a wide push button outside at exactly the place Macaffee is disposed to scratch. Now the first day after the installation it is true that Macaffee has a disposition to get the door open for himself if he tries under the circumstances he is in or is likely to be in when he tries. But he no more knows how, at this point, to let himself in than you know how to win the state lottery if, happily, you happen to do so. What one does successfully only by pure chance is not something one knows how to do.

Now in the preface to this book I promised not to play Counter Examples in Queer Possible Worlds, and I do not mean to be doing so here. My point is not that we wouldn't say Macaffee knows how to let himself in, that this is not how "our concept of an ability" is fashioned. Rather, my point is that the same principle is involved here as in cruder cases where dispositions very obviously do not equal abilities. Abilities are distinct from dispositions in having a necessary involvement in the purposive and non-accidental order. A consistent theoretical definition of abilities will consistently take this into account. What is interesting and central about abilities will not be captured if we let in cases where results are achieved coincidentally, or if we exclude cases where failures to achieve results are coincidental.

Notice also that the attempt to turn abilities into dispositions with the antecedent "if one tried," as we came to interpret that phrase, was a cheat all along. It is not just that trying requires an analysis and is unlikely itself to unpack into dispositions. Rather, the conditionals corresponding to dispositions are supposed to refer to causes in their antecedents and effects in their consequents, and we did not use "if one tries" in that way. If the match lit, then it was struck may be a true conditional, but it does not correspond to a disposition for the match to be struck if it lights. Similarly, our "if one tries" was not used merely as a reference to a cause. The idea was that if one tries that must be because one believes one can succeed, and that in turn will likely be because one recognizes one is in conditions that will make success possible. The only true disposition here is the disposition to succeed if one tries under certain conditions, namely, the ones one attempts to recognize. Moreover one attempts to recognize these conditions under that very description, namely, as certain conditions under which one can succeed. The description of this disposition thus appears to be empty. One has a disposition to succeed if one is right that conditions under which one will succeed are present.

' 4.5 Distinguishing Abilities by Means or Ends

We are still seeking the relation of an ability to the conditions of its successful exercise. Let me start fresh with an example that may bring out this relation more clearly.

Suppose that you once learned how to use WordPrefab 2.2, but for the last 15 years you have had a stenographer and haven't looked at any more recent word processing programs. WordPrefab 2.2 is now completely extinct and nobody anywhere has kept it either on their computers or on disk. It is clear then that you will no longer succeed in word processing if (without retraining) you try, no matter how we interpret the antecedent of the conditional. But do you still know how to do word processing?

A temptation is to reply that you don't know how to do word processing period; you only know how to do word processing in or by means of WordPrefab2.2. But consider: no one knows how to do word processing without employing some program or other, some means or other. So from that sort of reasoning it would follow that nobody knows how to do word processing period. Similarly, since everyone employs some stroke or other in order to swim, and since if we cut strokes finely enough, surely no one knows how to employ every possible swimming stroke there is, it would follow that nobody knows how to swim period. Indeed, since everything accomplished in, or in relation to, the world outside one is accomplished by some means or another, requiring some definite conditions or other to be in place in that outside world, it will turn out that no one has any abilities to affect the world outside period. Rather, all abilities are hedged with specific means and conditions.

In ' 1.9 I pointed out that we have two ways of distinguishing or designating abilities: by ends, or by ends plus means. Where abilities to identify substances are concerned, this is the distinction between what I called "concepts" and what I called "conceptions." What has happened just above is that we have distinguished abilities by means rather than merely by ends. If we distinguish abilities by means, then of course no one ever knows how to do something fullstop. The various means they know must always be described in saying what they know how to do. It is also true that if we cut designations by means finely enough, it may well be that two people seldom if ever have exactly "the same ability."

Now there is nothing wrong with distinguishing abilities in this manner for some purposes. But that should not blind us to the legitimacy of distinguishing them by ends alone for other purposes. For many purposes, it is quite legitimate to say that some people know how to swim whereas others don't, and that all those people who know how to swim have the same ability. Similarly, it is quite legitimate to say of the person who knows how to do word processing using WordPrefab2.2 that they know how to do word processing. They can have an ability to do word processing, even though, unfortunately, the conditions required for the means they employ (presence of a WordPrefab2.2 program on an accessible computer) are unavailable, so they are not in a position to exercise that ability. Having an ability is not the same thing as being able.

To see the distinction between counting abilities by their ends and counting them by their ends-plus-their-means clearly, it is important to keep another distinction made earlier in mind (' 4.2). This is the distinction between aiming at an end while also knowing

to use certain means to that end in certain circumstances, and aiming at getting to the end in those very circumstances or by the use of those very means. Aiming at word processing and aiming at word-processing-with-WordPrefab2.2 are two different things to aim at; knowing how to do these two things corresponds to two different abilities when abilities are counted by ends only. This is very easy to confuse with the fact that there are two different ways to count abilities, either by reference to ends only or by reference also to means. Similarly, abilities as counted by ends often have, for a given individual, many alternative means, employed under different conditions. This is not the same as the person having many alternative ends. It is not the same as the person's having many alternative abilities as counted by ends. For it may be that the person never aims at any of the specific means, or at being in conditions in which these means specifically can be used. Thus you walk on a mountain trail constantly adjusting your means to the conditions coming up underfoot, but you aim at none of these means nor at being in any of the conditions requiring these means.

' 4.6 Abilities Are Not Dispositions but Do Imply Dispositions

We have settled then that the person who learned how to use WordPrefab2.2 learned and still knows how to do word processing period whether or not WordPrefab2.2 still exists in any form. But there is no possibility at all in their current situation that their ability to do word processing will manifest itself, whether or not they try. Thus it appears that knowing how to do word processing has nothing at all to do with this person's dispositions in their current situation. At the same time, it is clear under what possible conditions this ability could be manifested Cnamely, under the condition that this person would have WordPrefab2.2 on the computer they used. What then IS the relation between this person's ability and the circumstances under which it would necessarily be manifested?

The relation, it seems, is historical. The conditions in which this ability would be manifested are the conditions in which it was historically designed as an ability. In general, the conditions under which any ability will manifest itself are the conditions under which it was historically designed as an ability. These are conditions in which it was learned, or conditions in which it was naturally selected for. They are conditions necessary to completing the mechanisms by which past successes were reached by the systems or programs responsible for the abilities. Past presence of these conditions helped account for the selection or maintenance of the systems or programs constituting the abilities. What I know how to do I must once have learned how to do. Otherwise it would not be knowledge but mere luck. Or what I have an ability to do is what my systems were maintained or selected for doing. That is my suggestion. I know of no other way to understand what abilities are that is consonant with the case advanced above against their merely being current dispositions of some kind. (This can also be taken as a challenge, of course, to come up with alternative suggestions.)

To an ability there always corresponds a disposition, but it does not follow that an ability IS a disposition. If an ability to do A were a disposition to do A under specified conditions, then we should be able to specify the conditions under which anyone must be disposed to do A if and only if they have an ability to do A. But these conditions cannot

be specified for the general case. Because different people who have learned to do the same thing may have learned to do it in different ways, relying on different conditions, the conditions under which a given ability (defined by its end) might manifest itself may be entirely different for different people. There is no such thing as the set of conditions under which it is necessary that any person be able to manifest their ability to do word processing or to swim. For each person, there is an independent reference here to personal history. Each person's ability to do A rests on a disposition defined through their particular past. Each has a disposition to do A if they try to do A under the conditions that accounted for their own past successes in doing A. If they have no such disposition, of course, then they have lost the ability to do A. It does not follow that their ability IS a disposition. Rather, which disposition(s) it is that can manifest the ability to do A is determined by which conditions helped account for the acquisition of this particular person's ability. To attempt to define the ability by reference its historically enabling conditions would move one in a circle.¹

Another question that has been running in the background is now easily answered as well. How should we understand the notion "trying" such that an ability is manifested when one succeeds in what one tries under historically enabling conditions, yet such that there can be mental abilities the goals of which we do not represent to ourselves or even understand, hence do not "try" to reach in the most ordinary sense of "try"? "Trying" to do A, in the needed sense, is simply the initiation and running to the point of success or failure of a mechanism or program that is designed to do A.

' 4.7 What Determines the Content of an Ability?

What then determines what a learned ability is an ability to do? The difficulty, recall, was that the kinds of parts, systems, and programs that embody abilities have innumerable dispositions that are not abilities, and also that real and even strong abilities can fail, the

1 The conditions under which a person with the ability to do A necessarily has a disposition to do A if they try can be thought of as a peculiar sort of normal conditions, namely, conditions that were normally present and active on the occasions of their past successes. These are conditions of the sort that I labeled with the capitalized term "Normal" in (Millikan 1984). "Normal" conditions for proper functioning of any mechanism are conditions that obtained and were active on occasions of successful past performance leading to selection of the mechanism.

outer conditions necessary to support them being absent and the absence remaining undetected. So a look only at the mechanism that has or embodies abilities will not reveal what its various dispositions are. What then does determine what its abilities are? I had originally raised this question as a more general form of the question: What determines, in the particular case, what particular substance one's perhaps stumbling, sketchy and inadequate conception is aiming at?

In the case of innate abilities, no matter what dispositions a mechanism happens to have, what determines its abilities is what it was selected for doing.² In the case of learned abilities, what natural selection selected for was the ability to learn in a certain way. It selected for mechanisms that became tuned through interaction with the environment to do things of useful kinds. For an organism to know how to do A as a result of learning is for it to possess an intact mechanism that is biologically designed to be tuned to do things like A and that has been tuned to do A as designed. That is, it became tuned in the same manner, following the same principles, as its successful ancestors when they were learning to do similar kinds of things.³ We humans possess at least a number, possibly very numerous, different kinds of learning mechanisms, including various mechanisms for trial-and-error learning, learning by association, by imitation, by figuring something out, and so forth. Each of these mechanisms works in accordance with its own principles, tailoring learned abilities in its own manner. To know how to do something as a result of learning, one must have a disposition to succeed in doing it under the conditions one learned under that afforded previous successes, or under the conditions that helped to tune one to do it by affording successes, and this learning or tuning must have been of the kind the learning mechanisms involved were selected for doing. Otherwise one does not know how, but does what one does only by accident.

Some learning mechanisms are extremely general in their possible applications and some are extremely quick. Of very general application, for example, are our abilities to learn something by figuring it out or thinking it through, using materials gained from previous experience, including knowledge of our prior abilities. So we often claim correctly that we know how to do something that we have never tried to do, for example, that we know how to get to Plainfield, or that we know how to make a certain kind of repair. If conservative, we may say only that we think we know how, but we are often right that we know. Similarly, observing only once that you have a certain capacity can immediately turn it into an ability. Anything that you find out you can effect immediately

2 Things that a mechanism was selected for doing are its "proper functions," as defined in (Millikan 1984, Chapters 1 and 2). Those concerned about the heavy but informal use of the notions of "function" and "design" throughout this book are referred to these chapters, where a tight definition is offered,

3 Abilities that have been learned correspond to mechanisms that have "adapted and derived" proper functions. For details on adapted and derived proper functions, see (Millikan 1984, Chapter 2) and (Millikan in press b).

becomes an ability. Having observed that stirring the red strawberries into the vanilla ice cream turns it pink, the child knows how to make pink ice cream.

We need not think of abilities, then, as characteristically derived from elaborate specialized training or tuning histories. They are, however, often derived from very numerous prior abilities that have been strung or blended together. The idea that one might count the number of a person's abilities, or count the abilities that go into a certain activity, often is not really coherent. Like patterns, however, or like patches of ground, abilities can be clearly distinguished and designated even when they have no clear criteria of individuation.

' 4.8 the Extensions of Substance Concepts

In the last chapter (' 3.6) I claimed that it is not the purposes of individuals but the subpersonal biological functions of their inborn concept-tuning mechanisms that connect their substance concepts with certain extensions. Similarly, the function of my eyes or my liver is not determined by my intentions for their use. We must suppose that natural selection has endowed us, specifically, with the ability to learn to identify substances, or, more likely, with a variety of abilities, each specific to learning to identify members of a different domain of substances (see Chapter five) or, at the very least, the ability to acquire abilities of the latter kind through learning. Considering the centrality of the ability to reidentify substances for any animal that acquires and applies any kind either of practical or theoretical knowledge (' 1.4), the likelihood that natural selection has tailored capacities very specifically to this purpose in such animals is about as high as the likelihood that it has tailored their capacities specifically for obtaining nourishment and mates.

Granted this, we can begin to answer the question how specific substances become assigned to specific substance conceptions. How the extensions of substance concepts are fixed. Chapter Five will fill the picture out more with some speculations about how one's abilities to reidentify specific substances are learned or developed, and Chapter Fourteen will complete the discussion of how conceptual abilities are focused on corresponding substances.

We encounter substances, better, we encounter natural signs of substances (see Appendix B) in perception (or in the speech of others (see Chapter Six)). We bring to bear certain innate abilities, or we bring to bear general skills acquired earlier in life (grasp of substance templates, grasp of general methods applicable to specific substance domains), or we proceed by trial and error learning in attempting inductions over encounters. In these ways we often manage to keep track, and we learn better how to keep track, of when we are encountering the same substances again. Thus the information gathered about each substance is brought together and brought to bear on future encounters with it. The important point is that the way this all happens is no accident. That is what it means to say that it is done using "abilities" and "skills" and by "learning." It is all done according to principles by which feats of the same abstract sort were accomplished by our ancestors, thus accounting for selection of the innate abilities and learning mechanisms responsible.

Now there will be times when a substance is misidentified, and for every

substance, of course, innumerable unrealized dispositions to misidentify it under adverse conditions. There also may accumulate considerable misinformation about certain substances. These mistakes occur because conditions on which past successes in keeping track or learning to keep track depended (historically enabling conditions for the cognitive mechanisms' development and use) are not currently present. Unless the cognitive mechanisms are malformed or damaged,⁴ mistakes are always caused by the presence of abnormal conditions, not necessarily in the sense of unusual conditions, but merely of conditions other than those under which the utilized mechanisms were operating on those specific occasions that historically afforded successes, hence accounted for their selection or for their tuning through learning. If current conditions were exactly the same as the historically enabling conditions for these mechanisms, they would obviously succeed. This is what distinguishes one's dispositions to identify from one's abilities to identify. It separates acts of misidentification from correct identifications and dispositions to misidentify from dispositions to identify correctly.

In clear cases, then (I will soon discuss some that are less clear), a substance concept has as its referent or extension a substance encountered at the very start of such a process of keeping track nonaccidentally and that fits the general abilities to keep track that have been brought to bear. It causally originates from encounters with the substance to which it refers, and the general abilities that are brought to bear in attempting reidentification determine the substance's ontological category. They determine, for example, whether it is the function of the concept to reidentify a stuff or an individual or a kind of individual, and if the latter, what sort of substance template is involved (see ' 14.1). The concept is a concept of A, rather than B, not because the thinker will always succeed in reidentifying A, never confusing it with B, but because A is what the thinker has been roughly keeping track of and picking up information about not accidentally but in accordance with abilities, skills, and know-hows.

Here is another way to approach the same matter, but that leaves it more open whether to buy into the historical theory of the nature of cognitive abilities I have proposed. Traditionally, it is supposed that what a substance concept is of is whatever fits certain features or properties it represents its extension as having. Its extension is what fits these properties, even though you may not always be able to recognize exactly when something does. In this way the concept can have a definite extension, even

⁴ A definition of normal structure for a device with proper functions including the functions corresponding to abilities is offered in (Millikan 1984) Chapter One.

though you sometimes make mistakes is recognizing that extension. But then the prior question can be raised, what determines which features or properties are the ones you are representing and trying to recognize? This cannot be determined by your representing prior properties that the represented properties must themselves have, without regress. Nor can it be determined merely by your dispositions to recognize these properties, since no one is infallible at recognizing properties any more than substances.

A standard reply is that we do recognize certain properties infallibly when we are in "normal conditions," and that this is what defines what properties our property concepts represent. For example, red is whatever appears red to me under normal conditions for seeing colors. But how then do we define "normal conditions," such that they are, for example, appropriately different for seeing the shapes of big things like mountains and the shapes of small things like fleas, appropriately different for hearing loud sounds and soft ones, and different for hearing sounds, for seeing colors and for identifying tastes, and so forth (consider how tea tasters must prepare themselves)? We must take care that "normal conditions" do not turn out to be, just, the conditions under which one perceives each of these various properties correctly, for that would be marching in place. On the other hand, if there is any noncircular way of defining "normal conditions" for the perception of various properties, we should be able to use the exactly same technique to define "normal conditions" for keeping track of various kinds of substances. The two problems are exactly parallel. How then should we understand these "normal conditions"?

Biologists are usually concerned, first, to understand normal function. They may be interested in diseases or other abnormal functions, too, but these are defined relative to normal function. I have proposed that normal function, in this context, is best defined relative to a history of natural selection,⁵ but you may supply your own favorite theory of normal function, should you have one, and it will serve the argument as well. My claim here is only that whatever normal function is, philosophers and cognitive psychologists, too, are, or should be, interested in it as well as biologists. I add that for the most part biological items require to be in rather definite conditions in order to perform normally. My own preference is then to define "normal conditions" relative to selectionist history also. As conditions under which that function was performed historically such as to be selected for (Millikan 1984, 1993a). But if you can supply a better definition, there is no objection. The point is that if we can give a definition of normal biological function and normal conditions for performance of this or that function, we can apply it also to performance of psychological functions, such as the development of substance concepts and their application. We can say what normal conditions for keeping track of substances are.

Grant then that there is a normal way or normal ways for development of substance concepts to occur, perhaps different for different substance domains. Or what is the same thing, grant that normal developmental cognitive psychology is a viable

5 1984, chapters 1 and 2; 1993a chapters 1 and 2.

field. There will be a normal way or ways that the child or adult first recognizes the manifestations of a new substance impinging on their perceptual organs and a normal way or ways that they attempt to keep track of that substance, or learn to keep track of it better, and so forth. And there will be normal conditions for success in keeping track, in building conceptions adequate to the substance. When everything goes exactly normally, then, there will be no question what the concept is of, even if there is a disposition to apply it incorrectly under conditions that are not normal for expression of these abilities.

But problems can arise when things do not go exactly, when they deviate from the ideal. It can happen that each of two substances is kept track of in a normal way over a variety of encounters, but that there are also mistakes made so that information gathered from both gets collected under the same concept. For example, one might have two people "mixed up" or "confused" in one's mind. Similarly, mass and weight were not distinguished throughout most of the history of science. More than two substances might also be entwined under one concept. If it is not definite which among various similar, closely related, overlapping or nested substances was the one primarily responsible for the information that has been gathered and/or for the tuning of the (would-be) tracking dispositions, then the concept is equivocal or vague. Two or more are being thought of as one. It is likely that normal development of many kinds of concepts involves a process of differentiating between substances originally confused together. Perner calls this a process of "focusing reference" (Perner 1998). It is tempting to interpret much of the history of science as an attempt to focus reference, for example, distinguishing weight from mass and oxygen from other oxidizers.

We also can imagine much more serious confusions than simple equivocation where it is not clear what if any substance or substances have been kept track of at all. Biological items are, in general, defined relative to an ideal. A diseased or damaged or malformed heart is a heart none the less because of historical relations it has to hearts that performed normally. But having described how normal hearts are structured and how they function, it is of no interest to biologists how far away from that ideal a thing has to be before one stops calling it a "heart." There are no exact borders of the substance heart in nature, and the biologist is concerned with nature. Similarly, I suggest, to press the question, in sufficiently abnormal cases, "But please, really, what is the extension of this person's substance concept?" is pointless. What, if anything, for example, was the referent of "phlogiston"?^{6, 7}

6 I take these matters up again at the end of chapter 14.

7 I am grateful to Andrew Milne for some very helpful suggestions on this chapter.