

Chapter Three

Classifying, identifying, and the function of substance concepts

§3.1 Orientation

Substances, as I have described them, are whatever one can learn from given only one or a few encounters, various skills or information that will apply on other encounters. Further, this possibility must be grounded in some kind of natural necessity. The function of a substance concept is to make possible this sort of learning and use of knowledge for a specific substance. For this, the cognizing organism must be able to recognize the specific substance under a variety of different conditions, as many as possible. It needs to do this, first, to grasp that the substance it is learning about over various encounters is one and the same so that knowledge of it can accumulate and, second, so that the accumulated knowledge can be applied. For substance concepts to be employed in the service of theoretical knowing —employed for knowing that rather than knowing how— the substance must be represented in thought in a univocal way, the same substance always represented as being the same. This makes possible a stable, unequivocal and nonredundant inner representational system.

The ability to recognize what is objectively the same substance again as the same despite wide variations in the faces it shows to the senses is necessarily fallible. Although you surely have many ways of identifying each member of your immediate family —similarly for water and for cats— there will always be possible conditions under which you would misidentify them, mistaking them for someone or something else. If a concept is genuinely a substance concept, if its extension is really a substance, this extension is not determined by one's fallible dispositions to recognize it. These dispositions, given any of numerous adverse conditions, would break up the extension of the substance or mix it with other things. The extension of a substance concept is determined not by one's dispositions (rightly or wrongly) to recognize it but, first, by the real extent of the substance in nature.

The purpose of this chapter is to clear the ground of certain rubble left by classical and contemporary theories according to which dispositions to identify do determine the extensions of substance concepts.¹ Chapter Four will then begin the foundation for another explanation of how a substance concept hooks onto its extension —onto a substance, whole— and how it is determined onto which substance it is hooked.² It will begin to establish a version of what is sometimes called "direct reference" theory.

To understand how the extension of a substance concept is determined, first we must understand more exactly what the functions of substance concepts

¹ Appendix B also concerns this matter. It concerns information theories of mental content according to which to have a concept of Xes, one must be able to "discriminate" between Xes and all other things..

² The explanation will not be completely finished, however, until the full implications of the theory of abilities to be introduced in Chapter Four have been drawn out in Chapters Thirteen and Fourteen.

are, hence to what sort of things in nature they need to be hitched in order to serve these functions. The standard view has been that terms for kinds and stuffs correspond to capacities to classify instances falling under these terms. I claim instead that most such terms correspond in the first instance to abilities to identify substances, and that they are only secondarily used as classifiers. The result of the standard view has been a thorough confusion between two quite different kinds of functions, the functions of identifying and the functions of classifying.

§3.2 the Functions of Classifying

Sharpening the distinction between the terms "identify" and "classify" somewhat for expository purposes, the difference between identifying and classifying lies both in purpose and in the psychological structure of these acts. The purposes of classification and identification are hopelessly entangled, for example, in the following contemporary descriptions by psychologist's of the functions of what the authors call "categorization" and "concepts":

Categorization...is a means of simplifying the environment, of reducing the load on memory, and of helping us to store and retrieve information efficiently. (Markman 1989, p. 11)

Without concepts, mental life would be chaotic. If we perceived each entity as unique, we would be overwhelmed by the sheer diversity of what we experience and unable to remember more than a minute fraction of what we encounter. And if each individual entity needed a distinct name, our language would be staggeringly complex and communication virtually impossible." (Smith and Medin 1981, p. 1.)

...concepts are used to classify...if you know nothing about a novel object but are told it is an instance of X, you can infer that the object has all or many of X's properties... (Smith and Medin 1981, p. 8).

A good classification system aids efficient storage, retrieval and transfer of information, or efficient storage and retrieval of the objects classified, or, in a different but related way, efficient communication. It aids the efficient storage and retrieval of what we already know when we use dictionaries, encyclopedias, telephone books, guide books, filing systems, classification systems in libraries, and so forth. It aids in putting objects away where they can be found again in grocery stores, hardware stores, museums, home workshops, and again, libraries. It aids in communication in the following way. Shared classification systems allow one person efficiently to convey enough information about a thing for another to retrieve it, either literally or from memory. For example, saying it is "the red book on my table in my study" that I want you to bring will get me what I want only if we share a way of classifying things into those that are studies and those that are not, those that are mine and those that are not, and into books

and nonbooks, red things and nonred things, tables and nontables. Then I can swiftly convey enough for you to retrieve the object I intend.

These being its principle uses, an ideal general classification system, designed to cover a general domain, will draw sharp lines around the classes it contains, so that each member in its domain falls determinately either in or out of each class. General domain classification systems are used, for example, by libraries and grocery stores. These lines need not cut between items in all possible worlds, however, but only in the actual domain where the classification system is to be used. Classification for purposes of communication, on the other hand, does not correspond to a single general domain. The domains that are involved in ordinary informal communication typically are severely restricted by context, varying radically from one speaker-hearer pair and from one occasion to another. For this reason, words whose natural extensions have very vague boundaries can still be used in specific communicative contexts to classify objects precisely. I refer to what I want simply as "red," but given the books on my table it is clear enough what object I want, even though the entire domain of red things shades off gradually into pink things, purple things, orange things, and so forth. If both my hearer and I know my intended local domain and I choose my words well with reference to it, I can often use very vague words to effect accurate hearer classifications in context. It remains true, however, that clear boundaries between classes relative to the actual members in the domain of its use is helpful for all of the functions of classification.

An ideal general classification system also has a tree or a grid structure, so that each item can be located within it by answering a determinate set of questions in order. This assures that each member of the relevant domain has one and only one location within the classification system, hence that it can be efficiently put away and retrieved. This is true also for classification used for communication. In asking for "the red book on the table in my study," each of my descriptive words, used here as classifiers, partitions my intended domain into two classes in such a way that, taken together, they separate off just the intended (unit) class.

The initial data for a classification task must include a specification of each property of each object to be classified that is relevant to its classification. It must be possible to answer each of the questions determining its classification. A librarian would not try to classify a book, for example, without carefully examining its contents, and to classify an object as a red book on the table in my study you must know it is red, a book, on a table, in a study, and that the study is mine —every one of these. On the other hand, most of the myriad properties of any object will not be relevant to a given classification task. Librarians don't need to note the colors, numbers of pages, numbers of illustrations, publishers, type fonts or, usually, the shapes and sizes of the books they classify. But the properties that define the classes in the classification system being used obviously do need to be determined, whether by observation or by inference, as either present or absent for each instance of a domain prior to classifying it. Consonant with this, in classical "categorization" experiments in psychology, since these are usually taken to be studies of classification, all properties of each

"stimulus" and each "test item" considered to be relevant to the classification to be learned are clearly exhibited to the learner.

Now consider Smith and Medin's claim quoted above that "...concepts are used to classify...if you know nothing about a novel object but are told it is an instance of X, you can infer that the object has all or many of X's properties...." Suppose that the extension of "X" is, as suggested, merely a class and not also a substance. In that case, you will not have to do any inferring in order to know that the object "has all or most of X's properties." For in that case, the properties of the object must include all of those properties used to determine that it falls in class X, and no other properties will be inferable in a grounded way. This is because grounded inferences from properties of some members of a class to other members are possible only if the extension of the class is a substance. Being a substance is the same thing as being something that grounds inferences of this kind; if no substance, then no grounded inference. Concepts that merely classify, and do not also identify substances, contain exactly as much information as is analytically put into them, no more and no less. If without concepts of this kind we would be "unable to remember more than a minute fraction of what we encounter" we would be just as unable with them. Use of words that are shorthand for strings of memorized properties may make classification and communication more efficient, but it cannot make memory or thought more efficient.

Similarly, when a classification system is used for storing away information in encyclopedias or libraries, it is not the classification system that contains the information. The information is in the encyclopedia entries and in the books, not in the classification system. In order to retrieve the information one must first find the right entries or retrieve the right books. And it is exactly the same when I use words to classify what I am talking about. By classifying my subject matter as a "red book on the table in my study" I hand you tools with which to circumscribe that object, but if you are to retrieve any more information about it than is already contained in my classification, you will either have to go and find the extension itself and inspect it, or find it in your memory so as to retrieve things you already happen to know about it. (On the other hand, exactly because the word "book" is not merely a classifier but corresponds to a rough substance, you also may retrieve from memory what you know about books generally and apply it to this case.)

§3.3 the Functions of Reidentifying

Contrast now the functions of reidentifying. Reidentifying is required not primarily for information storage, retrieval and transfer, but for information acquisition and information use. Rather than knowing its relevant properties in advance, as when one classifies something, one identifies a substance in order to come to know its properties. Identifying is necessary in order to collect together over time knowledge of a thing's properties, hence in order to know its properties on particular occasions, since many of these properties are manifest on encounter only some of the time. Identifying is also necessary in order to apply one's knowledge of things. One applies one's knowledge by managing to recognize a substance on the basis of whatever properties do happen to be

currently manifest, and then applying one's prior knowledge of others of its properties, properties not currently observed.

Notice first that these functions do not require the substance recognized to have sharp boundaries. Grasp of substances very often affords knowledge that is not invariant over the substance's entire extension. If the edges of the substance are vague, the variance is more marked toward these edges. Well-fashioned substance concepts, well-fashioned abilities to exploit substances as objects of knowledge, will include the ability to recognize a difference between more central and less central areas in the extent of the substance, and to portion out degrees of reliance on accumulated knowledge of the substance accordingly. When, as is often the case, the boundaries of substances really are vague, obviously there is nothing, other than taking this into account, that the organism can do about it. Artificially imposing precision will not help. By contrast, ideal general classification systems, I have said, are more efficient if precise relative to the entities actually in their domains. Where substance boundaries are vague in nature, the purposes of classification are sometimes served by drawing artificial boundaries around the extensions of these substances. For certain classificatory purposes, for example, what counts as war and who counts as a member of the working class or as a full-time student may be quite sharply but artificially defined.

Notice, second, that substances are not generally organized into tree or grid structures. This was argued in §2.6.

Notice, third, that unlike a task of classifying, the task of identifying a substance doesn't require that any one particular set of the substance's properties be known or manifest to one, or that different people should use the same properties of the substance in order to identify it. Any of very numerous means of recognizing the substance may be applied. Each of these methods will be fallible in principle. Identifying a substance and exploiting its possibilities is as fallible as any other practical activity one engages in. One may always stumble and fall.

§3.4 Understanding Extensions as Classes Versus as Substances

Although substance concepts, hence words for substances, can be used for the purpose of classifying, the reverse does not hold in general. There is a big difference between understanding something merely as a class and understanding it as a substance. Conceptions used to classify need only carve out some clear unequivocal extension within the domain to be classified. Conceptions that govern substance concepts must locate genuine ontological grounds of induction. A substance concept is distinguished by the role it is ready to play, accumulating additional means of identification, and anticipating certain kinds of inductions as likely to hold. A substance concept will be successful only if there really is some substance out there it is hooked into. One reason it is an error to place great value on operational definitions in science, for example, is that operational definitions, as such, are merely classifiers, hence do not necessarily correlate with substances. But it is substances rather than classes that are of interest to science.

Because conceptions filling out substance concepts can sometimes be used also as conceptions of classes, words for substances can vacillate between being understood as standing for substances and as standing only for classes. When confidence is lost in the reality of a substance or in the univocity of a substance term, it may begin to be used in a strictly classificatory way. For example, terms for many mental disorders have vacillated over the years between being understood as capturing substances, naming single diseases for which single etiologies and therapies might eventually be discovered, and as being merely classificatory, defining useful groupings of symptoms for efficient transfer of information.

Nominals that are used only to express concepts of classes typically are complexes built out of prior terms. They wear their analytical natures on their sleeves. Their extensions are functions of unions and disjuncts of the extensions of the prior terms that compose them. Typically, this sort of construction will be built up in the same way by all who understand the syntactic forms of the language. Thus although prior conceptions attaching to the element terms in the complex may differ from one speaker of the language to another, the mode of construction of complex conceptions out of prior conceptions will be common. These terms express analytical concepts, concepts only of classes. Substance concepts, on the other hand, are synthetical. A person's conception of a substance may employ prior concepts used in the process of identifying, but the substance concept is not equivalent to any mere function of prior concepts.

But there are, of course, many exceptions to the rule that synthetical concepts are expressed with simple nouns, analytical concepts with compounds. Consider the term "red sulphur." Red sulphur is not just sulphur that is red, but an allotrope of sulphur, a substance in its own right with its own suite of properties different from other forms of sulphur. On the other hand, red sulphur also happens to be the only substance that is both red and (pure) sulphur. Does the nominal "red sulphur" correspond, then, to a synthetical or an analytical concept?

Whether a word expresses an analytical or a synthetical concept may sometimes depend on the user. For some people the concept for "red sulfur" may be synthetical and for others analytical. A person who did not understand that sulfur that is red happens to be a substance in its own right would only have an analytical concept corresponding to the term "red sulfur." Accordingly, they would never recognize any part of the extension of "red sulfur" in any way other than by noting that it was sulfur and also noting that it was red. And they would not attempt inductions from samples of red sulfur to other samples of red sulfur that they would not have attempted either from samples of red to other samples of red or from samples of sulfur to other samples of sulfur. On the other hand, a person might instead have a synthetical concept of "red sulfur." That is the kind chemists have, for example. More interesting, it would be possible to have a concept of this stuff, this substance, without knowing either that it is red or that it is sulfur. One might recognize it as the sticky so-smelling substance typically found in such-and-such context, and be surprised to learn that it is always red, and that it is a kind of sulfur.

Similarly, should "Californian" correspond to a vague sort of substance, as suggested in §2.4, then although the "-ian" suffix suggests an analytical concept, there will be at least two ways to have a concept for "Californian," one analytical and the other synthetic. It is less plausible, of course, that one might have a reliable way of identifying Californians that did not depend on first determining that they came from California. Not every legitimate substance is reliably identifiable in multiple ways, if one counts as "identifying in the same way," using descriptions employing concepts of the same things in the same way. (These prior concepts might each be governed by variable conceptions, however.) Especially, as is true in this particular case, ways of identifying that relate directly to the real ontological ground of a substance may be uniquely reliable. Whether a substance concept is legitimately a one criterion concept or not does not rest on something ephemeral called "the rules of our language" (wherever they live) but lies in nature. If it is necessary that a vixen be a female fox, the deep reason is that it is sets of fox chromosomes that include two X chromosomes, copied from prior sets of this sort (close enough), that are responsible for causing the characteristic likenesses among vixens. Being a female fox is the real essence of vixenhood, nor is this a purely a priori matter.

One last difference between identifying and classifying. Classifying requires recognizing that a predicate applies to some definite subject. Suppose, for example, that we tried to model the act of classifying individual objects into the red ones and the green ones as, just, responding in one way to the red individuals, in another to the green individuals. What would determine that these responses constituted the classifying of individuals, rather than of time slices of individuals, or facing surfaces of individuals, or dye stuffs found on the surfaces of individuals, or patterns of ambient light impinging on retinas? You can't classify without some grasp of what you are classifying. You have to be able independently to think of the object you want to classify before you can classify it. Identifying an object, on the other hand, does not always require something conceptually prior. Identifying a substance as animals do, merely for practical purposes, requires only that behavioral responses the animal is disposed to learn by employing the concept should be appropriate to that substance, that is, they should be responses that are effective because of the properties or dispositions of that substance. True, identifying a substance for theoretical use does require that you have some appropriate predicate concepts, ones that you understand as applicable to, whether or not they are true of, the substance. I am not claiming that the only things we reidentify are substances.³ But whatever it is that one classifies, it is clear that the capacity to think of members within the domain to be classified is more fundamental than the ability to classify. Identifying is a skill prior to classifying.

³ In (Millikan 1984) Chapter 16, I talked quite a lot about reidentifying properties, and the analysis of the act of reidentifying to be offered below in Chapters 9 through 12 applies equally to concepts of substances and concepts of properties.

The traditional view among philosophers and psychologists has been that the ability to apply a term for a kind or a stuff is an ability to classify. This view has often taken the form of assuming that kind terms and stuff terms are descriptive, each corresponding to some sort of configuration of properties. If the concepts corresponding to these terms are concepts of configurations of properties, their extensions are naturally determined analytically as a direct function of the extensions of those properties. Both among contemporary psychologists and also in some philosophical circles, this view is still much the most common form of a more general position I will call "conceptionism." Conceptionism is the view that the extension of a concept or term is determined by some aspect of the thinker's conception of its extension, that is, by some method that the thinker has of identifying it. I am fully in charge of the extensions of my concepts; whatever I am, after due consideration, disposed to apply them to is what they are concepts of.

One way of identifying a substance, of course, is by means of the knowledge that it has certain properties or falls under a certain description peculiar to that substance. The classic form of conceptionism holds that the conception that determines the extension of a substance term is such a set of properties or such a description. We can call this classic view "descriptionism." Another form of conceptionism holds that the extension of certain concepts or terms is determined by means of identification procedures not employing prior concepts of properties. Concepts whose extensions are thought to be determined in this direct manner are sometimes called "recognitional concepts." Thus Fodor remarks "if a concept is recognitional, then having certain kinds of experience would, in principle, show with the force of conceptual necessity that the concept applies" (Fodor 1999). "Conceptionism" in either form contrasts, of course, with "direct reference" theories of conceptual content.⁴

§3.5 Descriptionism in the Psychological Literature

Conceptionist views—the view that substance concepts are basically classifiers, their extensions being determined by dispositions to apply them—underlie a surprising proportion of the masses of recent work on "concepts" and "categories" in the psychological literature. It will be worth spelling this out in some detail. I will try to show, indeed, that throughout the changing variety of competing theories of concepts and categorization developed by psychologists in the last half century, the theoretical assumption of conceptionism, generally in the form of descriptionism, has managed to go completely unchallenged. This is true despite the fact that Putnam's and Kripke's famous arguments (or at least their conclusions) against descriptionism (Putnam 1975, Kripke 1972) have been rehearsed numerous times in the psychological literature, and despite a number of brave attempts to integrate these insights into the psychological tradition (Neisser 1987 chapter 2, Lakoff 1987, Markman 1989, Keil 1989, Gelman &

⁴ In (Millikan 1998a and 1998b), I referred to both kinds of conceptionism as "descriptionism." This resulted in understandable confusion on the part of several commentators. Here I am shifting to what I hope is a more perspicuous terminology.

Coley 1991, Komatsu 1992). The difficulty, I believe, results from the fact that Putnam's and Kripke's insights were almost entirely negative. They told us how the extensions of certain substance concepts are not determined, but they supplied no adequate theories of how they are determined. Moreover, the tentative positive views that they offered focused more on the extensions of words in a public language than on the nature of concepts, leaving obscure the nature of the psychological states or processes that would constitute an understanding of the meanings of the words discussed. Thus they offered little aid to psychologists. One aim of this book is to help remedy that.

The descriptionist holds that the referent or extent of a substance term is determined by its falling under a description associated with the term by the term user. Certain properties, relations, facts about origins, facts about causes, similarities to prototypes, similarities to given exemplars, and so forth—certain "information" about each portion of the extent—determines it to be a portion of the extent, and the thinker or the thinker's "mental representation" determines which information is to play this role. In the psychological literature, this view is sometimes found in caricature in the statement that concepts are features or properties, for example, "many properties are concepts themselves" (Barsalou 1987, p. 129).

Using the concept chair as his example, Komatsu (1992) expresses what he claims is the most general question that psychological theories of concepts have attempted to answer thus: "...what information, very generally, is represented by the concept chair, so that people are able to reason about chairs, recognize instances of chairs, and understand complex concepts..." (1992, p.:500, italics mine). Building on (Medin and Smith 1981, 1984), he applies this descriptionist formula to each of five accounts of concepts:

the classical view (e.g., Katz 1972, Katz & Fodor 1963)...the family resemblance view (e.g., Rosch & Mervis 1975)...the exemplar view (e.g., Medin & Schaffer 1978)...the schema view [Komatsu later cites (Bartlett 1932, Piaget 1926, Minsky 1975, Rumelhardt 1980, Schank & Abelson 1977, Winograd 1975 and Neisser 1975)]...the explanation-based view (e.g., Johnson-Laird 1983, Lakoff 1987, Murphy & Medin 1985) [later he cites, e.g., Gelman and Keil] .

Descriptionism is most obviously compatible with nominalism, the view that the members of the kinds that words name are grouped together either conventionally, according to the dictates of culture, or according to patterns natural to human perception and thought. For example, heavily sprinkled throughout the literature we find references to "learning about people's categorization decisions." On this view, the descriptions that govern concepts have their source either in the conventions of society, or in peculiarities of the human perceptual and cognitive systems, that is, in ways it is natural to us to generalize. For example, in classical studies of concept learning, subjects were typically set the task of learning imaginary categories defined by arbitrarily

chosen sets of properties, and many studies exploring family resemblance or prototype or exemplar views of categorization have also set arbitrary tasks. The view that the human mind has its own ways of imposing various groupings of things into kinds, ways that languages must respect in order to be learnable, has been evident especially since Rosch's work on color categories (e.g., Rosch 1973, 1975). In this tradition, the psychological problem concerning categorization is understood to be that of ferreting out exactly what these psychologically imposed principles are—those principles in accordance with which children or adults "prefer to sort" (Markman 1989). Thus, for example, Lakoff subtitles his 1987 book, "What Categories Reveal about the Mind."

But descriptionism need not be allied with nominalism or conventionalism. It also has been combined with realism about human categories. The realist holds that many of our categories correspond to kinds that are grouped together by nature independently of mind. As we acquire categories we learn not merely, say, how to communicate with others, but how to grasp structures that were already there in nature. The view of substances that I am advocating is, of course, a variety of realist view. It might seem that there is an incompatibility between realism and descriptionism. If the extent of a category is determined by nature, then it cannot be determined by fitting a certain description associated with a word. But in fact there are a number of ways in which realism and descriptionism have been combined.

The simplest way is to take the extent of a substance term to be fixed by one, or a set, of definite descriptions of the substance. Whether it is supposed that the description is used rigidly or nonrigidly makes no difference in this context. In either case, the thinker entertains a prior description that determines the extent of his word or category. Thus the classical 20th Century view was that Aristotle himself was a natural unit in nature, and that to have a concept of Aristotle was to capture him in thought under a description such as "the teacher of Alexander," or under a suitable disjunct of descriptions. Similarly, there has been a tendency in the psychological literature to interpret Kripke's (1972) and Putnam's (1975) apparently anti-descriptionist views on the meaning of proper names and natural kind terms as invoking definite descriptions on a meta-level. Kripke is thought to have claimed that the referent of a proper name N is fixed in the user's mind by the description "whoever was originally baptized as N" and Putnam is thought to have claimed that the extent of a natural kind term is fixed for laymen by the description "whatever natural kind the experts have in mind when they use term T." This is what Fumerton calls "Russelling" a theory of direct reference (Fumerton 1989). It transforms it, of course, into a descriptionist theory instead.

Theories that language categories are organized "probablistically" (Medin 1989) by family resemblance or by reference to prototypes often combine realism with descriptionism. Families and prototypes are usually taken to center over highly correlated properties, and these correlations are taken to be empirically discovered. Thus prototype theory is naturally compatible with the view that many concepts end up paired with real kinds. But probablistic theories are regularly interpreted as explaining only how the learner's experience

generates the category. Then the actual extension of the category is taken to be determined, not by the real extent of a kind, but by how the learner is inclined to classify new examples. The same is true of exemplar theories and for variations on these two views. Thus Billman suggests that we should compare and test psychological models of structure and processing of concepts by examining the function from "learning instances plus the target items to categorize" to "the set of possible category judgments" (Billman 1992, p. 415, emphasis mine) and Ward and Becker state that "category structure" can mean "the set of items that the learner considers to be members of the category in question (i.e., the category extension)" (1992, p.454). Made explicit, the idea here seems to be that experience with a natural kind may inspire the category, but the category extent is determined by the thinker's potential decisions on exemplars. When all goes well, our psychologically determined kinds may contain the same members as the natural ones, that is all. Similarly, the realists Gelman and Byrnes tell us, explicitly making reference to Chomsky's theory of innate grammar, that "[w]e can determine how languages and conceptual systems are constrained by examining the forms and meanings that children construct, and which errors they fail to make" (1991, p. 3). That is, it is the child's inclinations that constrain the concepts.

Most explicitly realist in their approach to concepts are contemporary researchers holding what Komatsu calls an "explanation-based view" of concept structure. Komatsu characterizes this view by quoting Keil:

No individual concept can be understood without some understanding of how it relates to other concepts. Concepts are not probabilistic distributions of features or properties, or passive reflections of feature frequencies and correlations in the world; nor are they simple lists of necessary and sufficient features. They are mostly about things in the world, however, and bear nonarbitrary relations to feature frequencies and correlations, as well as providing explanations of those features and correlations. If it is the nature of concepts to provide such explanations, they can be considered to embody systematic sets of beliefs—beliefs that may be largely causal in nature. (Keil 1989, p. 1)

Note that the view is not just that concepts designate kinds for which there exist explanations of property correlations, but that the concept actually consists in essential part of an understanding or, looking beyond page 1 of Keil's text, a partial understanding of these explanations. Interpreting this in the terms of the last chapter, the concept consists in part of a partial understanding of the ontological ground of induction that underlies the concept. Of particular interest to the explanation theorists, for example, has been Medin's work showing that people behave as though believing that beneath their categories there are hidden essences making the things in the categories what they are (e.g., Medin and Ortony 1989). Keil, Carey, Gelman and Markman are among those who have done very interesting work tracing the development of children's natural

kind concepts and artifact concepts, for example, documenting the transition from reliance on superficial characteristic properties for identification of these kinds to use of rudimentary and then more sophisticated "theories" about the underlying causes of the unity of the kind.

These advocates of explanation-based views have remained strongly influenced by the characteristic mid- twentieth century doctrine that a "theory" is a set of inference connections among concepts, that the net of theory in which a concept is caught up determines its "meaning," and that the meaning of a concept determines its reference. Thus to introduce or change theories threatens to change meanings hence reference:

How can one be sure that one is even talking about the same concept at all if all concepts are relative to theories?...We do not want every change in theoretical beliefs to make the concepts embedded in them completely different from those that were embedded before the change; yet no precise method is offered [by Smith, Carey, and Wiser 1985] for making a decision....These are difficult issues, and it is hardly surprising that they are not yet resolved. (Keil 1989, p. 21-2)

Following Smith, Carey, and Wiser, Keil speaks of "'tracking' concepts across theory change" and agrees with them that probably "[d]escent can be traced...because of several properties of theories that stay fixed through change" (Smith, Carey and Wiser 1985, p. 182). And he agrees with Fodor that it is not obvious how the classical view could be true that "children and adults could have different kinds of concepts for the same terms," for that makes it seem as though [quoting Fodor 1972] "they must misunderstand each other essentially" (Fodor, p. 88; Keil 1989, p. 15-16). Again, the view here is conceptionist. There is no suggestion that the extent of the concept, its "meaning" in the most fundamental sense, might be directly fixed by the extent of a natural unit in nature, reference remaining the same while conceptions change. (For an exception to this, see Gopnik and Meltzoff 1996.)

My claim is that all this traditional work, supposedly on the nature of concepts, has actually concerned merely conceptions. For example, changes in theories about the underlying grounds of induction for specific kinds of substances are changes in conceptions, not in what the corresponding substance concepts are of. You cannot have a theory about something, say, about what makes dogs to be dogs (remember that Aristotle had a false theory about this) unless you can first think of the thing you would theorize about. Changes in your theory are not, then, changes in your concept or in its extension. It does not follow that these psychological studies, when they have concerned substance concepts, have not had, and will not continue to have, great value. Nor does it follow that they simply have not concerned meanings. Conception is definitely one of the things that "meaning" can mean. But these studies need to be reinterpreted as studies of conceptions if their value is to be secured. For example, their relevance to understanding word meanings will be

problematic should it be common for different speakers of the same language to have quite different conceptions of the same substance yet use the same word for it. And indeed, beginning in Chapter Six I will argue that radical divergence in speaker conceptions of exactly this sort are the rule rather than the exception.

§3.6 How Then Are the Extensions of Substance Concepts Determined?

If substance concepts are not just classifiers, so that conceptions of substances are not what determines the extensions of their corresponding concepts, how then are the extensions of these concepts determined? To argue that the extensions of substance words are not determined by conceptions and to explain how these words do hook onto their extensions instead was, of course, supposed to be part of Kripke's project in Naming and Necessity (1972) and Putnam's project in "The Meaning of 'Meaning'" (Putnam 1975). If Kripke was right, "Bill Clinton" does not attach to Bill Clinton by means of speakers associating with "Bill Clinton" any particular properties or relations, ipso facto not by means of associating with it any inner or outer causes, any essence, any particular kind of ontological ground. A proper name is not, as I have been putting it, a classifier but an identifier. Similarly, Putnam argued that to call a thing "water" or "Elm" is not to describe it. Natural kind terms do not work by being associated with properties. Rather, the extensions of "water" and "elm," like the extent of "Bill Clinton," are natural units in nature, units to which the concepts water and molybdenum do something like pointing, and to which they can continue to point despite large changes in the properties the thinker represents these units as having. Taking Lakoff's example discussed in §1.7, large changes can occur in the manner in which a child identifies cats, hence in the things the child is willing to call "kitty," without affecting the extension of the child's word "kitty." The difficulty, of course, is to cash out the metaphor of "pointing." Speaking literally, what determines which substance is the extension of a given substance term or concept on the present view?

Putnam spoke of "indexicality" rather than "pointing," and Kripke suggested (without actually endorsing it) a causal theory of the reference of proper names. As mentioned earlier, a difficulty with both of these suggestions was that they tended to collapse into more complicated descriptionist views (see Fumerton 1989). The situation was exacerbated by the fact that Kripke and Putnam both used a form of argument that seems to assume the very thing it is trying to disprove. Their arguments proceeded by offering examples, taken either from this world (Kripke) or from other possible worlds (Putnam), in which our intuitions cry out that wrong results follow from some particular classical view of the kind of description that determines reference. But if our intuitions are really the final judges here, that would certainly suggest that we have in mind what determines reference, and this brings us back either to a descriptionist view, or at least to a conceptionist view. Compare Russell's early view of demonstratives, where "this is a book" unpacks into something like "the thing at which I am pointing is a book." And compare Ned Block's view (1986) according to which the way we would make decisions about its extension in possible cases determines what sort of Kaplan-style character, what sort of function from possible worlds to extensions, a concept expresses, for example, whether its

referent is indexically determined as its cause, or as bearing some other relation to the concept.

Still, what else could possibly determine the extensions of our concepts if not our own intentions or dispositions? And if something else does determine the extension, what determines that this is what determines it, if not our prior intentions or dispositions? After all, what is an extension anyway? What is it for something to have an "extension"? Isn't an extension made into an extension by the fact that we apply the word "extension" in this way and not that? Aren't extensions determined by how we intend or are disposed to apply the term "extension"?

Is there a way to stop pulling at our bootstraps?

Consider: no one supposes that the function of vision is determined by the intentions of the individuals who happen to have eyes. Similarly, why should functions of the developmental processes responsible for concept formation and the functions of the concepts these processes shape be determined by the intentions of the individuals in which these processes happen to occur? It is not the purposes of individuals, but the biological functions—the unconscious purposes—of their inborn concept-tuning mechanisms that connect substance concepts with certain extensions.

I have proposed a theory telling what the most general function of substance concepts is. It is their job to make it possible to utilize substances as these are objectively defined in nature for purposes of gathering and applying information. In order to do this, they must include skills in reidentifying substances. Only in so far as they succeed in this task can they help us to proceed with successful inductions. My claim will then be that the extension of a substance concept is whatever substance in the world it is the job of that particular concept, given its particular phylogenetic and ontogenetic history of development, to be reidentifying or conceptually "tracking." Many mechanisms are involved in the development of a substance concept, and different kinds of mechanisms effect the development of concepts of different kinds of substances. Some speculations about these mechanisms, and about their particular phylogenetic and/or ontogenetic origins, will be offered in Chapter Five. But the rough idea is that the specific functions of the mechanisms or abilities responsible for originating a particular substance concept, whether these functions or abilities originated through evolution of the species or through individual learning, determine whether the concept is of anything definite and if so of what. This claim will need to be filled out a great deal before it will be plausible (or clearly understandable). The core of the project will be an analysis of what an ability, hence an ability to identify, is that does not equate abilities with dispositions (Chapter 4), and that explains how it is determined what a particular ability is an ability to do, even though the ability may be currently operating under conditions that fail to express it properly (Chapters 4 and 14).

You can call whatever the conception filling out a certain person's concept happens to corral, that is, whatever the person takes to be part of the concept's extension, by the name "the concept's extension" if you like. Humpty Dumpty was right about that. But then "extension" becomes a notion with very little

interest, and we will need to coin another term for the thing it was the real purpose of that particular conception to capture. A parallel would be to label whatever a frog happens to snap up with its tongue reflex—for example, beebees—as its "prey," and then have to coin another term to designate the things its reflex snap was designed to capture.