

ABSTRACT

I present a nondescriptionist theory of the nature of our most basic empirical concepts, concepts of what Aristotle called "substances": stuffs (gold, milk), real kinds (cat, chair) and individuals (Mama, Mount Washington). Surprisingly, our root concepts of these are identical in structure. Their extensions are natural units in nature, to which concepts do something like pointing, bypassing reference to properties the thinker represents them as having and theories about them. The difficulty, of course, is to cash out the metaphor of "pointing" in this context (c.f., Putnam's "indexicality"). A second aim is to show how this kind of concept interacts with language.

A Common Structure for Concepts of Individuals,
Stuffs, and Real Kinds:
More Mama, More Milk and More Mouse
Ruth Garrett Millikan

Philosophy Department
University of Connecticut
Storrs CT 06250

Key Words: categorization, concepts, externalism, natural kinds, names, basic-level categories, child language, theory of meaning, Putnam.

ABSTRACT

Concepts, taken as items that the psyche "acquires", are highly theoretical entities. There is no way to study them empirically without committing oneself to substantial preliminary assumptions about their nature. One aim of this paper is to show how, throughout the changing variety of competing theories of concepts and categorization developed by psychologists in the last thirty years, the implicit theoretical assumption of descriptionism has never seriously been challenged. I present a nondescriptionist theory of the nature of the most basic concepts that we possess, concepts of what I will call "substances," following Aristotle.

"Substances" include stuffs (gold, milk), real kinds (cat, chair) and also individuals (Mama, Bill Clinton, The Empire State Building). And yet, I argue, the category "substance" is univocal: stuffs, real kinds and individuals have something important in common. Based on this similarity, I propose, the earliest and most fundamental concepts that we have of these three sorts of things are identical in structure. The extension of "cat," like the extent of "Mama," is a natural unit in nature, a unit to which the concept cat does something like pointing, and to which it can continue to point despite large changes in the properties the thinker represents the unit as having. For example, large changes can occur in the manner in which a child identifies cats and in the selection of things the child is willing to call "cat" without affecting the extension of the child's word "cat." The difficulty, of course, and a main purpose of the essay, is to cash out the metaphor of "pointing" in this context (c.f., Putnam's "indexicality").

As I will describe substance concepts, having these does not need to depend on knowing words. On the other hand, language interacts with substance concepts in vigorous ways so as to completely transform the conceptual repertoire. I will discuss how the public language plays a crucial role both in the acquisition of substance concepts and in their completed structure.

A Common Structure for Concepts of Individuals, Stuffs, and Real Kinds:
More Mama, More Milk, and More Mouse
Ruth Garrett Millikan

I. Introduction

Frank Keil observes mildly, "...it is difficult to design and motivate empirical studies on concept acquisition without first committing oneself to a set of assumptions about what concepts are and how they are represented" (Keil 1989, p. 25). Indeed so! Concepts, taken as items that the psyche "acquires," are highly theoretical entities. Clearly it is not possible to study them empirically without committing oneself to substantial preliminary assumptions about their nature.

One aim of this paper will be to show how, throughout the changing variety of competing theories of concepts and categorization developed by psychologists in the last half century, the theoretical assumption of descriptionism has managed to go 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 core psychological literature, and despite a number of brave attempts to integrate these insights into the psychological tradition (Lakoff 1987, Markman 1989, Keil 1989, Neisser 1987 chapter 2, Gelman & Coley 1991, Komatsu 1992). The difficulty is that these insights were almost entirely negative. Moreover the tentative positive views offered concerned not the nature of concepts (something in the mind) but rather the extensions of words in a public language. Putnam and Kripke left obscure the nature of the psychological states or processes that would constitute an understanding of the meanings of the words they discussed, thus offering no aid to psychologists. I will try to help remedy that situation.

I will present a nondescriptionist theory of the nature of concepts of what (following Aristotle) I will call "substances." The category of substances includes (1) things we would ordinarily call "substances," namely, stuffs such as gold, and milk and mud, along with (2) things designated "primary substances" by Aristotle, namely, individuals such as Bill Clinton, Mama and the Empire State Building, along with (3) things designated "secondary substances" by Aristotle, namely, real (as opposed to nominal) kinds. Real kinds include, paradigmatically, both "natural kinds" and the correspondents of what Eleanor Rosch called "basic level" categories (Rosch 1975)-- those intermediate level categories such as shoe and mouse and house that children in all cultures learn first (Angelin 1977, Mervis and Crisafi 1982, Nelson 1974). My claim will be that these apparently quite different types of concepts have an identical root structure, and that this is possible because the various kinds of "substances" I have listed have an identical ontological structure when considered at a suitably abstract level. That is, surprisingly to us moderns, the Aristotelian term "substance" is univocal. Unlike the Aristotelian tradition, in modern times, concepts of stuffs and real kinds have traditionally been treated as predicate concepts. That is, to call a thing "gold" or "mouse" has been taken to involve saying or thinking that it bears a certain description.

One understands it as being gold or a mouse by representing it as having a certain set or appropriate sampling of properties or certain relations to other things, or a certain kind of inner nature or structure, or a certain origin or cause. I will argue, on the contrary, that the earliest and most basic concepts that we have of gold and mouse and so forth are subject concepts. Their structure is exactly the same as for concepts of individuals like Mama and Bill Clinton.

To call a person "Mama" is not to attribute to her any properties, relations, or inner or outer causes. It is not to classify her but to identify her. Similarly, Putnam argued, to call a thing "gold" or "mouse" is not to describe it. Neither concept consists of a representation of properties. Rather, the extensions of "gold" and "mouse," like the extent of "Mama," are natural units in nature, units to which the concepts gold and mouse 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. For example, large changes can occur in the manner in which a child identifies gold, hence in the things the child is willing to call "gold," without affecting the extension of the child's word "gold." The difficulty, of course, is to cash out the metaphor of "pointing" (Putnam said "indexicality"). Speaking literally, what is the structure of a substance concept on this view?

As I will describe substance concepts, having these need not depend on knowing words. Preverbal humans, indeed, any animal that collects practical knowledge over time of how relate to specific stuffs, individuals and real kinds, must performe wield concepts of these. On the other hand, language interacts with substance concepts in vigorous ways so as to completely transform the conceptual repertoire. Putnam argued for what he called "the division of linguistic labor." Though giving rather different reasons, I will argue similarly, that the public language plays a crucial role both in the acquisition of substance concepts and in their completed structure.

I will begin with a positive statement of what I take substances and substance concepts to be (II and III). From this nondescriptionist vantage it will be easier to see just how descriptionism continues to be ingredient in contemporary experimental work on concepts (IV). Then I will discuss the nature of concept development from a nondescriptionist perspective (V) and finally the crucial involvement of language in the acquisition and use of substance concepts (VI).

II. Substances

The bulk of a child's earliest words are concrete nouns, including names of individuals, names of basic-level kinds, and some names for stuffs ("milk," "juice"). These are acquired in a rush by the dozens between about one and a half and two years: "this vocabulary spurt is often called the naming explosion to reflect the large preponderance of nouns that are learned" (Markman 1991, p. 81; see Gentner 1982

and Ingram 1989 for reviews, Dromi 1987 for some reservations.)¹ Adjectives come later and more slowly, and abstract nouns later still. This suggests that the ability to distinguish concrete individuals in thought and the ability to distinguish basic kinds and stuffs may have something in common, and that concepts of properties and of other abstract objects may not be required for these tasks. There is much independent evidence that children come to appreciate separable dimensions, such as color, shape, and size, only after a considerable period in which "holistic similarities" dominate their attention (see Keil 1989 for discussion). Thus concepts of properties again appear as less fundamental than those expressed with simple concrete nouns. I propose that individuals, basic-level kinds, and stuffs have something in common that makes them all knowable in a similar way, and prior to properties.

We can begin with kinds. In recent years, a number of researchers have been interested in the structure of concepts of "natural kinds" and in the development of children's understanding of these kinds (e.g., Carey 1985; Keil 1989, Markman 1989, Gelman and Coley 1991). Natural kinds are said to be distinguished in part by the fact that many true generalizations can be made about them. Concepts of natural kinds thus provide an indispensable key to the acquisition of inductive knowledge. According to Gelman and Coley (1991), people develop natural kind concepts

...with the implicit...goal of learning as much as possible about the objects being classified.... For example, if we learn that X is a "cat," we infer that it has many important properties in common with other cats, including diet, body temperature, genetic structure, and internal organs. We can even induce previously unknown properties. For example, if we discover that one cat has a substance called "cytosine" inside, we may then decide that other cats also contain this substance...(p. 151)

Gelman and Coley (1991) call this feature "rich inductive potential." They, and especially Keil and Markman, are explicit, however, that "natural kinds" are not sharply set apart from artifactual or even, in Markman's view, from purely nominal kinds. "Bird" and "white thing," Markman tells us

...should be viewed as endpoints on a continuum from natural kind categories, which have rich correlated structure and are embedded in scientific theories, to arbitrary categories, which have impoverished correlated structure. Many other categories fall somewhere between

¹ There is evidence that Korean children may usually have a "verb spurt" a month or two before their "noun spurt" begins. Still the number of nouns soon overtakes the number of verbs (Choi and Gopnik 1993).

these two extremes. "Chair," for example, is an intermediate type of category. Once we know an object is a chair, we know a fair amount about its physical appearance, construction and typical function. (1989, p. 114)

If there is indeed such a continuum, basic-level categories would seem to be closer to the "natural kind" end. Thus Mervis tells us,

[basic-level] categories are based on large clusters of (subjectively) correlated attributes that overlap very little from category to category. In our world, these basic-level categories are the most general categories whose members share similar overall shapes (or similar parts in particular configurations; Tversky & Hemenway 1984) and similar functions and characteristic actions. (Mervis 1987, p. 202, citing Rosch et al 1976)

Many basic level kinds do not, however, figure in scientific theories, because they do not figure in universal laws. They are of special interest because they afford so many inductive inferences, not because they afford totally reliable ones. In this way they differ from "natural kinds" in the strong sense employed by some philosophers (e.g., Putnam 1975). I will lump these two varieties of kinds together, speaking only of "real kinds" as opposed to "nominal kinds." There are two continua from richer to poorer among real kinds, reflecting (1) the multiplicity of inferences supported, and (2) their reliability.

What I want to do, now, is to generalize the notion "rich inductive potential," showing how it applies not just to real kinds but also to stuffs and individuals.

Classically, induction is described as a movement from knowledge about certain instances of a kind to conclusions about other instances of the same kind. Consider, now, generalizations made over instances of the second order kind meetings with kind K, for example, over the kind meetings with mice. Compare this with making generalizations over meetings with the stuff milk, and meetings with the individual Mama. These are equally easy and equally productive ways to generalize.

The ontological category "substances," as I use this term, is roughly (more precision later) that widely extensive category consisting of items about which it is possible to learn from one encounter something about what to expect on other encounters.² Thus I can discover on one encounter (temporal or spatial) that cats eat fish and the knowledge will remain good on other encounters with cats. Or I can discover that Xavier knows Greek and this will remain good on other encounters with Xavier. Or I can discover that ice is slippery and this will remain good when I encounter

² An in depth discussion of the ontological category of substances is in (Millikan 1984), chapters 16 and 17.

ice again in that other puddle over there or next winter. For cats, I can also discover numerous other anatomical, physiological, and behavioral facts that will carry over. There is the entire subject of cat physiology and behavior studied by those attending veterinary schools. Even more carries over for Xavier, since he carries over all or most of what can be discovered about humans, as well as having many of his own stable properties. And for any determinate kind of stuff, there is a vast array of questions, such as "what is its chemistry?", "what is its melting point?", "what is its specific gravity?", "what is its tensile strength?" that can sensibly be asked about it and answered, once and for all, on the basis of one careful observation. Roughly for these reasons, cat(kind), Xavier, and ice are each "substances." Besides stuffs, real kinds, and individuals, the category substances may include such things as certain event types³, cultural artifacts, musical compositions, and so forth, but I will ignore these latter in the present essay.

It is not a matter of logic, of course, but rather of the make up of the world, that I can learn from one observation what color Xavier's eyes are or, say, how the water spider propels itself. It is not a matter of logic that these things will not vary from meeting to meeting. And indeed, the discovery on one meeting that cat is black does not carry over; next time cat may be striped or white. Nor does the discovery that Xavier is talking or asleep carry over; next time he may be quiet or awake. Nor does discovering that ice is cubical or thin carry over. Although substances are, as such, items about which enduring knowledge can be acquired from one or a few encounters, for each substance or broad category of substances, only certain types of knowledge are available. Moreover, most of the knowledge that carries over about ordinary substances is not certain knowledge, but merely probable knowledge. Some cats don't like fish, perhaps, and a stroke could erase Xavier's Greek. But compare: no knowledge whatever carries over about nonsubstance kinds, such as the red square or the two-inch malleable object, or the opaque liquid, except what applies to one or another of the analytical parts of these complexes taken separately. (Similarly for Markman's "white thing" above. It is not on the scale with substances, for there is nothing to be learned about it.)

There are various contemporary interpretations of the underlying reasons why there are such things as real kinds in nature, including, especially, more than one thesis on the nature of "natural kinds" (Putnam 1972, Boyd 1989, 1991; Hacking 1991; Kornblith 1993). Everyone agrees, however, that what makes something a natural kind

³ "There appears to be a basic or generic level of categorization for events, again just as for object categories (see Abbot et al; John 1985; Rifkin 1985; Rosch 1978; Tversky & Hemenway 1984)." (Clark 1991, p. 39).

is that there is some such reason: kinds are not natural if they yield inductive knowledge by accident. Similarly, I suggest, for real kinds generally. If a term is to have genuine "rich inductive potential" it had better attach not just to a pattern of correlated properties, but to an univocal explanatory ground of correlation.

My own position is that there are numerous different types of reasons for the occurrence in nature of real kinds, these accounting for success in generalizing over encounters in different ways (Millikan 1984, chapter 16). Sometimes there is a single underlying cause or inner structure (c.f., Putnam's "natural kinds") that results, always, or under common conditions, in a certain selection of surface properties, as is the case with the various chemical substances. In such cases, the kinds have real, not merely nominal, essences, discoverable by empirical investigation. Sometimes, rather than having a single unifying essence, the properties of a real kind may cluster because of a sort of homeostasis among them or among their causes (Boyd 1991). Then there is no essence at all. Nothing in common to all members of the kind. Nor is there an essence in the other cases I shall now mention.

Sometimes the unifying cause of a real kind may be largely external, as in the case of many artifact categories. Keil tells us,

Chairs have a number of properties, features, and functions that are normally used to identify them, and although there may not be internal causal homeostatic mechanisms of chairs that lead them to have these properties, there may well be external mechanisms having to do with the form and functions of the human body and with typical social and cultural activities of humans. For example, certain dimensions of chairs are determined by the normal length of human limbs and torsos...the causal homeostatic mechanisms for natural kinds are closely related to various domains of science, such as biology, chemistry, and physics, whereas those for artifacts and natural kinds involve more social and psychological domains of causality. (Keil 1989, p. 46-7)

Another very common explanatory ground determining similarities among members of a real kind is copying or reproduction. For example, a factor often accounting for limited variety within artifact categories is that the same design is copied over and over. Similarly, the animals or plants in a species are alike, not only because of homeostasis in the gene pool, but because they (their genes) are reproduced from one another. Another variety of real kinds are (fully or partially) socially constructed kinds, for example, school teacher, doctor, and father. People falling in these categories act similarly as a result of similar training handed down from person to person (reproduction), custom (more reproduction), social pressures to conform (reproduction again) or law. Sometimes members form a social kind "only because people class them together," but the "because" here may be causal not logical, hence the kind may be real not nominal.

Turning now to what holds individuals together over encounters over time, Xavier today is much like Xavier yesterday because Xavier today directly resulted from Xavier yesterday, in accordance with certain kinds of conservation laws and certain patterns of homeostasis. Similarly, Ghiselin and Hull have claimed that a species is really just a big scattered individual, causing itself to continue over time much as an individual does (Ghiselin 1974, 1981; Hull 1978). A dog is a member of the species dog because it was born of a dog, not because it is like other dogs. Conversely, some philosophers have thought of Xavier as a class consisting of Xavier time-slices each of which causes the next. Either way there is a deep similarity between individuals and many real kinds, and either way neither individuals nor real kinds need have essences.

Philosophers interested in such questions have thought up numerous bizarre examples where it would not be clear whether to say that this individual thing was numerically the same as that individual thing occurring later in time. However, in the usual case we take it, quite rightly, that whether a correct identification of an individual has been made is a matter of how the world is, not of how we humans (or we English speakers) like to identify. Similarly, the question whether a seemingly marginal item is or is not of a certain real kind is most often a straightforward substantive question about how the world is, not a question of how we humans (or we English speakers) like to classify. If it is not like other members of the kind for the very same reason they are like one another, it is not a member of the kind. On the other hand, because the occurrence of causative factors accounting for similarities among members of a group can be more or less irregular, and because numbers of grounded similarities can be larger or smaller, whether a real kind exists at all is sometimes a marginal matter.

In sum, a "substance" is something that one can learn things about from one encounter to apply on other occasions where this possibility is not coincidental but grounded. There is an explanation or cause of the samenesses.

I wish to make plausible now that despite the many different kinds of groundings that account for the unity of various types of substances, the basic structure of a concept that grasps a substance is always the same. This is possible because there is no need to understand what the ground of a substance is, or even that a substance has a ground, in order to have a concept of that substance. Throughout the history of philosophy and psychology, the tendency has been to project the structure of the object grasped by thought into the mind itself. On the contrary, I will argue, substances are grasped not by understanding the principles that hold them together, but by knowing how to exploit them for information gathering purposes.

III. Concepts of Substances

The "concept" of a substance, as I use that term⁴, is the capacity to represent it

⁴ I do not recommend generalizing this description of a concept

in thought for purposes of information gathering and storage, inference, and ultimately guidance of action. We wish to know the structure of this ability. To describe the structure of an ability is to tell what it is an ability to do, what sub-abilities are contained in it and, ultimately, by what means it is exercised--exactly how it accomplishes what it does. Using largely a priori means we cannot hope to travel very far, but Frank Keil was surely right that in order to engage in empirical research, one must have some idea what one is looking for. Experimental results are worthless without an approximation, at least, to a sound theoretical framework in which to interpret them.

From the standpoint of an organism that wishes to learn, the most useful and accessible subjects of knowledge are things that retain many of their properties, hence potentials for use, over numerous encounters with them. This makes it possible for the organism to store away knowledge about the thing collected on earlier encounters for use on later occasions, the knowledge retaining its validity over time. Substances are (by definition) what afford this sort of opportunity to a learner. In the experience of a child, for example, Mama retains many of her properties over various encounters with her just as milk and mouse do. Given this, we might expect the child, indeed we might expect any animal, to learn how to relate to, and what to expect from, these various items in much the same way. For example, ontologically speaking, individuals are space-time worms while real kinds are collections of similar space-time worms, but to have the capacity to understand this ontological distinction would require a grasp of space-time structure and temporal relations of a sort not acquired by children until years after they are proficient in the use of both proper and common names (Nelson 1991). Putting it Quine's way, the child's first recognitions (and those of the dog) must be merely of more Mama, more milk, and more mouse (Quine 1960 P. 92). The child observes things about Mama when she encounters her, not about samples or instances of Mama. Similarly, to learn things about milk, she need not understand what it would be to think of or keep track of portions of milk as individuals. And the very point of having the concept mouse would seem to be that under it, one does not distinguish Amos from Amos's brother, but conceives them both as the same. Note that I am talking here about applying substance concepts, not about acquiring them. My claim here is only that early substance concepts, even when what they are of, ontologically, is kinds, need not be predicate concepts applied to prior subject concepts. They need not be understood as descriptions of anything.

The various substances differ, of course, in the types of knowledge they afford. The child's individual Highchair retains its overall shape hence its sitting-on capacity over encounters but Mama does not (you cannot sit on Mama when she is standing).

indiscriminately, for example, to "mathematical concepts", "logical concepts", "modal concepts" and so forth. The idea that every word corresponds to a concept in some univocal sense of the term "concept" is surely mistaken.

Milk and Mama retain their color while cat does not. But these primitive subjects of knowledge are grouped into rough ontological categories. Even for the very young child, surely, a casual look at a new piece of furniture on the one hand and a new uncle on the other, easily reveals which can be counted on to retain its current climbing-up-on affordance and which may grow tired of the sport. Similarly, preschoolers know that what is sleepy might also be hungry, but not made of metal or in need of fixing (Keil 1983). An important question for psychologists, of course, is when and why and how these basic ontological category distinctions are grasped by the developing child.

Now think why a child, or animal, needs to carry knowledge of the properties of a substance from one encounter to another. If all of a substance's properties were immediately manifest to the child upon every encounter there would be no need to learn and remember what these properties were. Carrying knowledge of substances about is useful only because most of a substance's properties are not manifest but hidden from us most of the time. This is not, in general, because they are "deep" or "theoretical" properties, but because observing a property always requires a particular perspective on it. To observe that the sugar is sweet it must be in your mouth, to observe that the milk is drinkable and filling you must tip it and drink. You do not find out that the cat scratches until you disturb it, or that the fire burns unless you near it, and the pretty design on the front of the quilt is not seen from the back. Different properties and utilities of a substance show themselves on different encounters. That is WHY it is useful to collect knowledge of a substance over time.

Yet there is a sort of paradox here. For it won't help to lug knowledge of a substance about with you unless you can recognize that substance when you encounter it again as the one you have knowledge about. If different properties of a substance show themselves at different encounters, how is one to know when one has encountered the same substance again? The very reason you needed to carry knowledge about in the first place shows up as a barrier to applying it. Moreover, not only substances but their properties reveal themselves differently at different encounters. The enduring properties of substances are distal not proximal and they affect the external senses quite differently under different conditions and when bearing different relations to the perceiver.

Clearly, then, a most complex but central skill required for any organism that employs knowledge of substances will be the ability to reidentify these substances with fair reliability under a wide variety of conditions. This will be necessary, first, in order to develop practical skills in the use of various substances. It will be necessary, also, for any animal that uses representations of facts about substances as a basis for practical and theoretical inference. For example, suppose that I am hungry and I know that yogurt is good to eat and that there is yogurt in the refrigerator. This is of no use unless I also grasp that these two bits of knowledge are about the same stuff, yogurt. To caricature⁵, if I represent yogurt to myself in one way, with a mental heart, as I store

⁵ To model the act of reidentifying a substance in thought as using the

away the knowledge that yogurt is good to eat, but represent it another way, with a mental diamond, as I store away that it is in the refrigerator, these bits of information will not help me when I am hungry. A fundamental sub-capacity involved in having a concept of any substance must be the capacity to store away information gathered about it such that it is always represented again with what one is disposed to employ as⁵ a representation with the same semantic value. This capacity is the capacity to maintain a coherent inner representational system, which means that it is essential for representing something in thought at all!

The ideal capacity to identify a substance would allow reidentification under every physically possible condition, regardless of intervening media and the relation of the substance to the perceiver. The ideal capacity would also be infallible. Obviously there are no such capacities. If the cost of never making an error in identifying Mama or milk or mice is almost never managing to identify any of them at all, it will pay to be less cautious. If one is to recognize a substance when one encounters it a reasonable proportion of the time, one needs to become sensitive to a variety of relatively reliable indicators of the substance—indeed, to as many as possible, so as to recognize the substance under as many conditions as possible. Counted as indicators here would be, in the first instance, the various appearances of the substance to each of the various senses, under varying conditions, at varying distances, given varying intervening media, or resulting from various kinds of probing and testing. In the second instance would be pieces of information about the presented substance that it has these or those properties that indicate it reliably enough.

In the case of familiar substances, typically we collect over time very numerous means of identification, all of which are fallible, certainly none of which are "definitional," of the substances being identified. The purpose of a substance concept is not to sustain what Wettstein (1988) aptly calls "a cognitive fix" on the substance, but the practical one of facilitating information gathering and use for an organism navigating in a changing and cluttered environment. Consider, for example, how many ways you can recognize each of the various members of your immediate family by looks of various body parts from each of dozens of angles, by characteristic postures, by voice, by footsteps, by handwriting, by various characteristic activities, by clothes and other possessions. None of these ways nor any subset defines for you any family member, and probably all are fallible. There are, for example, conditions under which you would fail to identify even your spouse, conditions under which you would misidentify him and

same mental term again, as I do here, is actually a crude and misleading expedient. On this, see (Millikan 1991, 1993b and 1994).

5

conditions under which you might mistake another for him. The same is true of your ability to identify squirrels or wood. To be skilled in identifying a substance no more implies that one never misidentifies it than skill in walking implies that one never trips. Nor does it imply that one has in reserve some infallible defining method of identification, some ultimate method of verification, that determines the extension of each of ones (idiolect) substance terms. No more than the ability to walk implies knowing some special way to walk that could never let one trip.

It follows that it cannot be ones dispositions to apply a substance term that determines what its extension is. In a passage characteristic of the literature, Lakoff remarks, "It is known, for example, that two-year-olds have different categories than adults. Lions and tigers as well as cats are commonly called "kitty" by two-year-olds...." (1987, p.50). And how does Lakoff know that two-year-olds are not thinking of lions and tigers that they are kitties. Kitties grown big, why not? A little more experience and the child may change her mind. Not on the question what "cat" means, but on reliable ways to recognize cats. At age three, my mother stoutly insisted that her father was "Uncle Albert" when he came home one night without his beard. Surely it does not follow that "Uncle Albert," for her, referred also to her father? A child who has got only part way toward knowing how to ride a bicycle has not learned something other than bicycle riding, but partially learned how to ride a bicycle. It is the same for a child who has got only part way towards recognizing cats, or father, or Uncle Albert.

The practical ability to reidentify a substance when one encounters it, so as to collect information about it over time, and so as to know when it is possible to apply that information, has to be complemented, however, with another equally important ability. Having a concept of a substance also requires that one have some grasp of what kinds of things can be learned about it. For example, one must have some ability to tell which kinds of practical successes can be expected to carry over to new encounters with the substance. If the concept is to be used for gathering theoretical knowledge, one must know something of the range of predicates, i.e., determinables, that are applicable to the substance. That is, one must understand what some of the meaningful questions are to ask about it.⁶ You can ask how tall Mama is, but not how tall gold is. You can ask at what temperature gold melts, but not at what temperature chairs (as such) do. The latter is a question that can be answered only for (some) individual chairs. There is much that you can find out about the internal organs of each species of animal but not about the (visible) internal parts of gold or mud. Having a concept of a substance is not knowing an essence, but it must involve understanding something of what recognition of the substance might be good for in the context either of developing practical skills or theoretical knowledge.

IV Contrast with Descriptionism

⁶ See (Millikan 1984) chapter 15, p.252 ff., and chapters 16 and 17.

In contrast to the position just sketched, 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 frequently found caricatured in the statement that concepts are features or properties: "many properties are concepts themselves" (Barsalou 1987, p. 129).

Using the concept chair as his example, Komatsu (1992) describes 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..." (1992:500, italics mine). Building on (Medin & Smith 1981, 1984), he applies this 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 human perceptual and cognitive systems, 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 Lakoff subtitles his 1987 book, "What Categories Reveal about the Mind."

But descriptionism is not always allied with nominalism or conventionalism. It

has also 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 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 is not determined by fitting a certain description associated with a word. But in fact there are a number of ways in which these 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.⁷ 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 misinterpret Kripke's (1972) and Putnam's (1975) antidescriptionist views on the meaning of proper names and natural kind terms as invoking definite descriptions at one level removed. (No, Kripke did not claim that the referent of a proper name N is fixed in the user's mind by the description "whoever was originally baptized as N" nor did Putnam claim 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" (but see also Fumerton 1989).)

The theory that language categories are organized "probabilistically" (Medin 1989) by family resemblance or by reference to prototypes may 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 probabilistic theories are regularly interpreted as explaining how the learner's experience generates the category, the actual extension of the category being determined not by the real extent of the 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, italics mine) and Ward and Becker state that "category structure" can mean "the set of items that the learner considers to be

⁷ 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.

members of the category in question (i.e., the category extension)" (1992, p.454). It is assumed, that is, that although experience with a natural kind may inspire the category, 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 called an "explanation-based view" of concept structure. Komatsu characterizes this view by quoting Keil (1989, p. 1):

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.

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. 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. But these advocates of explanation-based views have remained strongly influenced by the characteristic mid twentieth century doctrine that the "meaning" of a term or concept is a matter of its connections with other terms or concepts, so that introducing or changing theories threatens to change meanings:

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, p. 15-16). Again, the view here is descriptionist. There is no suggestion here 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, see Gopnik&Meltzoff 1996.)

In the alternative to descriptionism that I am suggesting, having a concept of a substance is not having a defining description of it or a theory about it. To have a theory about a substance you have to be able to think of it, and it is this capacity that is the concept. To think of it one must be able to represent it in a stable representational system, where what is in fact the same substance again is represented as being the same again. To maintain such a representational system requires that one have the capacity to recognize the substance under varying conditions so as to know what incoming information to store as information about the same. Thus the core of a substance concept is a (necessarily fallible) capacity to recognize what is objectively the same substance again as the same despite wide variation in the faces it shows to the senses. The extension of one's concept is then determined, not by ones fallible dispositions to recognize portions of its extent, but by the real extent of the substance that has governed the development of these dispositions.

The standard descriptionist view takes the substance concept to be an ability to classify instances of the substance. Forcing the distinction, perhaps, between these two for expository purposes, the difference between identifying and classifying lies both in purpose and in psychological structure. The purpose of a classification system is nicely captured by the following contemporary descriptions of "categorization," and of "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 information storage and transfer: the efficient organizing of what we already know (encyclopedias), putting things away where we can find them again (libraries, grocery shelves), communication (briefly telling enough about the object for someone else to identify it.) The initial data for a paradigm classification task include a specification of all the properties of the object to be classified that are relevant to its classification. A librarian would not try to classify a book, for example, without carefully examining its contents. Similarly, in classical categorization experiments, all relevant properties of each "stimulus" and each "test item" are clearly exhibited to the learner.

Reidentifying is required, on the contrary, not for information storage and transfer, but for its acquisition and use. One needs to be able to identify a substance under diverse circumstances in order to come to know its properties, properties that happen not to be currently manifest. This one does by managing to recognize the substance on the basis of whatever properties do happen to be currently manifest, then applying one's prior knowledge of others of its properties to the current encounter. Only in this way can prior knowledge of the substance find a use.

The psychological structure of classification is the structure of subject-predicate judgment. To classify an item requires differentiating the item to be classified in thought and applying a predicate to it. For example, classifying animals as dogs, cats, or mice involves thoughts of Fidos and Spots, Amoses and brothers of Amoses, each individual to be judged a member of its proper category. But when the child recognizes Mama, "Mama" is not a predicate term: surely the child is not categorizing instances of Mama. Nor need the child conceive of mice as individuals in order to recognize the substance mouse again.

V The Development of Substance Concepts

Viewing a substance concept as an ability to reidentify which a motile person comes to exercise within a supporting but changing environment, the study of concept development is also seen in a new light. What sub-skills are involved in this ability?

What is the characteristic progression toward acquisition of these skills? The answers here are mainly for psychologists to find, but I can try to make the questions clearer.

According to various estimates, children acquire from five to nine words daily between the ages of two and six (Waxman 1991, Clark 1991, Byrnes and Gelman 1991) Chomsky says, "about a word an hour from ages two to eight with lexical items typically acquired on a single exposure..." (Chomsky 1995, p. 15). How is this possible? An obvious hypothesis here is that many concepts are developed prior to language, and indeed, at least some must be, for the infant recognizes her mother and the dog recognizes his master. Each has the capacity to reidentify the relevant individual under diverse conditions, thus making it possible to learn how to behave appropriately in their presences.

Some of the skills needed to accomplish the task of reidentifying ordinary substances have traditionally been classified as "motor" and "perceptual" rather than "cognitive". Perhaps the most fundamental of these is the ability to track objects with the eyes, head, feet, hands, ears, and nose, etc.. Objects tracked in this way are not merely conceived to be the same but are perceived as the same under certain conditions, the perception of sameness bridging, for example, over motions of perceived and perceiver, over changes in properties of the object, and over temporary disappearances of the object behind other objects. The mechanisms responsible for the ability to track and for perceptual "identity-" or "existence-constancy" may well be largely endogenous (Nelson&Horowitz 1987, Dodwell et al 1987, Spelke 1993) and certainly are "cognitively impenetrable" (Shepard 1976, 1983). These basic abilities are surely the bottom layer on which conceptions of substances are built.

Tracking allows the accumulation of information about a substance over a period of time, information perceived as about the same substance. Nor is it only individual objects that are tracked in this way. If I am tracking Fido, I am also tracking the species dog, and also fur and bone. Which of these I am tracking with my mind depends upon which I am learning about or registering information about as I go. And that is determined by which substance I identify on other occasions as the one this learning concerns—as being the same substance again. As I dissect my specimen frog in the zoology laboratory, whether I am conceptually tracking the individual, Kermit, or tracking just frogs, depends on whether I attempt to apply what I have learned from my experience only to later meetings with Kermit or whether to frogs in general.

For the usefulness of ones knowledge of a substance to last, however, one must also know how to reidentify the substance after a lengthy break, say, next day or next week. Call this "conceptual tracking": one understands rather than perceives that the substance is the same one again. Out of what materials are our abilities conceptually to track substances built?

By tracking a substance one can learn many different ways to recognize it Show it

looks, how it sounds, how it feels, the manner in which it moves and changes. The mechanisms of perceptual constancy for properties can then be brought into play. These mechanisms may be fashioned in part, and certainly are tuned, through experience, but much of their structure also may be endogenous (Dodwell et al 1987, cf., Gallistel et al 1993, Marler 1993). They cause distal qualities to appear as the same through wide variation in proximal manifestations. For example, they allow the same shape and size to be registered as the same despite alterations in angle of observation and distance, colors to appear as the same under widely varying lighting conditions, and voices to sound as the same voice through distortions and superimposed noise.

Involvement of the mechanisms of perceptual constancy should not be thought to imply, however, that actual concepts of properties are always involved in conceptual tracking of substances. Not if having concepts of properties means being able to represent properties, as such, in thought. For example, being caused to token mental squirrel again when prompted by the same distal configuration of shape, color, texture and motion is not, as such, to token any thoughts of the shapes colors or textures themselves. The thought of a property is not just a reaction caused by a property; it must play an appropriate representational role. This accords, of course, with the finding that children appreciate holistic similarities before appreciating separate property dimensions such as color and shape.

When perceptual tracking is coupled with exploratory manipulation, probing and testing, this may reveal properties and dispositions which prove to be better tracers, better aids to achieving conceptual constancy. An easy example is the tool bag of tests and routines that chemists use in order to reidentify chemical stuffs. In the end, indeed, any knowledge at all that one has of a substance can help to identify it, if not positively, then negatively. No, we think, that can't be Sally after all because Sally doesn't know French, or that can't be real gold in the window because real gold would cost more than that. It is because knowledge of the properties of substances is often used in the process of identifying them that it is easy to confuse having a concept of a substance with having knowledge of properties that would identify it.

But how does the child know which aspects of the substances it is learning to track can be relied on for reidentification? And how does she know what questions she should expect to be answerable about each substance? Just as the child has built-in perceptual tracking abilities and built in perceptual constancies, we might expect her to have certain built-in conceptual tracking abilities.

There is evidence, for example, that infants may have built into them systems designed, specifically, to recognize human faces. And it is well known that they have a strong disposition from the earliest days to track and study human faces (e.g., Johnson et al 1991). Also, many species that recognize conspecifics as individuals instinctively use smell for this purpose, and human infants also know Mama by smell in the early

months (MacFarlane 1977).⁸ It appears that the infant may know innately at least two good ways conceptually to track individual conspecifics. Faces and personal odors are indicative of individual identity; clothes, postures, and so forth, are not.

The mechanisms by which infants reidentify individuals perceptually apparently do not rely upon properties of the tracked object remaining the same but upon common movement, spatial location and trajectory (Gopnik&Meltzoff 1996). Xu and Carey (1996) have recently produced experimental evidence that 10 month infants, unlike 12 month infants, are not surprised if an object of one kind apparently turns into an object of another kind, say, a yellow rubber duck into a white Styrofoam ball, though they are surprised if an object they are tracking apparently turns into two objects. Tracking in this property-blind way would make it possible to observe, for various broad kinds of objects, what sorts of things tend to remain the same and what sorts may change within a short period, yielding clues for later conceptual tracking.

Whether we have built in ways of conceptually tracking stuffs or real kinds of any particular sort, such as physical kinds, animal kinds, plant kinds, artifacts, social kinds and so forth, is clearly a matter for empirical research--research of the sort that Spelke, Carey, Keil, Gelman, Markman and others have recently been doing, though I am suggesting a somewhat different framework for interpretation of experimental results. Without doubt, the results of more traditional studies of concept formation may also cast light on how conceptual tracking develops. Examining "the function" from "learning instances plus the target items to categorize" to "the set of possible category judgments," as Billman put it (1992) should help us to discern what kinds of traces are followed in attempting conceptual tracking, at various ages, and for different domains of real kinds. That we should be acutely sensitive to correlations among properties, likely among specific kinds of properties in specific domains (cf., Gallistel et al 1993, Marler 1993, Atran 1989; Keil 1979, 1989; Carey 1985; Markman 1989; Gelman and Coley 1991, Spelke 1989, 1993) seems an obvious way for us to attempt to track many kinds of substances. But experiments need to be designed and interpreted with it in mind that the cognitive systems are designed by evolution and tuned by experience to find real world substances, not random logically possible ones. Close attention should be paid to the details of real world ontology, to the principles that hold real substances together, and the relevance of experiments using artificial objects and kinds should be carefully justified.

The most accurate and sophisticated ways of tracking substances conceptually emerge only as insight is slowly gained into the ontological principles that ground them. The psychologists Medin, Gelman, Keil, and Gopnik&Meltzoff (1996), especially, have been interested in tracing the origin and development of children's understanding of these principles. I much admire this research. My suggestion is only that we should be

⁸ Thank you, Dan Dennett

clear that understanding of this sort is not necessary to having a concept of a substance, and that having or lacking such understanding need make no difference to the extensions of ones substance concepts.

A substance concept causally originates from the substance that it denotes. It 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 conceptually, hence physically, tracking and picking up information about, and because the concept has been tuned to its present accuracy by causal interaction with members of A's specific domain and/or with A itself, during the evolutionary history of the species and/or through the learning history of the individual. If it is not definite which among various 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, simply, equivocal. For example, to have two people "mixed up" or "confused" in one's mind is to have an equivocal substance concept (Millikan 1984 chapter 15, 1991, 1993a chapter 14, 1994).

One more extremely fundamental medium through which conceptual tracking is achieved is language...

VI Substance Concepts and Language

The story I have been telling about substance concepts apparently runs headlong into the blatant fact that many of these concepts, both for children and adults, have been acquired without encountering the substances "themselves" but only by "hearing of them." With regard to these same substances, however, we are often in the position that Kripke (1972) and Putnam (1975) observed, knowing neither how to identify these substances in the flesh, nor by any unique or defining descriptions. That is, neither verificationist nor descriptionist theories of concept extension explain these cases either. This entire problem falls away, however, if we view speech as a direct medium for the perception of objects in the same way that, say, light is.

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

There are two things that distinguish direct perception quite sharply from the acquisition of information through language, but neither implies a difference in

immediacy. In direct perception, the spatial and temporal relation of the perceiver to the object perceived is given, whereas it is not normally given through language. On the other hand, when watching television, the spatial relation of perceiver to perceived is not given either, nor, unless the program is live, is the temporal relation, yet one perceives that the newscaster frowns or smiles just as immediately as one would in his presence. The second feature that distinguishes perception is its near infallibility. For the most part, it takes a modern understanding of the mechanisms of perception and a substantial technology to manage materially to fool the human eye or ear. False appearances are easily arranged, however, using modern communications media, offering the most common (though generally overlooked) illustration of the persistence of perceptual illusion. Similarly, through language, persistent illusions are very easy to arrange, hence abundant. That is, sentences are often false, and even when you know they are false, they continue to present the same false appearances. They do not shift and appear to say something different. In sum, hearing sentences may be quite a lot like watching the media, which in turn is quite a lot like watching the original.

Think of the matter this way. There are many ways to recognize, for example, rain. There is a way that rain feels when it falls on you, and a way that it looks out the window. There is a way that it sounds falling on the rooftop, "retetetetetet" and a way that it sounds falling on the ground, "shshshshsh". And falling on English speakers, here is another way it can sound: "Hey, guys, it's raining!"⁹ Nor should you object that it is not rain you hear in the last case but rather "a sentence". A sound? Is it then a sound that you hear rather than rain on the roof? Is it a TV screen that you see rather than Dan Rather? A pattern of ambient light rather than the TV screen? Best of all, perhaps all you see is a visual impression. You can, if you like, hear or see any of these things. What you see when you look depends, first, on where you focus your eyes; it depends, second, on where you focus your mind, your attention.

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

⁹ Thank you, Crawford Elder

substance as manifestations of it; one learns how to translate information arriving in one more kind of package at ones sensory surfaces into beliefs.

Learning a language is, in part, learning more ways to pick up information through the senses and put it away in the right boxes. A difference, of course, is that this way of picking up information is much more fallible than in the case of ordinary perception. But no human ability is infallible. Further, just as substances are sometimes look-alikes in the flesh (twin brothers), many substances are sound-alikes in words (John_(Doe) and John_(Roe)). But substances are tracked through the medium of words not merely by means of the same words manifesting the same substances. Like more direct manifestations of substances, words and sentences occur in context, allowing methods of tracking to be used that are analogous to more ordinary tracking in that they rely in large part on expected spatial, temporal and causal relations (c.f., trajectory) rather than persistence of properties. (How do I recognize that as John's elbow poking out over there behind the lamp? I saw John head that way with a book just a moment ago.) Some of these relations are natural, as the natural relation between a speaker's experience and context of speech and his expressed knowledge. One will usually know which "John" a speaker is talking about in this way. Other such relations are conventional, as in the interpretation of certain anaphoric pronouns and certain indexicals.

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

It is even possible, indeed it is common, to have a substance concept entirely through the medium of language. It is possible to have it, that is, while lacking any ability to recognize the substance in the flesh. For most of us, that is how we have a concept of Aristotle, of molybdenum and, say, of African dormice. CThere, I just handed you a concept of African dormice, in case you had none before. Now you can think of them nights if you like, wondering what they are likeCon the assumption, of course, that you gathered from their name what sorts of questions you might reasonably ask about them (animal questions, not vegetable or mineral or social artifact questions). In many cases there is not much more to having a substance concept than having a word. To have a word is to have a handle on tracking a substance via manifestations of it produced in a particular language community. Simply grasping the phonemic structure of a language and the rudiments of how to parse it enables one to help oneself to an embryo concept of every substance named in that language. That, I suppose, is why it is possible for small children to learn a new word every hour. The

basic phenomenon here is the same as that underlying Putnam's "Division of Linguistic Labor" (1975) and Burge's claim that constitution of the very content of ones thought sometimes passes through the word usages of a surrounding language community (1979, 1982, 1986).

Acquiring adequate substance concepts involves learning to focus one's thought, such that all of the incoming information scattered over time about each substance is put into one slot, and associated with the right categories of properties (determinables). Earlier, I suggested that preschoolers who take tigers to be "kitties" may be confused, not about the meaning of the word "kitty," but about how to identify cats. From our present perspective, however, thinking tigers are "kitties," that is, putting tiger information away in the same slot as information gotten from hearing about "kitties," is being confused about tigers as well as about domestic cats.¹⁰ But Gelman and Coley (1991) are surely right that "a word can serve to stake out a new category, which then must be explored in more depth" (p. 184; see also Gopnik & Meltzoff 1993). Words are handles to hang onto, helping to stabilize concepts so as gradually to eliminate equivocation in thought. Granted that those who speak to one have adequate concepts themselves.

But has not a blatant distinction has been overlooked here between merely knowing of a word and knowing what that word means? On the present view there is an interesting question what it is for a child to learn the meaning of a word that names a substance. Traditionally, this is supposed to involve coming to exercise the same concept in connection with the word that adults do. But since a concept is an ability, there is an ambiguity here in the notion "same concept," derived in turn from a natural ambiguity in the notion "same ability." Suppose, for example, that you tie your shoes by looping one lace into a bow, encircling it with the other, and pulling through, while I tie mine by looping each lace separately, then tying them together. The results that we get will be exactly the same, but do we exercise the same ability? Sometimes what counts as the same ability is what accomplishes the same; other times it is what accomplishes the same by the same means.

Similarly, both the organic chemist and the child identify sugar and collect knowledge about it. Does it follow that there is a concept that they both have, hence that they have "the same concept"? In one sense they do, for each has the ability, fallibly, to identify sugar. But in another sense they do not, for the chemist has much more sophisticated and reliable means at her disposal for identifying sugar than the child. Similarly, we could ask, did Helen Keller have many of the same concepts as you and I, or did she have largely different ones, and again the answer would be equivocal. Suppose we say that the child has the "same concept" as the chemist, namely, the

¹⁰ For a full discussion of equivocation in concepts, see (Millikan 1993a chapter 14, 1993b, 1994).

concept of sugar, but that she has a very different "conception" of sugar than does the chemist. Similarly, Helen Keller had very many of the same concepts as you and I but quite different conceptions of their objects. This fits with the ordinary view that people having very different information or beliefs about a thing have "different conceptions" of it, for information one has about a substance is often used to help identify it.

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

On the other hand, for some (how many?) real kinds, it may be that there are core methods by which nearly every adult (the "nearly" is for Helen Keller) knows to reidentify them. Or there may be certain conditions under which any adult would recognize the substance, or examples of the substance that any adult would recognize given a chance to examine them. Then there may be a sense in which the child does not fully understand "the meaning" of the word for that substance until her competence at identifying the substance has been filled out to match adult standards. In this sense of "the meaning," knowing how to track a substance only by tracking its name would not be nearly enough for "knowing the meaning". But is it in this sense that you "know the meaning" of the word "molybdenum," or "brisket," or "African dormouse"? Indeed, do you know what these words mean? Best not to fall into a verbal dispute over what gets to count as "knowing the meaning".

BIBLIOGRAPHY

- Abbot, V., Black, J.B. & Smith, E.E. 1985 "The Representation of Scripts in Memory" Journal of Memory and Language 24:179-99.
- , Angelin, J. M. 1977 "Classifiers" Language 53:285-311
- Atran, S. 1989, "Basic Conceptual Domains" Mind & Language 4. nos. 1 and 2:7-16.
- Barsalou, L.W. 1987 "The Instability of Graded Structure: Implications for the Nature of Concepts". In U. Neisser 1987:101-140.
- Bartlett, F. C. 1932 Remembering (London: Cambridge University Press).
- Billman, D. 1992 "Modeling Category Learning and Use". In Burns 1992:413-448.
- 3Boyd, R. 1989 "What Realism Implies and What it Does Not" Dialectica 43:5-29.
- Boyd, R. 1991 "Realism, Anti-Foundationalism and the Enthusiasm for Natural Kinds" Philosophical Studies 61:127-48.
- Burge, T. 1979 "Individualism and the Mental". In P. French, T. Uehling & H. Wettstein

- eds., Studies in Metaphysics: Midwest Studies in Philosophy IV: 73-121 (Minneapolis: University of Minnesota).
- Burge, T. 1982 "Other Bodies". In A. Woodfield, ed., Thought and Object, 97-120 (Oxford: Clarendon).
- Burge, T. 1986 "Individualism and Psychology" Philosophical Review 95:3-45.
- Burns, B. ed. 1992 Percepts, Concepts, and Categories ; The Representation and Processing of Information (North-Holland).
- Byrnes, J.P. & Gelman, S.A. 1991 "Perspectives on Thought and Language: Traditional and Contemporary Views". In Gelman & Byrnes 1991:3-27.
- Carey, S. 1985 Conceptual Change in Children. (Cambridge MA: The MIT Press)
- Carey, S. & Gelman, R. eds. 1993 The Epigenesis of Mind: Essays on Biology and Cognition (Hillsdale NJ: Erlbaum).
- Choi, S. Gopnik, A. 1993 "Nouns Are Not Always Learned Before Verbs: An Early Verb Spurt in Korean. In The Proceedings of the Twenty-fifth Annual Child Language Forum, 96-105 (Stanford: Center for the Study of Language and Information).
- Chomsky, N. 1995 "Language and Nature," Mind 104:1-61.
- Clark, E.V. 1991 "Acquisitional Principles. In Lexical Development". In Gelman & Byrnes 1991.
- Dodwell, P., Humphrey, K & Muir, D. 1987 "Shape and Pattern Perception". In Salapatek, P. & Cohen, L. Eds., Handbook of Infant Perception Vol 2: From Perception to Cognition: 1-77 (Harcourt Brace Jovanovich).
- Dromi, E. 1987 Early Lexical Development (Cambridge: Cambridge University Press)
- Fodor, J. A. 1972 "Some reflections on L.S. Vygotsky's Thought and Language" Cognition 1:83-95.
- Fumerton, R. 1989 "Russelling Causal Theories". In Rereading Russell:Essays in Bertrand Russell's Metaphysics and Epistemology, Minnesota Studies in the Philosophy of Science, vol.12., ed. C.W. Savage and C.A. Anderson, 108-118 (Minneapolis:University of Minnesota Press).
- Gallistel, C., Brown, A., Carey, S., Gelman, R. & Keil, F. 1993 "Lessons from Animal Learning for the Study of Cognitive Development". In Carey&Gelman 1993, 3-36.
- Gelman, S.A. & Byrnes, J.P. eds. 1991 Perspectives on Language and Thought (Cambridge:Cambridge University Press).
- Gelman, S.A. & Coley, J.D. 1991 "Language and Categorization:The Acquisition of natural kind terms". In Gelman Byrnes, 1991:146-196.
- Gentner, D. 1982 "Why Nouns are Learned Before Verbs; Linguistic Relativity versus Natural Partitioning". In S. Kuczaj, ed., Language, Thought and Culture; Language Development volume 2 , 301-334 (London:Erlbaum).
- Ghiselin, M. 1974 "A radical Solution to the Species Problem" Systematic Zoology 23:536-44.
- Ghiselin, M. 1981 "Categories, Life and Thinking" Behavioral and Brain Sciences 4: 269-283.
- Gilbert, D. 1993 "The assent of man:mental representation and the control of belief". In D.M.Wegner & J.W.Pennebaker eds., Handbook of Mental Control (Englewood Cliffs, NJ:Prentice-Hall).

- Gopnik, A. & Meltzoff, A. 1993 "Words and Thoughts in Infancy: The Specificity Hypothesis and the Development of Categorization and Naming". In C. Rovee-Collier & L. Lipsitt eds., Advances in Infancy Research 8:217-249 (Norwood NJ: Ablex).
- Gopnik, A. & Meltzoff, A. 1996 Words, Thoughts and Theories (Cambridge MA: MIT Press).
- Hacking, I.A. 1991 "On Boyd" Philosophical Studies 61:149-54.
- Hull, D.L. 1978 "A matter of individuality" Philosophy of Science 45:335-360. Reprinted in Sober 1994, 193-215.
- Ingram, D. 1989 First Language Acquisition; Method, Description and Explanation (Cambridge:Cambridge University Press).
- John, O. "Actions, Verbs and the Role of Context: Differences Between Categories of Objects and Those of Actions and Events" unpublished manuscript, University of Oregon and Oregon Research Institute, Eugene.
- Johnson, M.H., Dziuawiec, S., Ellis, H.D. & Morton, J. 1991 "Newborns' preferential tracking of face-like stimuli and its subsequent decline" Cognition 40:1-21.
- Johnson-Laird, P.N. 1983 Mental Models (Cambridge, Mass:Harvard University Press).
- Katz, J. 1972 Semantic Theory (New York: Harper & Row).
- Katz, J. & Fodor, J. A. 1963 "The Structure of a Semantic Theory" Language 39:170-210.
- Keil, F.C. 1983 "Semantic inferences and the Acquisition of Word Meaning". In T.B. Seiler W. Wannemacher, eds., Concept Development and the Development of Word Meaning (Berlin:Springer-Verlag).
- Keil, F.C. 1979 Semantic and Conceptual Development :An Ontological Perspective (Cambridge MA:Harvard University Press).
- Keil, F.C. 1987 "Conceptual Development and Category Structure". In U. Neisser ed., Concepts and Conceptual Development:Ecological and intellectual Factors in Categorization, 175-200. (Cambridge:Cambridge University Press).
- Keil, F.C. 1989 Concept, Kinds and Cognitive Development (Cambridge, MA:MIT Press).
- Keil, F.C. 1991 "Theories, Concepts, and the Acquisition of Word Meaning". In Gelman Byrnes 1991:197-221.
- Komatsu, L.K. 1992 "Recent Views of Conceptual Structure" Psychological Bulletin 112.3:500-526.
- Kornblith, H. 1993 Inductive Inference and its Natural Ground (Cambridge MA:MIT press).
- Kripke, S. 1972 Naming and Necessity (Cambridge MA:Harvard University Press).
- Lakoff, G. 1987 Women, Fire, and Dangerous Things (Chicago:Chicago University Press).
- MacFarlane, A. 1977 The Psychology of Childbirth (Cambridge MA: Harvard University Press)
- Markman, E. 1989 Categorization and Naming in Children (Cambridge, MA:MIT Press).
- Markman, E. 1991 "The Whole-object, Taxonomic and Mutual Exclusivity Assumptions as Initial Constraints on Word Meaning". In Gelman & Byrnes 1991, 72-106.
- Marler, P. 1993 "The Instinct to Learn". In Carey & Gelman 1993, 37-66.

- Medin, D.L. 1989 "Concepts and Conceptual Structure" American Psychologist 44.12:1469-1481.
- Medin, D.L. & Ortony, A. 1989 "Psychological Essentialism". In S. Vosniadou A. Ortony, eds., Similarity and Analogical Reasoning, 179-195 (New York:Cambridge University Press).
- Medin, D.L. & Schaffer, M.M. 1978 "Context Theory of Classification Learning" Psychological Review 85:207-238.
- Medin, D.L. & Smith, E.E. 1981 "Strategies and Classification Learning" Journal of Experimental Psychology:Human Learning and Memory 7:241-253.
- Medin, D.L. & Smith, E.E. 1984 "Concepts and Concept Formation" Annual Review of Psychology 35:113-138.
- Mervis, C.B. 1987 "Child-basic Categories and Early Lexical Development". In U. Neisser, ed., 1978:201-233.
- Mervis, C.B. Crisafi, M.A. 1982 "Order of Acquisition of subordinate-, Basic-, and Superordinate-level Categories" Child Development 53:258-66.
- Millikan, R. 1984 Language, Thought and Other Biological Categories (Cambridge MA:MIT Press).
- Millikan, R. 1986 "The Price of Correspondence Truth" Nous 20.4:453-468.
- Millikan, R. 1991 "Perceptual Content and Fregean Myth" Mind 100.4:439-59.
- Millikan, R. 1993a White Queen Psychology and Other Essays for Alice (Cambridge MA:MIT Press).
- Millikan, R. 1993b "On Mentalese Orthography". In Bo Dahlbom, ed., Dennett and His Critics, 97-123. (Oxford:Blackwell).
- Millikan, R. 1994 "On Unclear and indistinct Ideas". In James Tomberlin, ed., Philosophical Perspectives vol. VIII, 75-100 (Atascadero CA:Ridgeview Publishing).
- Minsky, M. 1975 "A Framework for Representing Knowledge". In P.H. Winston, ed., The Psychology of Computer Vision , 211-277 (New York:McGraw-Hill).
- Murphy, G.L & Medin, D.L. 1985 "The Role of Theories in Conceptual Coherence" Psychological Review 92:289-316.
- Nelson, C. & Horowitz, F., 1987 "Visual Motion Perception in Infancy: A Review and Synthesis". In Salapatek&Cohen 198:123-153.
- Nelson, K. 1974 "Variations in Children's Concepts by Age and Category" Child Development 45:577-84.
- Nelson, K. 1991 "The Matter of Time: Interdependencies between Language and Thought in Development". In Gelman 1991:278-318.
- Neisser, U. 1975 Cognition and Reality (San Francisco: W.H. Freeman).
- Neisser, U., ed. 1987 Concepts and Conceptual Development (Cambridge:Cambridge University Press).
- Piaget, J. 1926 The Language and Thought of the Child (New York: Harcourt Brace).
- Putnam, H. 1975 "The Meaning of 'Meaning'". In Language, Mind and Knowledge, Minnesota Studies in the Philosophy of Science, vol. 7, ed. Keith Gunderson (Minneapolis:University of Minnesota Press).
- Quine, W.V. 1960 Word and Object, (Cambridge, MA:MIT Press).
- Rifkin, A. 1985 "Evidence for a Basic-level in Event Taxonomies" Memory and

- Cognition 13:538-556
- Rosch, E. 1973 "Natural Categories" Cognitive Psychology 4:328-50.
- Rosch, E. 1975 "Universals and Cultural Specifics in Human Categorization". In R. Brislin, S. Bochner & W. Honner eds., Cross Cultural Perspectives on Learning, 177-206 (New York:Halsted).
- Rosch, E. 1978 "Principles of Categorization". In E. Rosch B.B. Lloyd eds., Cognition and Categorization, 27-48 (Hillsdale NJ: Erlbaum).
- Rosch, E. & Mervis, C.B 1975 "Family Resemblances: Studies in the Internal Structure of Categories" Cognitive Psychology 7: 573-605.
- Rosch, E., Mervis, C.B, Gray, W.D., Johnson, D.M. & Boyes-Braem, P. 1976 "Basic Objects in Natural Categories" Cognitive Psychology 8:382-439.
- Rumelhardt, D.E. 1980 "Schemata: The Building Blocks of Cognition". In R.J. Spiro, B.C. Bruce & W.F. Brewer, eds., Theoretical Issues in Reading Comprehension, 33-58 (Hillsdale NJ: Erlbaum).
- Schank, R.C. & Abelson, R.P. 1977 Scripts, Plans, Goals and Understanding (Hillsdale NJ: Erlbaum).
- Shepard, R. 1976 "Perceptual Illusion of Rotation of Three-Dimensional Objects" Science 191:952-54.
- Shepard, R. 1983 "Path-Guided Apparent Motion" Science 220:632-34.
- Smith, Edward E. & Medin, Douglas L. 1981 Categories and Concepts (Cambridge:Harvard University Press).
- Smith, C., Carey, S. & Wiser, M. 1985 "On Differentiation; A Case Study of the Development of the Concepts of Size, Weight, and Density" Cognition 21:177-237.
- Spelke, E. 1989 "The Origins of Physical Knowledge". In L. Weiskrantz, ed., Thought Without Language (Oxford: Oxford university Press).
- Spelke, E. 1993 "Physical Knowledge in Infancy: Reflections of Piaget's Theory". In Carey&Gelman 1993, 133-170.
- Tversky, B. & Hemenway, K. 1984 "Objects, Parts and Categories" Journal of Experimental Psychology: General 113:169-193.
- Ward, T.B. & Becker, A.H. 1992 "Intentional and Incidental Learning". In B. Burns 1992:451-491.
- Waxman, S.R. 1991 "Semantic and Conceptual Organization in Preschoolers". In Gelman&Byrnes 1991: 107-145.
- Wettstein, H. 1988 "Cognitive Significance Without Cognitive Content" Mind 97:1-28.
- Winograd, T. 1975 "Frame Representations and the Declarative-procedural Controversy". In D.G. Bobrow & A. Collins eds., Representation and Understanding: Studies in Cognitive Science 185-210 (San Diego:1953).
- Xi, F. & Carey, S. 1996 "Infant Metaphysics: The Case of Numerical Identity" Cognitive Psychology 30:111-153.