

A common structure for concepts of individuals, stuffs, and real kinds: More Mama, more milk, and more mouse

Ruth Garrett Millikan

Department of Philosophy, University of Connecticut, Storrs, CT 06250-2054

Electronic mail: millikan@uconnvm.uconn.edu

Abstract: Concepts are highly theoretical entities. One cannot study them empirically without committing oneself to substantial preliminary assumptions. Among the competing theories of concepts and categorization developed by psychologists in the last thirty years, the implicit theoretical assumption that what falls under a concept is determined by description (“descriptionism”) has never been seriously challenged. I present a nondescriptionist theory of our most basic concepts, “substances,” which include (1) stuffs (gold, milk), (2) real kinds (cat, chair), and (3) individuals (Mama, Bill Clinton, the Empire State Building). On the basis of something important that all three have in common, our earliest and most basic concepts of substances are identical in structure. The membership of the category “cat,” like that of “Mama,” is a natural unit in nature, to which the concept “cat” does something like pointing, and continues to point despite large changes in the properties the thinker represents the unit as having. For example, large changes can occur in the way a child identifies cats and the things it is willing to call “cat” without affecting the extension of its word “cat.” The difficulty is to cash in the metaphor of “pointing” in this context. Having substance concepts need not depend on knowing words, but language interacts with substance concepts, completely transforming the conceptual repertoire. I will discuss how public language plays a crucial role in both the acquisition of substance concepts and their completed structure.

Keywords: basic-level categories; categorization; child language; concepts; externalism; names; natural kinds; Putnam; theory of meaning.

1. 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 target article is 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 that the extensions of concepts (the set of things that fall under the concept) are determined by descriptions has managed to go unchallenged. This is true despite the fact that Putnam’s (1975a) and Kripke’s (1972) famous arguments against descriptionism (or at least their conclusions) have been rehearsed numerous times in the core psychological literature, and despite the fact that there have been a number of brave attempts to integrate these insights into the psychological tradition (Gelman & Coley 1991; Keil 1989; Komatsu 1992; Lakoff 1987; Markman 1989; Neisser 1987, Ch. 2). The difficulty is that these insights were almost

entirely negative. Moreover, the tentative positive views offered have 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 na-



RUTH GARRETT MILLIKAN is Professor of Philosophy at the University of Connecticut. She returned to professional philosophy in 1983 after raising children. Since then she has published *Language, Thought, and Other Biological Categories* (MIT Press), *White Queen Psychology and Other*

Essays for Alice (MIT Press) and thirty five papers in the philosophies of psychology, language, and biology. She has been a fellow at the Center for Advanced Study in the Behavioral Sciences and President of the American Society for Philosophy and Psychology. She delivered the Gareth Evans Memorial Lecture for the University of Oxford in 1991 and the Romanell Lecture for the American Philosophical Association in 1997. Currently she is Visiting Research Professor at University of Stockholm and is preparing a book on Concepts for Cambridge University Press.

ture 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's *Categories*) I will call "substances." The category of substances includes (1) things we would ordinarily call "substances," namely, stuffs such as gold, milk, and mud, (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 & Crisafi 1982; Nelson 1974b). 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 (having one meaning only). 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 sample 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 that, on the contrary, the earliest and most basic concepts 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 extension 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 them as having. For example, large changes can occur in the way 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 in the metaphor of "pointing" (Putnam said "indexicality"). Speaking literally, what is the structure of a substance concept in this view?

Having substance concepts need not depend on knowing words. Preverbal humans, indeed, any animal that collects practical knowledge over time of how to relate to specific stuffs, individuals, and real kinds, must have concepts of them. On the other hand, language interacts with substance concepts in vigorous ways to completely transform the conceptual repertoire. Putnam (1975) argued for what he called "the division of linguistic labor." On the basis of rather different reasons, I will argue similarly, that 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 (sects. 2 and 3). From this nondescriptionist vantage it will be easier to see just how descriptionism continues to be an ingredient in contemporary experimental work on concepts (sect. 4). Then I will discuss the nature of concept development from a nondescriptionist perspective (sect. 5) and finally the crucial involvement of language in the acquisition and use of substance concepts (sect. 6).

2. 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 of age: "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; Gelman & Coley 1991; Keil 1989; Markman 1989). 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, 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 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 figure in scientific theories, however, 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 intended 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 of "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 *encounters with kind K*, for example, over the kind *encounters with mice*. Compare this with making generalizations over *encounters with the stuff milk* and *encounters 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 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 in one encounter (temporal or spatial) that cats eat fish, and that 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 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, including all or most of what can be discovered about humans, along with many of his own stable properties. 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 and answered, once and for all, on the basis of one careful observation. 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 types,³ cultural artifacts, musical compositions, and so forth, but I will ignore these in the present essay.

It is not a matter of logic, of course, but rather of the makeup 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 cats are black does not carry over; next time a 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 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 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" (Boyd 1989; 1991; Hacking 1991a; Kornblith 1993; Putnam 1975). Everyone agrees, however, that what makes something a natural kind 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 (Millikan 1984, Ch. 16; forthcoming) is that there are many different reasons for the existence of real kinds. These reasons account for successes in generalizing over encounters in a variety of different ways. Sometimes there is a single underlying cause or inner structure (cf. Putnam's "natural kinds") that results, always, or under common conditions, in a certain selection of surface properties, as in the case of the various chemical substances. In such cases, the kinds have real essences, not merely nominal ones, 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 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, pp. 46–47)

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 is the (fully or partially) socially constructed kind, 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 (Millikan, forthcoming).

Turning now to what holds individuals together over encounters (over time), Xavier today is much like Xavier yesterday because Xavier today resulted directly 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 a standard 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 we should say that this individual thing was numerically the same as that individual thing occurring later in time. However, in the usual case we assume, quite rightly, that whether or not a correct identification of an individual has been made depends on how the world is, not on how we humans (or we English speakers) like to identify things. 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 about how we humans (or English speakers) like to classify things. If it is not like other members of the kind for the reason the other members are like one another, it is not a member of the kind. On the other hand, because the occurrence of causal factors accounting for similarities among members of a group can be more or less irregular, and because numbers of causally 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 about which one can learn from one encounter things to apply on other occasions *where this possibility is not coincidental but grounded*. That is, there is an explanation or cause of the samenesses.

I now wish to show that it is plausible that despite the many different kinds of groundings that account for the unity of various types of substances, the basic structure of a concept of 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 into the mind itself the structure of the object grasped by thought. I will argue the contrary, namely that substances are grasped not by understanding the structures or principles that hold them together but by knowing how to exploit these substances for information gathering purposes. Knowing how to use a thing is not knowing facts about its structure.

3. Concepts of substances

The “concept” of a substance, as I use that term,⁴ is the capacity to represent the substance in thought for the purpose of information gathering and storage, inference, and ultimately the guidance of action. We wish to know the structure of this ability. To describe the structure of an ability is to say what it is an ability to do, what sub-abilities are contained in it and, ultimately, by what means it is exercised, that is, how exactly 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 of 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 theoretical or practical use, over numerous encounters with them. This makes it possible for the organism to store 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 (and the dog’s) first recognitions must be merely of *more Mama*, *more milk*, and *more mouse* (Quine 1960, p. 92). Children observe things about *Mama* when they encounter her, not about samples or instances of *Mama*. Similarly, to learn things about milk, they 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 using it, one does not distinguish Amos from Amos’s brother; one conceives them 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 concepts 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 affordance of sittability-upon, across multiple encounters, but *Mama* does not (you cannot sit on *Mama* when she is standing). *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, 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 relation to 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 the glass 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. 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 uses 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 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 thing (yogurt). To caricature,⁵ if I represent yogurt to myself in one way, with a mental heart, as I store away the knowledge that yogurt is good to eat, but represent it another way, with a mental diamond, as I store away the information that it is in the refrigerator, these bits of information will not help me when I am hungry. A fundamental subcapacity 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 understands to be another 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 indicators would be pieces of information about the presented substance – that it has this or that property that marks it reliably enough.

In the case of familiar substances, we typically collect numerous means of identification over time, all of them fallible, and certainly none of them "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 the look 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 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.

It follows that it cannot be one's 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). How does Lakoff know that two-year-olds are not thinking

of lions and tigers that they *are* kitties – kitties grown big? A little more experience and children may change their minds – 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. The same is true of a child who has got only part way toward 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, that is, “determinables,” that are applicable to the substance. That is, one must understand what are some of the meaningful questions to ask about it (see Millikan 1984, Ch. 15, pp. 252ff., and Chs. 16 and 17). 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 recognizing the substance might be good for in the context of developing either practical skills or theoretical knowledge.

4. Contrast with descriptionism

In contrast to the position just sketched, the descriptionist is one who holds that the referent or extension of a substance term is determined by its falling under a description associated with the term by the 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 extension – determine it to be a portion of the extension, 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). But it takes many other forms as well.

Thus, using the concept *chair* as his example, Komatsu (1992) describes the most general question that psychological theories of concepts have attempted to answer as follows: “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). Building on Medin and Smith (1981; 1984), Komatsu 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; Minsky 1975; Neisser 1975; Piaget 1926; Rumelhardt 1980; Schank & Abelson 1977; Winograd 1975] . . . the explanation-based view (e.g., Johnson-Laird 1983; Lakoff 1987; Murphy & Medin 1985) [later he cites Carey 1985; Gelman 1988a; 1988b; Gelman & Markman 1987; Keil 1989].

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; 1975b). 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 the 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 a realist view. Realism and descriptionism might seem incompatible. If the extension 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 realism and descriptionism have been combined.

The simplest way to combine them is to take the extent of a substance term to be fixed by one, or a set, of definite descriptions of the substance.⁶ Thus the classical twentieth 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 combination of either/or 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 extension 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, my emphasis) 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). In other words, it is assumed 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 “we 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 & 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 transi-

tion 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 et al. 1985] for making a decision. . . . These are difficult issues, and it is hardly surprising that they are not yet resolved. (Keil 1989, pp. 21–22)

Following Smith et al., Keil speaks of “tracking” concepts across theory change” and agrees with them that probably “descent can be traced . . . because of several properties of theories that stay fixed through change” (Smith et al. 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, pp. 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 thing. 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 one’s 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 & 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 & 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.

5. The development of substance concepts

Viewing a substance concept as an ability to reidentify, which a mobile person comes to exercise within a supporting but changing environment, the study of concept development is also seen in a new light. What subskills are involved in this ability? What is the characteristic progression toward acquiring 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 (Byrnes & Gelman 1991; Clark 1991; Waxman 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 infants recognize their mothers and dogs recognize their masters. Each has the capacity to reidentify the relevant individual

under diverse conditions, thus making it possible to learn how to behave appropriately in their presence.

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, and so forth. 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 (Dodwell et al. 1987; Nelson & Horowitz 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 being 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 one's 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. Let us call this "conceptual tracking": one understands rather than perceives that the substance is the same one again. Out of what materials is it that our abilities conceptually to track substances are built?

By tracking a substance perceptually one can learn many different ways to recognize it: how it looks, how it sounds, how it feels, the way 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 can appear as the same under widely varying lighting conditions, and voices can sound like 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 that prove to be better trackers, 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 do children know which aspects of the substances they are learning to track can be relied on for reidentification? And how do they know what questions they should expect to be answerable about each substance? Just as children have built-in perceptual tracking abilities and built-in perceptual constancies, we might expect them to have certain built-in conceptual tracking abilities.

There is evidence, for example, that infants may have inborn 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). In addition, many species that recognize conspecifics as individuals instinctively use smell for this purpose; and (Dan Dennett has reminded me) human infants also know Mama by smell in the early months (MacFarlane 1977). It appears that innately the infant may know at least two good ways to track individual conspecifics conceptually. Faces and personal odors are indicative of individual identity; clothes, postures, and so forth, are not.

The mechanisms by which infants reidentify individuals perceptually do not appear to rely on the invariance of properties of the tracked object but upon common movement, spatial location, and trajectory (Gopnik & Meltzoff 1996). Xu and Carey (1996) have recently produced experimental evidence that 10-month-old infants, unlike 12-month-olds, 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 experi-

mental 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 discern what kinds of traces are followed in attempting conceptual tracking at various ages and for different domains of real kinds. To be acutely sensitive to correlations among properties, probably among specific kinds of properties in specific domains (cf. Atran 1989; Carey 1985; Gallistel et al. 1993; Gelman & Coley 1991; Keil 1979; 1989; Markman 1989; Marler 1993; Spelke 1989; 1993) seems an obvious way to track many kinds of substances. But experiments need to be designed and interpreted bearing 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 and 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 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 one’s 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 either the members of A’s specific domain or with A itself, during the evolutionary history of the species 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 or for the tuning of the (would-be) tracking dispositions, then the concept is equivocal. For example, to have two people “mixed up” or “confused” in one’s mind is to have an equivocal substance concept (Millikan 1984, Ch. 15; 1991; 1993a, Ch. 14; 1994; 1997).

We now move to a still more fundamental medium through which conceptual tracking is achieved, namely, language.

6. 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) noted: knowing neither how to identify these substances in the flesh, nor by any unique or defining descriptions. That is, neither verifica-

tionist nor descriptionist theories of concept extension explain these cases. 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 exaggerated – that 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 directly generates a belief, unless cognitive work is done to prevent this, just as with what one perceives through other media. Loading the cognitive systems with other simultaneous tasks, such as having to count backwards by threes, has the effect of facilitating belief fixation regarding whatever one hears or reads (Gilbert 1993).

There are two things that distinguish direct perception quite sharply from the acquisition of information through language, but neither implies a difference in immediacy. In direct perception, the spatial and temporal relation of the perceiver to the object perceived is given, whereas it is not normally given through language. On the other hand, in 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 a lot like perceiving through the media, which in turn is a lot like directly perceiving the original.

Think of the matter this way. There are many ways to recognize, say, rain. There is a way that rain feels when it falls on you, and a way that it looks through the window. There is a way that it sounds falling on the rooftop, “retetetetet,” 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!” (Thanks to Crawford Elder for this example.) 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 television 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; second, it depends on where you focus your mind, your attention.

But there is no need to belabor 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 television 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, although 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 substance as manifestations of it; one learns how to translate information arriving in one more kind of sensory package 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 (cf. 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, such as the natural relation of a speaker’s experience and the context of his speech to 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.⁷

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, that is, in the absence of any ability to recognize the substance in the flesh. For most of us, that is how we have a concept of Aristotle, of molybdenum, and, say, of African dormice. There, I just handed you a concept of African dormice, in case you had none before. Now you can think of them at night if you want to, wondering what they are like – on the assumption, of course, that you gathered from their

name what sorts of questions you might reasonably ask about them (animal questions, not vegetable or mineral or social artifact questions). In many cases there is not much more to having a substance concept than having a word. To have a word is to have a handle on tracking a substance via manifestations of it produced in a particular language community. 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 phrase, "The Division of Linguistic Labor" (Putnam 1975) and Burge's claim that the constitution of the very content of one's *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 in such a way 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 (for a full discussion of equivocation in concepts, see Millikan 1993a, Ch. 14; 1993b; 1994). But Gelman and Coley (1991, p. 184) are surely right that "a word can serve to stake out a new category, which then must be explored in more depth" (see also Gopnik & Meltzoff 1993). Words are handles to hang onto, helping to stabilize concepts so as gradually to eliminate equivocation in thought, as long as those who speak to us have adequate concepts themselves.

But have we not overlooked an obvious distinction here between merely knowing a word and knowing what that word means? In the present view there is an interesting question about 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 as 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 outcome; other times it is what accomplishes the same outcome 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 chemists have much more sophisticated and reliable means at their disposal for identifying sugar than do children. Similarly, we could ask, did Helen Keller have many of the same concepts as you and I, or did she have different ones, and again the answer would be equivocal. Suppose we say that children have the "same concept" as

chemists, namely, the concept of sugar, but that their "conception" of sugar is very different from that of chemists, for children use very different methods to identify it. 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 simply that it is referring to *that* substance. To know what the word means is just to have a concept of the substance that includes knowing how 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?) substances, it may be that there are core methods by which nearly every adult (the "nearly" is for Helen Keller) knows how 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 children do not fully understand "the meaning" of the word for that substance until their 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."

NOTES

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

2. An in-depth discussion of the ontological category of substances can be found in Millikan 1984, Chapters 16 and 17.

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

4. I do not recommend generalizing this description of a concept 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.

5. To model the act of reidentifying a substance in thought as using the same mental term again, as I do here, is actually a crude and misleading expedient (see Millikan 1991; 1993b; 1994; 1997).

6. 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.

7. For a full discussion of equivocation in concepts, see Millikan 1993, Chapter 14; 1993b; 1994; 1997.

Open Peer Commentary

Commentary submitted by the qualified professional readership of this journal will be considered for publication in a later issue as Continuing Commentary on this article. Integrative overviews and syntheses are especially encouraged.

Animal concepts

Colin Allen

Department of Philosophy, Texas A&M University, College Station, TX
77843-4237. colin.allen@tamu.edu snafell.tamu.edu/~colin/

Abstract: Millikan's account of concepts is applicable to questions about concepts in nonhuman animals. I raise three questions in this context: (1) Does classical conditioning entail the possession of simple concepts? (2) Are movement property concepts more basic than substance concepts? (3) What is the empirical content of claiming that concept meanings do not necessarily change as dispositions change?

If having concepts does not depend on knowing words, and if a substance concept is an ability to reidentify units in nature, then we might reasonably expect to derive insights into the nature and evolution of concepts from ethological studies of animal cognition. Although I will focus on challenges to some of Millikan's claims arising from a consideration of ethological research, I think that overall her proposal to orient questions about concepts to questions about capacities is extremely useful.

The ability to recognize predators and learn about them is of obvious survival value. Vervet monkeys and chickens are among the species that have evolved signal systems that are specific to types of predator. Vervets distinguish at least four categories of predator by producing different alarm calls (eagle, snake, leopard, primate), and as infants they learn when to produce these calls, for example, refining production of the "eagle" call from a broad category, including objects such as falling leaves, to very specific identification of just those species of raptor that actually prey on vervets (Cheney & Seyfarth 1990). Chickens produce distinct calls for aerial and terrestrial predators, but an artificially projected image of a terrestrial predator on an overhead monitor elicits aerial-predator vocalizations, albeit at a lower rate than an image of an aerial predator (Evans & Marler 1995).

These brief sketches may be insufficient to settle the question of whether the capacities of vervets and chickens entail that they have predator concepts. However, I believe that Millikan would willingly attribute concepts to these organisms. Indeed, even the widespread capacity for classical conditioning, which requires the reidentification of both conditioned and unconditioned stimuli (CS/UCS) in order to form the association between CS and response, seems to fit her account of concept possession. I expect that Millikan would wish to distinguish different levels of sophistication with respect to concepts, and an articulation of those levels would be of great interest to those pursuing comparative and phylogenetic studies of cognition.

Granting, for the sake of argument, that concepts are implicated in the alarm-calling behavior of both vervets and chickens, what else can these data tell us about Millikan's proposals? One suggestion that is less than convincing is that "concepts of properties . . . [are] less fundamental than those expressed with simple concrete nouns" (sect. 2, para. 1). Consider the chickens' tendency to give aerial-predator calls to images of raccoons shown circling overhead. What seems most salient to the chickens is not what natural kind the object belongs to – although, as Evans and Marler (1995) showed, this is not entirely ignored – but what the object is doing, namely, circling overhead (Allen & Sidel 1997). For infant vervets, the first basis for categorization also seems to be a dynamic property, such as motion overhead. Only later are the vervets

tuned in to the kinds involved. My hunch, contra Millikan, is that dynamic properties may be knowable prior to the kinds, individuals, and stuffs that she regards as most basic. Although I agree that the substance concepts may be more basic than static property concepts such as color or shape, I suspect that the tendency to forget about the importance of dynamic properties may be as a result of the focus, shared with most psychologists, on human concept-acquisition during word learning, a development that is relatively late both phylogenetically and ontogenetically and that depends on substantial prelinguistic conceptual abilities on the part of infants.

Finally, Millikan's suggestion that a child's ability to reidentify a substance may be partial, and thus there is no need to treat a change in application of a term over time as a change of concept is a good one. It is not clear, however, when, if ever, she thinks it would be appropriate to describe learning as conceptual change. Here, the ethological research can provide examples to consider. Specifically, one might ask how, if at all, we settle the question of whether a vervet infant's vocalization changes meaning as it evolves from a response to any object moving in the sky to a very specific response to just those eagles that prey on vervets. Of course one can make the case that the real extent of the set of objects that are aerial predators (mediated by the responses of adult vervets to those predators) has governed the development of these dispositions in infant vervets. But can anything falsify (or otherwise provide empirical content to) the claim that the ontogeny of the vervet aerial-predator call involves the refinement of their capacity to recognize the members of this category rather than a change of concept?

Different structures for concepts of individuals, stuffs, and real kinds: One Mama, more milk, and many mice

Paul Bloom

Department of Psychology, University of Arizona, Tucson, AZ 85721.
bloom@u.arizona.edu

Abstract: Although our concepts of "Mama," "milk," and "mice" have much in common, the suggestion that they are identical in structure in the mind of the prelinguistic child is mistaken. Even infants think about objects as different from substances and appreciate the distinction between kinds (e.g., mice) and individuals (e.g., Mama). Such cognitive capacities exist in other animals as well, and have important adaptive consequences.

In her interesting and provocative target article, Millikan addresses the question of what different types of substance concepts (such as "Mama," "milk," and "mouse") have in common. The commonalities that she finds motivate a developmental claim. Following Quine, Millikan suggests that our earliest and most basic concepts of substances "have an identical ontological structure" (sect. 1, para. 3). Although I am sympathetic to her overall theoretical framework and conclusions, I will devote this commentary to arguing against this developmental claim, on both empirical and theoretical grounds.

In his original formulation, Quine (1960) suggested that the acquisition of the grammatical count/mass distinction leads children to think differently about objects and about stuff. This claim was addressed by Soja et al. (1991) who taught new words to 2-year-olds who had not yet learned this grammatical distinction. These children nevertheless treated names for objects very differently from names for stuff: names of novel objects were extended to other entities of the same shape, regardless of substance, whereas names of novel substances were extended to other entities of the same substance, regardless of shape. Other studies have found that even in their earliest utterances, 1- and 2-year-olds treat common nouns (e.g., "milk" and "mice") differently from proper

names (e.g., “Mama”), suggesting that they distinguish words that refer to kinds from those that refer to individuals (Bloom 1990; Bloom & Kelemen 1995; Katz et al. 1974).

What about infants? Some recent studies reveal an intriguing pattern of results. A pyramid is placed in front of an infant. A screen then rises and hides the pyramid. Then the screen drops, and the pyramid is gone. How does the infant react? It depends. If the pyramid was previously seen by the infant as a bounded and coherent entity, the infant is surprised when it disappears (Baillargeon et al. 1985), but the infant is not surprised if the pyramid was originally created by piling up several smaller objects (Chiang & Wynn 1997). Similarly, if an object is dropped behind a screen, infants expect the object to be there when the screen is moved, but if sand is poured behind a screen and then the screen is moved, they are unsurprised if the sand has disappeared (Huntley-Fenner & Carey 1995). These findings show that the capacities infants have are not tracking mechanisms but rather *object-tracking* mechanisms, and they apply only to entities that are seen as obeying certain principles that only objects have (Spelke 1994).

In another set of studies, a Mickey Mouse doll is placed in front of an infant, and then a screen rises to hide it. The infant then observes a hand place another identical Mickey Mouse behind the screen. What does the infant expect to see when the screen drops? Mouse? More mouse? Actually, the infant expects to see two Mickey Mouse dolls and is surprised if there is either one or three (Wynn 1992). Similar numerical abilities are present in other species, such as birds and rats, and underlie behaviors – for example, determining relative rate of return when foraging – that are of clear adaptive value (Gallistel 1990). Moreover, the result above suggests that infants appreciate that there can exist distinct individuals that belong to the same kind, just as they appreciate that something can be the same individual even if it changes kinds, say, from a duck to a ball (Xu & Carey 1996). Taken together, these findings suggest that infants’ judgments about what *kind* an object belongs to are distinct from their judgments about which *individual* the object is.

Millikan may be right that all substance concepts “point” to natural units. But even putting aside issues of enumeration and quantification, these concepts do their pointing in quite different ways. If there is a mouse in the car, there might well be another mouse in the bathroom. But if Mama just drove away, then the woman in the bathroom certainly is not Mama. Spatiotemporal continuity is irrelevant for judging what kind an object belongs to but is essential for judging which individual it is, and tracking a particular individual is often essential to an animal’s survival. Newly hatched ducklings will follow whatever moving object they first see and will form an attachment to that object. This tracking behavior is adaptive, since that specific object – but not other objects that look just like it – is typically the bird’s mother. Predators of herding animals will pick a single animal from the group and chase that specific animal, trying to wear it down. If predators were to switch quarries, they would get exhausted but their prey would not (Pinker 1997).

Given these functional advantages of thinking about certain entities as distinct individuals that can be counted, tracked, and categorized (see also Bloom 1996), the Quinean proposal that there is no difference between our initial thoughts of Mama, milk, and mice is not only bad developmental psychology, it is bad evolutionary biology as well.

If “tracking” is category-specific a “common structure” may be redundant

Pascal Boyer

Centre National de la Recherche Scientifique, 69363 Lyon, France.
pboyer@mrash.fr

Abstract: Identifying objects as members of ontological domains activates category-specific processes. There is evidence that these processes include particular ways of “tracking” substances and could do all the “tracking” necessary for concept acquisition. There may be no functional need or evolutionary scenario for a general tracking capacity of the kind described by Millikan.

Millikan’s nondescriptionist account requires a “general capacity to reidentify” substances, a “tracking” capacity (TC) common to stuffs, individuals, and kinds. Millikan also mentions that the child needs additional cognitive abilities to know, not just how to reidentify a substance, but also “what are some of the meaningful questions to ask about it” (sect. 3, para. 10). There is evidence for the early development of such abilities. Even infants seem to have some assumptions about which questions concerning individuals are more meaningful, depending on whether these individuals are living or artifactual, humanlike or not. Identifying objects as members of particular ontological domains activates category-specific principles.

I suggest that all the “tracking” that needs to be done may well be accomplished by such specialized principles. In this view, domain specificity would start at a low level of processing, that of “tracking” same-substance cues. Representing two faces as the same person with different emotions would activate cognitive resources entirely different from those involved in recognizing two sheep as same kind or two differently shaped telephones as same function. Note that this does not imply a return to the view that having a concept is having a description. All that is required is that (1) objects are assigned to ontological domains on the basis of perceptual cues (e.g., motion for animate-inanimate), (2) this makes the system attend to or expect “same substance” information in a way that is specific to the domain. There is evidence for this phenomenon from development, normal adult processes, and pathology.

In development, category-specific principles are often described as resulting in particular expectations about same-substance information. But the evidence supports the view that category-specific principles are the outcome of such specialized tracking rather than its causes. In face recognition, identifying two face-episodes as one person helps in representing the goals underlying people’s intentional action. In living kinds, children assume that same-species animals have the same “innards” (more important than same appearance) before they have representations of the causal role of those innards. For artifacts, same-use (namely, same action sequence) is more salient than same shape for sameness of stuff. These principles do not cross domains: two persons doing the same thing in the same way remain two, not one. A striking illustration of this specificity of tracking is Hall’s (1993) work on preschoolers’ intuitions of preserved identity for an object whose parts have all been replaced one by one. For many subjects, object identity is preserved if cues suggest that the object is a living thing, with opposite intuitions for an artifact. So even intuitions about whether something is the same over time (the very point of “tracking” capacities) is different for different categories.

The clinical literature seems to suggest that low-level “tracking” capacities underpin domains. Category-specific impairment used to be interpreted either as localized damage or as impairment of high-level conceptual resources; but it is now clear that it can result from general deficits such as Alzheimer’s disease. So differential use of perceptual cues (with consequent differences in what counts as same substance) may create category-specific impairment when the access to or the organization of those cues are damaged.

In normal adult subjects there is likewise evidence for such specialization at the level of encoding. To give but one illustration, consider Schacter and Cooper's (1993) study of novel artifact-like objects, primed for either functional or structural features. Differences in recognition and recall rates suggest that the recognition of objects as same-stuff differs for structure and function, and that these are differences of encoding, not high-level inference.

To sum up, there are low-level computational differences between ontological domains, in terms of the processes that build projectable "same-substance" criteria. Millikan herself suggests that "just as children have built-in perceptual tracking abilities and built-in perceptual constancies, we might expect them to have certain built-in conceptual tracking abilities" (sect. 5, para. 9), but she concludes 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 one's substance concepts" (sect. 5, para. 13). True, one does not need to have category-specific tracking in principle in order to have a substance concept. In actual fact, however, children seem to start with category-specific assumptions that include particular modes of tracking.

A consequence would be that there is no real functional need for a general tracking capacity of the kind described by Millikan. The proposal may seem unparsimonious. In different ontological domains, the different "tracking capacities" (TCs) have a similar function, namely, to produce reliable same-substance intuitions, as Millikan rightly argues. They may also have some computational properties in common, and processing would be more economical if those similar resources were shared instead of duplicated. From an evolutionary viewpoint, however, the "many TCs" description is more plausible. Evolution by gradual increments makes the emergence of a general TC less likely than the accretion of specialized TCs. Humans, in this domain as in others, may differ from other species in having more of those specialized capacities rather than having replaced them with one big domain-general processor by an evolutionary miracle.

Concepts in artificial organisms

Angelo Cangelosi^a and Domenico Parisi^b

^aCentre for Neural and Adaptive Systems, University of Plymouth, PL4 8AA Plymouth, United Kingdom

angelo@soc.plym.ac.uk www.tech.plym.ac.uk/soc/research/neural

^bInstitute of Psychology, CNR – National Research Council, 00137 Rome, Italy. domenico@kant.irmkant.rm.cnr.it gracco.irmkant.rm.cnr.it

Abstract: Simulations with neural networks living in a virtual environment can be used to explore and test hypotheses concerning concepts and language. The advantages that result from this approach include (1) the notion that a concept can be precisely defined and examined, (2) that concepts can be studied in both nonverbal and verbal artificial organisms, and (3) concepts have properties that depend on the environment as well as on the organism's adaptive behavior in response to the environment.

Talk about concepts is infused with language. We tend to identify concepts with the meanings of words and the use of concepts with applying a linguistic label to some entity or with responding appropriately to linguistic labels. But, as Millikan reminds us, we should be able to talk about concepts of preverbal humans or of "any animal that collects practical knowledge over time of how to relate to" (sect. 1, para. 5) its environment. If we are unable to do so we can also study how learning or evolving a language can make a difference to the concepts of a preverbal infant or an evolving population of organisms previously lacking a language. The problem is that studying concepts independently of language is difficult both empirically and theoretically. One way of overcoming this difficulty is to use simulations with artificial organisms.

In our simulations (Cangelosi et al. 1997; Cangelosi & Parisi, submitted) a population of artificial organisms lives in an environ-

ment with both edible and poisonous mushrooms. Each individual's behavior is controlled by a neural network that encodes the location and perceptual properties of the nearest mushroom in two separate sets of input units, and the organisms' movements in its output units. Because each mushroom is unique, upon encountering a mushroom an organism must first recognize the mushroom as either edible or poisonous and then generate the appropriate movements (e.g., approach and eat an edible mushroom, or go away from a poisonous one). A genetic algorithm is applied to the population of organisms. Organisms reproduce as a function of their energy level, which is increased when they eat edible mushrooms or decreased when they eat poisonous ones. Reproduction consists of generating one or more new neural networks (offspring) with connection weights identical to the connection weights of their parent network, except for a few random mutations. The network architecture is identical in all individuals and includes a layer of internal units to which only the input units encoding a mushroom's perceptual properties send their connections. Edible and poisonous mushrooms have perceptual properties encoded as binary patterns that deviate by one or two bits from two different prototype patterns. The pattern of activation observed in this layer of internal units when the organism encounters a mushroom is the internal representation of the mushroom. This internal representation is further elaborated in a second layer of internal units along with the input encoding the mushroom's location, to determine the network's output (movements).

By examining the internal representations of mushrooms we observe that at the beginning of the simulation – when the organisms have randomly assigned connection weights and hence their behavior is inefficient – these representations tend to be randomly distributed within the entire abstract space of possible activation patterns of the internal units. After a certain number of generations, however, the selective reproduction of the best individuals and the random mutations cause an increase in the average performance level of these organisms (namely, eating the edible mushrooms and avoiding the poisonous ones) and a parallel change in the internal representations of mushrooms: edible mushrooms tend to have similar internal representations, and the same is true for poisonous mushrooms, but edible mushrooms tend to have internal representations different from poisonous mushrooms. Internal representations of this can be thought of as "concepts," that is, internal (neural) representations that tend to minimize the differences between entities in the environment that require same type of behavior from the organisms and to maximize the differences between entities that require different behaviors. Our organisms possess two concepts: the concept of an "edible mushroom" and of a "poisonous mushroom."

In these simulations an individual mushroom can generate only a single input for the neural network. But suppose the same mushroom can generate a number of different inputs; for example, the input encoded in the network's input units when the mushroom is viewed from one side is different from the input encoded when the mushroom is viewed from another side. If the organisms are to survive, the internal representation evoked by the various inputs generated by the same mushroom should be the same. Evolved organisms should be able, for example, to approach and eat an edible mushroom on the basis of any current sensory input from the mushroom; this is facilitated if all inputs from the mushroom evoke the same internal representation. This is what makes it possible to interpret the "concepts" possessed by an organism as concepts of "substances," that is, internal representations that carry over from one encounter with the same mushroom to the next encounter. Internal representations of this type have "rich inductive potential" in that, although they are evoked by a single input, the organism "knows" from the internal representation that the entity with which it is dealing can also generate a number of further inputs.

Millikan is interested in concepts of entities that, in the English language, would be referred to by using mass, count, or proper nouns. However, for nonverbal organisms it is not clear that

concepts of entities that we would refer to with adjectives, such as white (thing), are really different. Suppose our organisms are trained (evolutionarily) to approach and eat white mushrooms and to avoid red ones. Properties such as white or red may have less inductive potential but they do have some. For example, we know that white things tend to become dirty more easily than darker color things.

What difference does it make for our organisms, in particular for their concepts, to have a language? In another set of simulations, each time an organism encounters a mushroom it also “hears” a linguistic label that is identical to all edible mushrooms but different from the linguistic label accompanying poisonous mushrooms. An organism’s neural network includes an additional set of input units that encode this linguistic label. These input units send their connections to the “conceptual” internal units. The results show that with “language” there is an improvement in the concepts of our organisms. The internal representations of edible mushrooms become even more similar among themselves, and so it is with poisonous mushrooms, whereas the distance is increased between the internal representations of the two types of mushrooms. At the same time, the practical behavior of the organisms with language is more efficient than the behavior of the non-linguistic organisms. As Millikan suggests, linguistic labels are just another input for the organism, along with the inputs from the mushrooms. But being identical for all cases requiring an identical (or very similar) internal representation, labels help construct these representations.

At the end of evolution, linguistic labels help evoke the appropriate internal representation and the appropriate behavior both when they are experienced together with a perceived mushroom and when they are experienced alone. The next step is to have linguistic labels acquire an ability to evoke an appropriate internal representation even if they are never experienced along with environmental entities but only in association with other linguistic labels. This would capture in the simulations of how children can acquire concepts directly via language.

In conclusion, simulations with neural networks living in an environment can be used to explore and test hypotheses concerning concepts and language. The advantages that result from this approach include (1) the notion that a concept can be precisely defined and examined, (2) that concepts can be studied in both nonverbal and verbal artificial organisms, and (3) concepts have properties that depend on the environment as well as on the organism’s adaptive behavior in response to the environment.

Names, and what they are names of

Greg Carlson

Department of Linguistics, University of Rochester, Rochester, NY 14627.
carlson@ling.rochester.edu

Abstract: Terms designating substances and kinds function grammatically much like proper names of individuals. This supports Ruth Millikan’s theory, but it also poses the question of how we can understand the reference of kind terms when the ontological status of the kind term is uncertain or disputed.

Millikan’s view that certain terms of natural language refer to “substances” in a nondescriptive way has a good deal to recommend it, but I do wish to pose one question later. First I wish to comment on some grammatical issues that expand on points made in the target article, and then turn to the “sticky” issue of dinosaurs.

The Aristotelian arrangement of “substances” as discussed here likens terms referring to individuals (proper names) to “substances” in the other sense used here, namely, as designated by mass terms – “gold,” “water,” and so on. The “secondary substances,” kinds of individuals such as lions and houses, are then classified differently. However, from a grammatical and semantic

point of view, it has been noted widely that mass terms such as “gold” and “water” have a great deal in common with count “kind” terms such as “lions” and “pencils,” which grammatically appear in the plural (in English); and that proper names of individuals differ grammatically and semantically from these. Thus, mass terms appear to designate entities that act much more like secondary substances than primary ones, from a grammatical point of view. What is of more interest from this point of view is that both mass terms (e.g., “water”) and kind terms (e.g., “lions”) have distributional and semantic properties in common with proper names, and are not at all like noun phrases occurring with determiners such as “this,” “all,” “many,” and so on. So, for instance, the phrase “so called” may be applied not only to proper names as in (1) below, but also to kind and mass terms, as in (2):

1. Slim is so called because of his build.
2. a. Cardinals are so called because of their color.
b. Coke is so called because it once contained cocaine.

Based on observations of this type, along with several others, the theory of reference to kinds presented in Carlson (1980) argues that mass and kind terms should also be regarded as names, although of sorts of things that may be regarded as individuals, differing from ordinary individuals such as you or the Eiffel Tower in certain ways but ontologically of the same sort, as Millikan outlines here. Thus, from a grammatical point of view, the facts would appear as Millikan suggests. This allows these terms to make direct, nondescriptive reference to substances in much the same way as Kripke (1980) suggests for individuals and Putnam (1975) for kinds.

The underpinnings of this point of view, however, require that the objects designated by such terms have objective significance, which may include, apart from real essences, utilitarian function, social function, and other matters as discussed in section 2. From what is presented here, it is a little difficult to determine just what sorts of terms are excluded from having the status of substances. The few concrete examples are complex (“red square”; “two-inch malleable object”) as opposed to the numerous simple substance-referring examples. Based on their description, non-basic-level terms also do not designate substances. So, whereas “chair” would designate one, “furniture” would not; or, “dog” would designate one, but “dachshund” would not. Yet, the target article provides us with no reason to think this might be so. What concerns me most – and this concern is by no means confined to this target article – is how, in a realist theory of the sort discussed here, we get “neutral” representations of the sort we must countenance. To illustrate, consider (my knowledge of) dinosaurs. First, I would surmise that the term “dinosaur,” as opposed to “*tyrannosaurus rex*” or “*pterodactyl*,” is the basic-level term, being far more frequent, earlier learned, monomorphemic, and more. Now, as a child I was, like many, fascinated by dinosaurs, reading about them in books, going to museums. I vaguely recall learning that they were very large and fearsome reptiles that lived way, way in the past some time, and are now extinct. I also vaguely recall learning that, since dinosaurs are extinct, present-day animals such as iguanas, crocodiles, and Komodo dragons are not dinosaurs. The logic of this seemed inescapable. More recently, I heard something about some theory of dinosaurs saying they were warm-blooded and much more mammalian creatures than previously thought. And, since a crocodile or a Komodo dragon are not like this, they are not a dinosaur; the logic of this, too, is quite clear. But we now have two competing ideas about dinosaurs, one where they are a nominal kind (reptile, typically big, that lived a long time ago), and the other where they are a natural kind.

So, when we say of dinosaurs that they once roamed the earth, what is the referent of “they” in this sentence? You cannot be assured that it is fixed by a real kind, even if you are a scientist who fervently believes this to be true; we cannot be assured that it is fixed by usage alone since the scientist might just be right. Note that in either case, the amount of knowledge that carries over from one instance to the next – dinosaurs being extinct and hence a closed finite class of individuals – is the same; so this

criterion is of no help, either. This is why I wish Millikan had developed the last tantalizing bit at the end of her article when she asks, "But have we not overlooked an obvious distinction here between merely knowing a word and knowing what that word means?" (sect. 6, para. 10). Achieving a separation between the two – if the are indeed different – might not only give us a better idea of what concepts are, it might also teach us something about what they are concepts of.

What is the point? Concepts, description, and rigid designation

Bradley Franks^a and Nick Braisby^b

^a*Department of Psychology, London School of Economics, London WC2A 2AE, United Kingdom. B.franks@lse.ac.uk;*

^b*Department of Psychology, London Guildhall University, London E1 7NT, United Kingdom. braisby@lgu.ac.uk*

Abstract: Millikan's nondescriptionist approach applies an account of meaning to concepts in terms of designation. The essentialism that provides the principal grounds for rigid designation, however, receives no empirical support from concepts. Whatever the grounding, this view not only faces the problems of rigid designation in theories of meaning, it also calls for a role for pragmatics more consonant with descriptionist theories of concepts.

Millikan's proposal that substance concepts point to their denotations applies the rigid designation theory of word meaning (Kripke 1980; Putnam 1975) to concepts: just as this theory rejects descriptionism about meaning, so Millikan rejects descriptionism about concepts. However, as Kripke acknowledges, rigid designation faces difficulties that descriptionism successfully meets; and we argue that such difficulties are all the more pressing for nondescriptionist accounts of concepts.

Rigid designation rests on essentialist intuitions about ordinary language use in counterfactual scenarios (Bealer 1987). Putnam suggests that if we discovered that all cats were robots, they would still be denoted by "cat" and called "cats" in ordinary language use, hence "meanings' just ain't in the head!" (Putnam 1975: p. 227); natural kind terms serve to point at their denotations rather than describe them.

Millikan cites empirical evidence that appears to support essentialist intuitions. However, this evidence does not examine the critical kinds of counterfactual scenario used by Kripke and Putnam. Braisby et al. (1996), investigating exactly these kinds of scenario, show that ordinary language use does not conform to essentialism. Two conclusions follow. First, the essentialist motivation for rigid designation is compromised, and hence Millikan's assumption of rigid designation needs more support. Second, the pattern of language use obtained points toward a descriptionist view. According to such theory, the perspectival view (Franks & Braisby 1997), concepts are descriptive but their content varies according to contextual and pragmatic constraints, including actions involving objects that concepts classify. Hence, we share Millikan's contention that concepts are implicated in action, though we suggest that this does not rule out descriptionism.

Even if rigid designation could be motivated without essentialism, however, Millikan's view of concepts faces challenges analogous to those facing a theory of meaning based on rigid designation.

Rigid designation fails to explain well-known failures of Leibniz's Law. If the meaning of "the morning star" is the object "pointed" to, then this is identical to the meaning of "the evening star"; however, it would then follow that "Venus is the morning star" means the same as "Venus is the evening star." The same goes for any sentences obtained by substitution of coreferential expressions. Frege (1892/1966) explained the differences in meaning between such sentences as differences in senses (characterizable in terms of descriptions).

The challenge for Millikan is to account for such differences without recourse to the senses/descriptions difference. If the concepts "evening star" and "morning star" point to the same substance, and this pointing exhausts the content of the concepts, the thought that "Venus is the evening star" has the same content as "Venus is the morning star." Millikan must demonstrate that these thoughts do not differ in content, and that any apparent differences reside in associated recognition or identification capacities.

Rigid designation also encounters difficulties with empty names (those without reference). Millikan suggests that "a substance concept causally originates from the substance that it denotes" (sect. 5, para. 14). Millikan's first problem is to explain how such concepts can figure in the truth-conditions of thoughts. Her second problem is to explain how there can be recognition capacities for such concepts when there is nothing to be recognised. Either there can be no empty concepts, or causally originating from their denotation is not a defining feature of substance concepts.

Finally, Millikan suggests how substance concepts may be causally linked to their denotations in the absence of perceptual acquaintance, by construing language in terms of direct perception. This raises the prospect of avoiding intentional terms in explaining reference. Such a view requires elaboration. Following rigid designation, reference might be supported by a causal-historical chain connecting language users to an initial baptism of a substance. However, such a chain depends upon people intending to use words to refer to the same kind as those to which others who precede them in the chain had referred. On this view, Millikan would require the "tracking" abilities of children to be sensitive to changes in referring intentions. The direct perception analogy would then break down, however, because referring intentions are not transparent in ordinary language.

Even if Millikan adopted an alternative construal of linguistic direct perception, the account may still presuppose the intentional. If pointing explained conceptual content, and could be explicated via direct perception, then pointing would be nonintentional. However, ostension (and sorting) appears to be irreducibly intentional (Fodor 1994); there is nothing in the act of pointing *per se* that is sufficient to indicate the referent. Similarly, what expressions refer to is not determined by the expression alone but is mediated by pragmatic constraints. Such mediation is more consonant with descriptionism and may depend on referring intentions (Kripke 1977). Such intentions may or may not accord with the conventional content of the description. The perspectival view of concepts aims to respect the inescapable intentionality of concepts (Franks & Braisby 1997): in determining a referent, it is not just pointing that matters, but the point of pointing.

In sum, Millikan assumes that the Kripke-Putnam position is unproblematic, but it faces difficulties that descriptionism overcomes; only when nondescriptionism addresses these difficulties will it be clear whether or not it is more viable than descriptionism.

Are there wordlike concepts too?

Christopher Gauker

*Department of Philosophy, University of Cincinnati, Cincinnati, OH 45221-0374. christopher.gauker@uc.edu
ucaswww.mcm.uc.edu/philosophy/gauker/*

Abstract: Millikan proposes that there are mapping functions through which spoken sentences represent reality. Such mappings seem to depend on thoughts that words express and on concepts as components of such thoughts, but such concepts would conflict with Millikan's other claims about concepts and language.

One starting point for a theory of concepts is that a concept is related to a whole thought as a word is related to a whole sentence. One reason to believe that concepts of this kind exist is that it

would make an important kind of mental process analogous to verbal discourse. Another reason is that the primary function of (declarative) sentences is to *express* thoughts and this too, would be a structural parallel between a sentence and a thought. From this approach, one might think that the content of a concept consists in part of its relations to other concepts, as the meaning of a word consists in part of its relations to other words. And one might think that some concepts are structured as descriptions, for a description can be a component of a sentence in the same way a word can be. However, one would not say that concepts are action schemata or that concepts are structures composed of theories, because such things cannot go together to form thoughts in the way words go together to form sentences. For the same reason, one would not say, with Millikan, that concepts are abilities to reidentify.

Will Millikan allow that in addition to concepts-as-abilities-to-reidentify there are the sorts of concepts that words might be thought to express, such as *horse*, *magnetic*, and *quickly*? Her remarks in the final section pull in that direction. What a speaker does, she tells us, is to pick up information-bearing patterns from the environment, “translate” them into language, and, in that form, “beam them” at us. This evokes a picture in which information passes from the world into a mind, is translated into speech and, finally, is translated back in the mind of a hearer. Moreover, words represent reality, she says, because there are “mapping functions” that map sentences into reality. As Millikan herself has stressed (1993a, p. 78), one cannot define a mapping function merely as any isomorphism between structures of sentences and structures of reality, because there are far too many such isomorphisms. One has to identify the right isomorphism. Those who wish to construe sentences as maps of reality usually suppose that the mapping is first of all a relation between thoughts and reality and pertains to spoken sentences only insofar as sentences express thoughts. In that case, thoughts must have some kind of structure that allows sentences to express them, and so there will be concepts that stand to thoughts as words stand to sentences.

It seems, however, that Millikan must reject any such rationale for positing wordlike concepts in addition to concepts used in her sense. A person can be said to have a concept, in her sense, by virtue of having a word, regardless of whether that person is skilled at making identifications with it, and there is no natural line between “knowing the meaning” and not knowing it. Indeed, “the very content of one’s *thought* sometimes passes through the word usages of a surrounding language community” (sect. 6, para. 8). But this idea should seem problematic to anyone who holds that sentences map onto reality only by virtue of expressing thoughts and that this is why there must be concepts that words express. What a person thinks may depend on what the person is told and in this way may depend on the way words are used; but, on the mapping theory, it should not turn out that the very content of a person’s thought depends on aspects of usage that that person has not personally detected.

A rejection of wordlike concepts seems to emerge even from Millikan’s explanation of the realism of our representations. Before the mind can classify, it must posit the thing to be classified. This ability to posit is equally an ability to reidentify. In light of this Millikan explains how our concepts can be about kinds not constituted by our classificatory dispositions. The explanation is that successful reidentification depends on objective facts and the cohesion of the pertinent kinds rests on those same facts. Someone might claim that the realism of our representations lies equally in our classifications, for success in classification likewise rests on what sorts of things really exist. Millikan would deny this because the reidentifications are more fundamental: so much so that the classifications themselves are perhaps best treated as methods of reidentification. Indeed, the ability to use a classificatory word is itself largely an ability to reidentify. In this way we can understand the use of classificatory words without having to posit the sorts of concepts that they might express.

In characterizing words as mapping reality, Millikan has created a need for a theory of concepts that she has not provided and that may not be feasible given her other right-minded commitments. She needs to postulate concepts that stand to thoughts as words stand to sentences so that words can map onto reality by expressing thoughts that map onto reality. But it is hard to see how she can accept concepts in this sense, given her commitment to thought content as something that is relative to prevailing linguistic usage and her explanation of the realism of our representations.

Some might counsel Millikan to go ahead and posit wordlike concepts as well and to embrace the conception of language as merely a tool for expressing thoughts. I hope for just the opposite. If, as Millikan claims, we can reidentify individuals and substances without classifying them, perhaps we can also compare them without classifying them. Perhaps we can even compare whole situations without classifying them. By thus comparing the situations in which words are used to other situations, we might even acquire a use for language in bringing about cooperation in the pursuit of collective goals. A theory of language along these lines will have no need for Millikan’s mapping functions.

Why language is not a “direct medium”

Tamar Szabó Gendler

Department of Philosophy, Syracuse University, Syracuse, NY 13244.
tgendler@syr.edu

Abstract: Millikan contrasts her substance-based view of concepts with “descriptionism” according to which description determines what falls under a concept. Focusing on her discussion of the role of language in the acquisition of concepts, I argue that descriptions cannot be separated from perception in the ways Millikan’s view requires.

Prologue. In the opening pages of Dickens’s *Hard Times*, the imperious schoolmaster Thomas Gradgrind is depicted quizzing his young charges on the subject of horses. Turning to one of the students, he bellows: “Give me your definition of a horse.” The girl remains silent, alarmed by the question. Piqued, the schoolmaster announces: “Girl number twenty unable to define a horse! Girl number twenty possessed of no facts, in reference to one of the commonest of animals.”

He next turns to one of her classmates, in an effort to obtain the answer he seeks. “Bitzer,” says Gradgrind, “Your definition of a horse.” And Bitzer replies: “Quadruped. Gramnivorous. Forty teeth, namely twenty-four grinders, four eye-teeth, and twelve incisive. Sheds coat in the spring; in marshy countries, sheds hoofs, too. Hoofs hard, but requiring to be shod with iron. Age known by marks in its mouth . . .” (and much more).

Pleased with this answer, Gradgrind announces triumphantly: “Now girl number twenty . . . you know what a horse is” (Dickens 1854/1955).

The Moral. Like Dickens, Millikan seeks to bring out the absurdity of the view that a child who is unable to define a term is possessed of no facts concerning the referent of that term. And like Dickens, she wishes to emphasize the absurdity of its counterpart – that in hearing the definition of a term, the child comes to know what the object, referred to by the term, is.

In this Millikan and Dickens are surely right. Whatever is required for having the concept “horse,” it is not the capacity to specify necessary and sufficient conditions for being one; nor is it the capacity to produce a list of characteristics (forty teeth, sheds coat) that (most) horses share.

But although these requirements for concept attribution are surely too strong, I am unconvinced that the proper alternative is the one Millikan proposes. My worries about her antidescriptionist project can be traced to a single source: Millikan underestimates the importance of what is often referred to as the “theory-ladenness of perception,” namely, that what we perceive in the world depends in crucial ways on the framework by which we

make sense of the world (cf., e.g., Gopnik & Wellman 1994; Kuhn 1962; Lakatos 1970; as well as Kant 1787/1929). My following comments focus on one specific point missing in Millikan's discussion (sect. 6) of the role of language in the acquisition of concepts, because its absence there is most problematic. In addition, I show how this has general implications for Millikan's view as a whole.

Why language is not transparent. Millikan argues that we should view speech "as a direct medium for the perception of objects in the same way that, say, light is" (sect. 6, para. 1). Although she notes certain asymmetries between perceptually acquired knowledge and knowledge acquired through the testimony of others, she neglects what I take to be the most important difference between the two modes of learning about the world, namely, that speech is produced only by intentional individuals, whereas ordinary perceptual media often arise nonintentionally. It is this, of course, that explains the relative fallibility of language. Indeed, Millikan herself recognizes this, for her prime example of perceptual illusion ("arranging" false appearances – by means of modern communications media) relies on precisely this feature (intentionality of production). The question is whether this difference in fallibility results in a difference in "immediacy," or, more generally, to what level of immediacy are our perceptions subject?

Millikan remarks that even when sentences are false, "they continue to present the same false appearances – they do not shift and appear to say something different" (sect. 6, para. 3). This is meant to be analogous to the way in which perceptual illusions – from the Müller-Lyer lines to Dan Rather on television – persist in *seeming* to tell us something about the way the world is, even when we know that it is not. And whereas Millikan cites data that suggest that language has a similarly direct route to belief (Gilbert 1993), experience tells against this as a general phenomenon. It seems false to say that the sentence "I am dying" presents the same "appearance" when uttered by a laughing 8-year-old as when uttered by a tearful cancer patient.

The reason for this is that part of what determines a sentence's "appearance" is the intention of the agent who produces that sentence. That a certain arrangement of words has a certain meaning depends upon a huge number of social and nonsocial facts, among them facts about the beliefs and intentions of the sentence's utterer. But if the production of sentences depends on this kind of intentionality, then it should not be surprising that their interpretation does too. And because this process of interpretation is precisely what it is to understand a sentence, it follows that speech cannot be – in the sense Millikan requires – a "direct medium."

General implications. There is, indeed, reason to think that the ways in which theory (description) enters into our "perception" of speech are paralleled by a similar entry of theory (description) into our perception of the sensory realm. That is, we might accept Millikan's suggestion that speech and light are similar in giving us access to objects in the world, but instead of saying that speech is more like light than it initially seems to be, we might say that light is more like speech than it initially seems to be. In section 4, Millikan distinguishes her view from descriptionism, saying, "In contrast to the position just sketched, the descriptionist is one who holds that the referent or extension of a substance term is determined by its falling under a description associated with the term by the user" (para. 1). But the description under which a term falls may affect the perception of the entity to which the term refers. The analogy of optical illusions may be no more relevant to our seeing a mouse (as a mouse) than to our hearing a sentence (as a sentence).

Etiological classification and the acquisition and structure of knowledge

Michael T. Ghiselin

Center for the History and Philosophy of Science, California Academy of Sciences, Golden Gate Park, San Francisco, CA 94118-4599.
mghiselin@calacademy.org

Abstract: Millikan's account of how we acquire our most basic concepts might be clarified by a better ontological taxonomy, especially one that distinguishes between natural kinds on the one hand and wholes composed of parts on the other. The two have a different causal basis, which is important because once classification goes beyond the stage of naive induction, it becomes fundamentally etiological.

The terminology of Millikan's ontological systematics strikes me as a bit odd. According to her definition, World War II is a substance, whereas I would classify it as an individual (event) that falls under another ontological category, which I would call "process" (Ghiselin 1981; 1997). It is also not clear why she treats entities to which mass nouns refer as "stuffs" rather than "real kinds." Consider, for example, *cow* and *beef*. Although Millikan acknowledges the point, now generally accepted in the philosophy of taxonomy, that biological species are individuals, she still tries to treat both these and social individuals as natural kinds, with members rather than components. There is indeed a deep similarity between natural kinds and individuals, including the sharing of many properties by the elements. Thus, for example, the parts of the University of Connecticut have much in common, the parts of the Department of Philosophy thereof have still more. But the University of Connecticut is not a kind of department nor is the Department of Philosophy a kind of professor. Neither is a kind of anything. They are wholes composed of parts.

I emphasize the distinction between the two kinds of groups because Millikan, although making the distinction herself, does not make it strongly enough, and, more important, because the tradition has been to conflate both groups under the rubric of natural kinds. The taxa of systematic biology, like the subunits of universities, are not kinds at all. It is intuitively obvious, of course, as well as downright wrong, that *Homo sapiens* is a synonym for *human being*. But that would make a sperm cell or a zygote a human being. To help alleviate such confusion I have introduced the term "componential sortal" for such words (Ghiselin 1997, pp. 64-65). A human being is an organism-level component, not an instance, of our species. The same is true of a kitty and its species, a point to which we shall return. For the moment it should be clear that conflating the two kinds of groups under "natural kind" conceals the fundamentally different causes of the properties held in common: laws of nature on the one hand and history on the other.

Another common mistake has been to assume that entities are related because they are similar, rather than similar because they are related. Such a notion is commonly reinforced by a naive inductionist model of knowledge acquisition. Human beings are perfectly capable of dealing with groups as if all they shared were perceived properties. But it is misleading to assume that such a primitive manner of thinking is somehow obligatory, and that we cannot do otherwise. Even scientists indulge in such phenomenal classification, but when they discover the underlying causes of natural order, they switch to etiological classification. Genealogical classification in systematic biology is a compelling instance. Such a shift often necessitates extensive intellectual restructuring, as in realizing that birds are dinosaurs.

It does seem likely that we come equipped with the capacity to identify individuals. Indeed, the ability of cells within an organism to tell self from nonself is fundamental to immune responses. Ontological judgments may come to us quite easily, even if we are very young: if we see a pair of identical twins at the same time, their similarity is utterly irrelevant to their being distinct individuals. Such connections may be deeper than what Millikan has in mind. Human beings are able to understand the realities of

procreation over and above the appearances of family resemblance and nutrition.

The notion that we may classify on the basis of a single instance suggests the sort of ostensive definition with respect to the naming of species that is laid out in the International Code of Zoological Nomenclature. A name is bestowed upon the group by designating a museum specimen as a nomenclatural type, and no other criterion fixes the reference. On the other hand, a great deal of background information is inevitably presupposed when people deal with terms and concepts. Zoologists know that species often display sexual dimorphism. Likewise, in the learning of language there is a tacit assumption that there already exists a language to be learned. For this and for many other reasons, treating how children acquire language as an example of knowledge in general can be seriously misleading.

We might reconsider the example of a child who calls a tiger a kitty. If his mother were a professional systematic zoologist, she would no doubt be delighted with his ability to place the organism in question in the correct position in the Linnean hierarchy. According to my dictionary, "kitty" is a playful synonym for "cat," and "cat" in the broad sense means "big cat" as well as "house cat." Even in the nontechnical language of adult laymen, "cat" is approximately equivalent to an organism-level component of the family Felidae, and not just of the species *Felis catus* Linnaeus, 1758. The confusion is in the mind of the philosopher, not of the child. Philosophers have a long history of dealing with tigers, alas, the ones that are supposed to inhabit Mars. So coming down to earth and discussing real tigers is a distinct breakthrough, even though the treatment is not altogether satisfactory.

What can externalism do for psychologists?

Alison Gopnik

Department of Psychology, University of California at Berkeley, Berkeley, CA 94720. gopnik@cmsa.berkeley.edu

Abstract: I suggest several ways that externalism could influence psychological theorizing. Externalism could just capture our everyday intuitions about concepts and meanings. More profoundly, it could enter into psychology through evolutionary theory, guide our hypotheses about conceptual abilities, and, most significantly, it could influence our accounts of learning and conceptual change.

What could Millikan's interesting new view of concepts do for psychologists? It does not seem necessary to adopt Millikan's strong externalist view to account for the intuitions that Kripke (1980) and Putnam (1975) originally appealed to. The kind of internalist "psychological essentialism" advocated by Keil (1989) and Gelman and Wellman (1991) can do that. The intuitive philosophical evidence is fully compatible with the possibility that we are psychologically constituted to understand the mind in an externalist way, but that the best scientific account of the mind really is internalist. The contribution of "psychological essentialism" is to recognize that the fact that we are externalists in our everyday life is important and explains much of our behavior and language. Externalist accounts have made a real contribution to an internalist psychology in this way.

Millikan, however, wants to argue not only that externalism is our intuitive view of concepts but that it is the true view of concepts. In particular, it is the view of concepts that an empirical psychology ought to adopt. On this view certain theoretical concepts in psychology (such as the notion of "concept" itself) are inherently relational; they inherently include information about the structure of the outside world as well as about our own internal states. Our understanding of vision inherently involves information about the physical properties of the world, and it does so because our best way of understanding how vision works is in terms of its function of veridically representing external objects. If the conceptual system is similar, then trying to construct an

internalist account of concepts is like trying to do psychophysics without the physics.

How could this picture be cashed out in an actual psychological theory? One possibility is that the external world only enters into the theory at the point in the distant past where evolution shaped the conceptual "tracking" abilities. However, these evolutionary facts might only enter very weakly into psychological explanations. After all, in some sense all psychological explanation will ultimately find its grounding in evolutionary facts, and so all psychological explanation will ultimately be externalist. In practice, though, this rarely influences our psychological theorizing in a deep way.

A more significant type of externalism might come if the analogy to perception is taken more seriously. If conceptual structures are supposed to track external kinds, and if we know something about the nature of those kinds, then that knowledge might inform our hypotheses about the tracking abilities that constitute conceptual structure, just as knowing physics and optics informs our hypotheses about the perceptual system. Here the inclusiveness of Millikan's account is a bit problematic. The kind of knowledge we will need about the actual nature of individuals, substances, natural kinds, and artifacts will be radically different, even if, at a broader level, all of these things are being tracked. For example, it is possible that certain patterns of correlation among properties are associated with particular patterns of underlying causal structure, and this might explain why we track natural kinds in certain ways. But this would not apply to our ability to track artifacts or individuals, which will be united in very different ways. Nevertheless, it might be promising to look for the characteristics of different types of real kinds that would, in fact, provide us with useful inductive purchase on those kinds, and to search for relevant cognitive mechanisms. The project might be like "ideal observer" theories of perception, where the best way of capturing visual information is worked out, and then we can compare that ideal to actual human capacities.

However, there are some significant dissimilarities between cognition and perception. Our perceptual mechanisms seem to change relatively little with development. As a result we can specify the relations between the real structure of objects and our visual algorithms with some authority. In contrast, our conceptual tracking mechanisms will always be bound up with our particular changing ontological convictions. Millikan is right in saying that the tracking mechanisms should not be identified with our ontology, that they may even predate such beliefs. Nevertheless, there seems to be a kind of bootstrapping. Early tracking mechanisms that are based on movement and trajectory will eventually be replaced by others, particularly the linguistic mechanisms Millikan discusses in her last section. By the time children reach preschool age, their inductive inferences will almost entirely be determined by their particular beliefs about the ontological structure of the world. If they believe that animal kinds are best indicated by their inheritance factors, that will be their concept of animals and will explain their use of animal terms. It is an open empirical question whether more basic evolutionarily determined tracking mechanisms are still in place then, or indeed in adulthood, or whether their function has been taken over entirely by particular beliefs; what we call theories. If the early mechanisms have been overturned, then it seems that we could indeed know all there is to know about concepts by specifying the particular beliefs subjects hold. What role will be left for externalism in psychological explanation?

Millikan, like Putnam and Kripke, suggests that the meaning of a particular word or the content of a particular concept is what relates that concept at that time to a specific aspect of the world. Notice that this is not the way that external objects enter into psychological theorizing in perception. Instead the relation is much more indirect; the external world plays a crucial role in fixing the mechanisms that allow us to form certain representations rather than others, given particular inputs. Similarly, the crucial role of the external world in conceptual development is not likely

to take the form of individual relations between reified “concepts” and real kinds.

We can expect, however, that the external objects will play a crucial role in learning in the very bootstrapping processes I just described. Somehow our interactions with the external world lead us to form one set of ontological convictions, one theory, rather than another. Even if knowing our theories is all we need to explain our current psychological functioning, we may well need to invoke relations to the external world to explain why we have those theories rather than others. These learning processes have been almost completely neglected in cognitive psychology. If Millikan’s externalism forces us to solve the problem of how our interactions with the world lead to conceptual change, she will have provided an important service to psychology as well as philosophy.

Staying in touch: Externalism needs descriptions

James A. Hampton

Department of Psychology, City University, Northampton Square, London EC1V 0HB, UK. j.a.hampton@city.ac.uk

Abstract: Externalism cannot work as a theory of concepts without explaining how we reidentify substances as being of the same kind. Yet this process implies just the level of descriptive content to which externalism seeks to deny a role in conceptual content.

The recent antidescriptionist, externalist program for concepts has impressed few empirical psychologists, as might be expected given that it takes the most interesting question – Why do we have the concepts that we do? – and seeks to provide an answer that has little use for (or of) psychological inquiry. The question is what truly constitutes or defines the content of a concept. Descriptionists look to answer the question by identifying the collected set of beliefs that a person has gathered about a class or type of thing. Thus, to a descriptionist the concept “milk” is the set of property descriptions which tend (generally) to hold true of a particular class of stuff. According to which descriptionist theory you espouse, these descriptions may involve a greater or lesser degree of theoretical elaboration or stored experience. They may also be partial or incomplete.

The problem with this descriptionist view, as identified by Millikan and others, appears to be that it puts the cart before the horse. How are we able to accumulate a set of beliefs about a substance concept, if we cannot first have a reasonably reliable means of knowing that a new experience is of the *same kind* as an earlier one? How can we add a new belief to the set without independently first identifying what our belief is about? To Millikan the ability to reidentify something as an instance of the same class is logically prior to the acquisition of descriptive properties that may be generally true of the class. Conceptual content is constituted by the real nature of the class that gives rise to these (fallible) acts of reidentification. All our internal concepts do is point to real classes in nature.

The problem of reidentification is clearly central to Millikan’s argument. Yet how exactly can the process occur without involving some form of descriptive representational format? As Millikan points out, the ways in which we reidentify Mama, a cat, or a glass of milk will all depend on different aspects of the perceptual and sensorimotor information presented to us. Yet how are we to know anything at all about a newly experienced object without transducing its physical characteristics through our sensory organs and perceptual apparatus, thereby deriving a mental representation of it? There can surely be no identification or reidentification without a concomitant creation of a mental description.

Millikan answers this question by stating that such a description need not imply property concepts since “the thought of a property is not just a reaction caused by a property; it must play an

appropriate representational role” (sect. 5, para. 7). She proposes that descriptions may be used in reidentification, but that they should not be representable in thought. It is puzzling how this distinction between the “aspects” of objects used to identify them and true “property concepts” that can also be represented in thought is meant to be cashed out. Perhaps being representable in thought means that a property can be isolated by attending to it as a separable dimension (i.e., making judgments about it independently of other dimensions), as well as being able to label it with a language term.

Yet why should descriptionist accounts of concepts be limited in their descriptive vocabulary to such a limited range of descriptive power? It should be clear that there are many crucial aspects of conceptual knowledge that fail this test. The way that we represent the shape of “cat,” or the taste of “milk,” or the face, voice, and smell of “Mama” are most unlikely to involve independently isolated features or dimensions representable in thought. But why should we not consider these concepts to involve descriptive information, and why should we suppose that we perform the recognition independently of the descriptive information we have already stored with the concept?

A descriptionist account of concepts could not possibly get off the ground if its representational power did not include the ability to store and represent this kind of imagistic information. We are clearly capable of representing a wide range of information that could not be expressed in language. Verbal concepts must be grounded in experience at some level. Given a realistic version of descriptionist concept representation, it is hard to make sense of Millikan’s distinction between unanalysed reidentification aspects and “concepts of properties.” Millikan’s argument relies on a notion of *description* that is so restrictive as to rule out any account of concepts other than those nominal concepts with explicit definitions such as *prime number*.

The target article also fails to address a major critique of externalist theories, which should be familiar to all followers of this debate. It is just not true that all of our concepts are tracking real classes in the world. We may concede that natural kinds have an independent existence as classes, but the majority of concepts involved in everyday thought are not natural kinds. Take some of the political and moral issues that have divided society in recent years. Externalism seems to require that in each and every case there be a real and objective answer as to which side is correct; there is an external “good,” which we attempt to track through our fallible concepts. This is an extreme position to be forced into. Other concepts are defined relative to the possessor. Thus “home town” and “favorite food” are defined relative to the person whose home town or favorite food it is. For these concepts it is clearly false to say that each concept possessor is tracking the same externally real class. Yet how could the concepts then be defined except through their descriptions, as they exist within that individual?

In summary, whatever the merits of the “pointing” view of concepts as an account of very young children’s first attempts to construct concepts, any psychologically adequate theory of the full range of adult concepts will need to incorporate a strong descriptionist component.

Reidentification and redescription

Marc D. Hauser^a and W. Tecumseh Fitch^b

^aDepartments of Anthropology and Psychology, Harvard University, Cambridge, MA 02138. hauser@wjh.harvard.edu
wjh.harvard.edu/~mnkylab

^bSpeech and Hearing Sciences Program, Harvard-MIT, Cambridge, MA 02138. tec@wjh.harvard.edu wjh.harvard.edu/~mnkylab

Abstract: Millikan’s account of substance concepts fails to do away with features. Her approach simply moves the suite of relevant features into an encapsulated module. The crux of the problem for scientists studying human infants and nonhuman animals is to determine how individuals reidentify objects and events in the world.

The difference between a descriptionist's and Millikan's account is one of focus, not of kind. Although the idea that concepts point to extensions determined by some objective reality of the real kind represented is philosophically interesting, the practically important problem is to understand how an organism can (more or less accurately) reidentify instances of a kind. Millikan posits that there are mechanisms underlying perceptual constancy for properties and that these are not specifically concepts of properties. This move simply buries the description (a list of properties) in a cognitively impenetrable module but does not do away with the description all together. Further, in both Millikan's and the descriptionist's view, the information necessary to reidentify a substance must be stored somewhere, and even if the means used to accomplish this are nonverbal, fuzzy, fallible, and domain-general, they can still be generally thought of as descriptions.

Consider an extremely simple case: how a living ant identifies a dead ant. Wilson (1970) has demonstrated that when an ant dies, its body is carried out of the nest and deposited at some distance. How is a dead ant identified? Simple. When ants die, they emit a chemical known as oleic acid. Take some oleic acid and dump it on a live ant, and it, too, will be dragged away. On Millikan's substance view, a dead ant is reidentified by a simple, perceptually invariant mechanism: the detection of oleic acid. Here, the substance view boils down to a feature, a single perceptually invariant feature that provides means of identification that is typically infallible (i.e., when devious experimenters are not about). In the ant's system, the "concept" of dead ant is "described" by firing of oleic acid receptors, although it seems quite doubtful that the ant "knows" this.

To take a more complex example, how does a monkey mother reidentify her offspring? Like most other vertebrates, monkeys are able to recruit a variety of sensory modalities, including olfaction, vision, and audition. The accuracy of identification in any one modality will depend upon (1) species-typical design features (e.g., how much cortical space is dedicated to processing in a particular modality), (2) environmental factors that alter the quality of the incoming signal, and (3) the experiences that the mother has had of the infant in various modalities. In some species, such as the vervet monkey, infants produce vocalizations that appear to include an "acoustic signature" revealing the caller's identity. Thus, when an infant's call is played back through a speaker to its mother and two other adult females, the mother looks toward the speaker and the two females look toward the mother – mom recognizes the call as "her infant" and the other females recognize the call as "that mother's infant" (Cheney & Seyfarth 1980).

There is substantial variation in the acoustic morphology of the infant's different calls (e.g., as a result of changes in its affective state), and these may occasionally lead to misidentification. Nonetheless, some portion or aspect of the calls is sufficient to guide reidentification and appropriate action (e.g., orientation and search). Although the precise acoustic characteristics that underlie this capability are poorly understood, let's assume for now that a characteristic warbling in all of Joey's various calls gives them away as "Joey's" (in fact, the processes underlying individual recognition are probably significantly more complex than this, but this is not critical for this point). Now, a mother monkey, certainly cannot tell us that "the warbling means Joey," nor can we play back just the warble and expect her to respond appropriately (because the warble is part and parcel of the entire call). Thus, the features responsible for reidentification need not be abstracted away from the stimulus array they are part of, nor do they need to be available to conscious introspection (whatever this would mean for a human infant or a monkey). Nonetheless, some part of the identifier's brain *must* have the warbling stored away and linked to the concept "Joey." To us, this stored-away information constitutes a description, in Millikan's terms, of the substance "Joey."

It is interesting that if an infant's call is played back a day or two after its death, the mother fails to respond – mom appears to be treating the call as "not my infant's" (Allen & Hauser 1991). Thus, unlike the ant case, reidentification in monkeys involves more than a single feature, and there is plasticity with respect to how

particular features are integrated into a system for guiding action. The infant's call is not an automatic trigger that guides maternal action. It is contingent upon the mother's knowledge of the infant's status – living or dead (and some other factors as well). How this sort of knowledge interacts with the more perceptual information used to identify a particular call is an interesting and open question, and we look forward to Millikan's thoughts.

One final challenge. Millikan argues that "language interacts with substance concepts, completely transforming the conceptual repertoire." But rather than articulating the transformation, what she offers instead is a simple verbal proposal for the way language provides a vehicle for concept exchange in the absence of direct perceptual experience. Therefore, although it seems clear that language enables us to build up our conceptual repertoire, it is unclear how, on Millikan's account, the rudimentary "substance" concept builds up into a more explicit concept of kinds – natural, artifactual, and so on. This is a critical issue, and many researchers working on infant cognitive development and animal cognition have a vested interest in the outcome of the debate (Cheney & Seyfarth 1990; Hauser & Carey, in press; Karmiloff-Smith 1992; Uller 1997; Uller et al., in press; Xu & Carey, 1996). [See also BBS multiple book reviews of Cheney & Seyfarth's "How Monkeys See the World" *BBS* 1992 15(1) and of Karmiloff-Smith's "Beyond Modularity" *BBS* 1994 17(4).] Thus, for some, language interacts in significant ways with our numerical concepts as well as our folk psychological concepts. But here, the evidence is generally correlational: a capacity for a nonlinguistic concept X comes in during child development, along with the acquisition of linguistic concepts. What is now required is that theoreticians develop an account of how language might transform conceptual foundations and that experimentalists design tests that clearly reveal the direction of causality. So much to do!

The most basic units of thought do more, and less, than point

Frank Keil

Department of Psychology, Cornell University, Ithaca, NY 14853.
fckl@cornell.edu

Abstract: Thinking of concepts as explicit lists of features used to pick out referents neatly is indeed mistaken; but there are other alternatives than making concepts mere pointers. These alternatives are suggested by the difference between meaning X and having the concept X, problems of conceptual change, implicit conceptual schemata, the conceptual requirements of the division of cognitive labor, and how concepts figure in perception versus language.

Millikan properly takes psychologists to task for not addressing how the Putnam/Kripke causal theory of reference seems to undermine descriptivist views of concepts. Her eloquent and provocative target article makes all of us rethink what concepts are; but some lingering issues suggest a different conclusion.

First, "having the concept X" seems different from "meaning X." I might learn to refer to "pyrosomes" by misunderstanding someone as saying that pyrosomes are often found in "grease" (instead of Greece). In such a case I would not know whether pyrosomes are instruments, vitamins, or sea animals, but I could still refer to them. When I say "pyrosome," I engage in a communicative act relying on causal links between pyrosomes and their name. I can "mean" pyrosome in Putnam's sense with totally incorrect knowledge of pyrosomes. What we mean to refer to can be something completely different from what we think we are referring to. For Millikan, I have a concept of pyrosome in these cases, but acts of reference may not always constitute having a concept. I can also have robust concepts without meaning to refer, even for nonlogical/mathematical concepts. I can learn about an "ideal gas," a substance concept that I know does not and cannot refer in any possible world. Thus, we can refer and intend to

convey a meaning without having a concept at all; and we can have a concept without ever intending to refer. This double dissociation of concepts from referents suggests that concepts do more, and less, than point.

Second, if concepts are pointers without content, conceptual development and change seem to evaporate. I can learn more and more facts about substances, but for Millikan, such facts are orthogonal to the concept. Increasing substance tracking skills do count as conceptual change, but why tracking and not facts? If an infant initially tracks dogs through barking, but later also uses furriness, why is that conceptual change when learning that dogs have fur is not? Ability to reidentify becomes reified as the new conceptual core for unclear reasons. Reidentification skills are said to be first "motor" and "perceptual," and "surely the bottom layer on which conceptions of substances are built" (sect. 5, para. 3). Why? If young infants reidentify things acting with reciprocity and at a distance, what makes that perceptual or motor? Indeed, Millikan later grants innate conceptual tracking abilities (sect. 5, para. 9), as long as they are not mediated by conscious memory for discrete properties. Having explicit awareness of properties as central to a kind is not necessary for having a concept, and Millikan correctly criticizes us for sometimes assuming so in our studies; but, there are many states between (1) explicitly knowing what properties are central and (2) not having properties that play any role in concept structure. An explicit understanding of the ontological principles that ground substance tracking may not be needed for a substance concept; but that does not make concepts representationally empty, or essentially the same early on across all kinds and individuals. Either conceptual change is illusory and occurs outside the concept proper, or, if changing tracking skills do constitute conceptual change, concepts become more than mere pointers and require richly structured skills and implicit cognitive schemas.

Third, it follows that explicit properties and pointing do not exhaust concept structure. I might understand bacteria in terms of their causal powers and patternings and not in terms of any specific features. Bacteria afford certain kinds of interactions and are adapted to certain niches. They occupy a causal nexus in a particular way that I can partially and distinctively apprehend but may not be able to state in propositional terms. Similarly, my concept of a bounded solid object may center on knowing the causal behavior exemplified in billiards without knowing a property list. When people see intentional agents in triangles and squares moving a certain way in a filmstrip, they are not consciously invoking feature sets, yet they are apprehending a patterning distinctive to social agents. As a realist who believes that humans and other animals have developed a diverse array of concepts so as to better pick up on and think about the heterogeneous cluster patterns in the world, I am uncomfortable limiting the conceptual repertoire to either property lists or simple pointers. Moreover, when concepts are considered as tracking skills, that account may end up allowing almost anything in concept structure, except explicit property knowledge.

Fourth, the division of linguistic labor and the implied division of cognitive labor does not free the laborers of all conceptual knowledge outside their area. Having the concept of "gold" is more than knowing that experts can reidentify gold much better than I. It is having skeletal knowledge of the terrain of expertise itself, to know what chemists are and do, what natural patterns they can expertly navigate and why. My concept of gold is situated within that larger sense of the functional anatomy of knowledge in a community and how it connects to reality.

Finally, language does not transmit structural isomorphisms and invariants the way perception does. Language transmits information about substances, but not as in perception. Even television preserves a structural isomorphism in the visual array in ways that language cannot. There is nothing in the sentence "dormice waddle" that maps informationally onto dormice waddling, but televised images do. Conversely, language puts dormice directly into propositional structures in ways that vision cannot.

Concepts are both units of thought that can be generatively combined into propositional structures and things that allow us to pick up on and respond selectively to chunks of causal patternings in the world. For me, "concept" refers not to ways of tracking entities over space and time, but to those basic units of thought that allow us to grasp reality and then mentally manipulate and explore the consequences of what we grasp. "Knowing how to use a thing" is indeed "not knowing facts about its structure" (sect. 2, para. 15); but knowing how to use a thing may well be knowing, in a different conceptual sense, much about its structure and the causal patternings in which it is embedded.

Mapping Millikan's conceptual work onto (empirical) work by psychologists

Lloyd K. Komatsu

Department of Psychology, Carleton College, Northfield, MN 55057.

lkomatsu@carleton.edu

Abstract: There are three points of difference between psychologists' assumptions and those that Millikan suggests: (1) concepts as representations versus concepts as reflecting a capacity; (2) concepts having a role in categorization and inference versus a role in reidentification; and (3) the "basic level" as an aspect of the "vertical" dimension of categories versus being a kind of category, on a par with natural kinds.

Millikan's target article promises to be the start of a lengthy conversation between philosophers and psychologists embodying the productive disciplinary cross-fertilization that characterizes cognitive science at its best. But for psychologists to benefit from Millikan's insights, we need to understand how her ideas map onto ideas common in psychology.

Psychologists generally assume, for example, that a concept is a representation (e.g., Smith & Medin 1981). In contrast, Millikan argues that we should assume that a concept is an ability, a capacity to represent. Similar shifts in emphasis from representation or structure to process have had a significant effect in other areas of psychology (e.g., in memory and attention). It is therefore very important to follow up this suggestion. But it is not immediately clear to me how to go about doing that: how such a shift might be realized in this area, or what the implications of such a shift might be.

Millikan suggests that substance kinds can be distinguished along two dimensions: the number and the certainty of the inferences the categories of a kind support. This idea is familiar to psychologists and has been important in work contrasting natural, artifact, and nominal kinds (e.g., Keil 1989). However, psychologists have generally focused more attention on a different pair of dimensions along which categories may vary: the horizontal and the vertical (e.g., Rosch 1978).

On this view, the horizontal dimension distinguishes kinds (e.g., "dogs" rather than "cats" or "tables"), as well as varieties of kinds (e.g., artifact kinds rather than natural kinds). The vertical dimension distinguishes levels of inclusiveness (e.g., "dogs" rather than "animals" or "dachshunds"). As used by psychologists, the "basic level" refers to that point along the vertical dimension that strikes the best compromise between having too few categories (i.e., sacrificing inductive richness) and having too many (i.e., sacrificing cognitive economy).

Millikan, on the other hand, regards the basic level as another variety of kind. Is this to suggest that she believes that the vertical/horizontal distinction is not important for an understanding of concepts, and/or that it can be replaced by the dimensions of number and certainty of inferences supported? Although it is clear that variation in number of inferences supported marks differences in both the horizontal and vertical dimensions, it is less clear whether certainty of inferences operates along the vertical dimension. Could this be one of the consequences of concepts-as-capacity as opposed to concepts-as-representation?

Another difference between Millikan's ideas and those familiar to psychologists concerns the role(s) in which concepts are thought to play. Whereas psychologists generally understand concepts as being involved in categorization and inference, Millikan proposes that they are in the service of "reidentification." The popularity among psychologists of "mixed" models of concepts, in which categorization and inference are supported by different representations (e.g., Miller & Johnson-Laird 1976; Smith 1995) might suggest that Millikan's ideas and those of psychologists are not terribly far apart, were we to equate reidentification with categorization. But Millikan appears to reject equating them in that way. Categorization as psychologists discuss it lies on a continuum with inferences and is often a matter of conscious reflection. In contrast, Millikan implies that reidentification is cognitively impenetrable and clearly distinct from inferences.

The reidentification/categorization distinction reminds me of another distinction made by psychologists: that between repetition and associative priming. Like Millikan's reidentification, the operation of repetition priming is often implicit, not consciously available. And like reidentification, repetition priming does not rely solely on physical identity. The exact parameters that are relevant to repetition priming (and that distinguish it from associative priming) are difficult to articulate, much as it is difficult to articulate the parameters that drive reidentification (and that distinguish it from categorization). I look forward to hearing Millikan's thoughts on the relationship between reidentification and repetition priming.

Concept acquisition and use occurs in (real) context

Kenneth R. Livingston

*Department of Psychology and Program in Cognitive Science, Vassar College, Poughkeepsie, NY 12604-0479. livingst@vassar.edu
depts.vassar.edu/~psych/facultypages/livingston.html*

Abstract: A realist story of concepts like Millikan's can and should accommodate facts about how the context of items available for comparison during concept formation affects just what concept is formed or reidentified. Similarly, the contribution of the goals and purposes of the conceptualizer are relevant to how concepts are acquired and deployed, but can be understood as entirely consistent with a view of concepts as objectively evaluable.

It is refreshing to read an account of concepts that pays proper attention to their ontological roots in the real world beyond the organism. It is equally refreshing to see the psychological literature about concepts and concept acquisition interpreted as supporting an indexical account of what it is exactly that concepts do: they point, they do not describe. Pointing amounts to picking out real and stable patterns of similarity and difference in the world. Ontologically, these patterns or correlations are grounded in the causal structure of the world, including causes external to the kind, as in many artifact categories. Epistemologically, it is because there are real categories in this ontological sense that concepts continue to point to the same entities across changes in the conceptualizer's understanding of the specific properties, correlations, and causes that characterize the category. Psychologically, the correlated structure of real kinds is sufficiently rich that we can exploit it to form concepts even if the ontological glue that explains the correlations is unknown. It follows, Millikan says, that questions of identification and classification are questions about how the world is, not about how we humans prefer to identify or classify things.

I am persuaded that a realist story is the right *kind* of story to tell about concepts. What needs clarification is what it means for a realist to say that our human preferences are not relevant to the correctness of our classifications. This assertion could be misunderstood to be urging a view of the mind as diaphanous, which is

certainly not Millikan's position. On the other hand, denying that the mind is diaphanous does not require that one accept a Kantian view of concepts as *a priori*. What follows is a condensed version of an argument for how concepts constitute objective judgments about the real world but must nevertheless be understood as products of the human mind. The argument helps to explain important empirical findings not addressed in Millikan's target article, findings that have in common a focus on the importance of the context, both internal and external, within which concept formation and use occur.

External factors first. How a particular entity is indexed depends critically upon the context of other entities available for consideration by the concept learner. Tversky's (1977) classic work on similarity points in this direction. More recently Medin et al. (1993) have shown that how a dimension of variation is coded for purposes of establishing similarity depends on what portion of the dimension is available for examination. Livingston and Andrews (1995) have demonstrated that whether a particular dimension is salient for a classification may depend on how instances that vary on that dimension are presented to the concept learner. These results, among others, suggest that whether a given pair of items is identified as belonging to the same category depends upon the contrast set within which they are encountered. So, for example, given the set [Amos, Minnie, Sylvester] Amos and Sylvester are less likely to emerge as the natural grouping than they are when they occur in the set [Amos, Sylvester, Tweety].¹ How an item is indexed is affected by the comparison set within which it is encountered.

This is all still a matter of how the world is, not how humans would like it to be, but it does highlight the relevance of the focus of attention during concept acquisition. For human beings (at least), the direction of attention is not determined in any simple way (not after the first months of infancy) by what is available to the senses on a given encounter with the external world. The concepts we have acquired prior to that encounter, along with our goals and purposes at the time, also matter. The human capacity to control the direction of attention makes it possible to treat some members of a set to the exclusion of others, thereby changing the (psychological) context of potentially confusable alternatives and altering the relative salience and judged importance of various dimensions of variation (see, e.g., Livingston & Andrews 1995; Medin et al. 1997). These selections constitute hypotheses about which groupings will be the right ones for our purposes, and are tested in various ways against facts about the consequences of deploying the concept-in-the-making in pursuit of those purposes.

In the end, facts about how the world works determine which concepts are valid for a given purpose and which are not. This is the sense in which it is correct to say that what we prefer is not relevant to which concepts are valid. But for a given entity, how it is classified on a given occasion is in part a function of what the conceptualizer is up to at the moment. For example, consider again the set [Amos, Minnie, Sylvester]. If I am interested in ridding myself of creatures that nest in my attic, Amos and Minnie belong together, in contrast to Sylvester. But if I am a psychologist interested in whether behavior varies as a function of reproductive anatomy, Amos and Sylvester constitute the right grouping from this set. Both concepts are valid, and both are abstract in the same way – they ignore some dimensions of similarity and difference in favor of others (e.g., sex for species or vice versa).

There is a great deal else to be said about how the mental machinery for abstracting from particulars to concepts actually works (see, e.g., Goldstone 1994; Kelley 1984; Livingston et al., in press), including the role of language and culture in the process (cf. Malt 1995), but the basic thrust of such discussions can and should return to a realist framework. There is a real structure of similarity and difference in the world, explainable by reference to various causal and functional roles of entities and events, as Millikan has argued. But these similarities and differences do not become categories until a human (or other appropriately organized) mind forms a concept that picks out one or more facets of

that similarity structure in preference to others that are available. A correct theory of concepts requires an understanding of how the conceiving instrument and the world it conceives interact, a view that I think is entirely consistent with Millikan's account.

NOTES

1. For those deprived of the American cartoon experience, Amos and Minnie are mice, Sylvester is a cat, and Tweety is a bird.

Finding order in our world: The primacy of the concrete in neural representations and the role of invariance in substance reidentification

Bruce J. MacLennan

Department of Computer Science, University of Tennessee, Knoxville, TN 37996-1301. maclellenn@cs.utk.edu www.cs.utk.edu/~mclennan

Abstract: I discuss neuroscientific and phenomenological arguments in support of Millikan's thesis. I then consider invariance as a unifying theme in perceptual and conceptual tracking, and how invariants may be extracted from the environment. Finally, some wider implications of Millikan's nondescriptionist approach to language are presented, with specific application to color terms.

Because I am in substantial agreement with Millikan's thesis, this commentary will explore connections between her thesis and neural network theories of knowledge representation.

In his Heideggerian critique of traditional ("symbolic") cognitive science and artificial intelligence, Dreyfus (1979; 1982; 1991, pp. 115-21) pointed out the futility of trying to represent our skillful coping in the world in terms of atomic, abstract, context-free predicates. Even Husserl acknowledged the "huge concreteness" of this hypothetical abstract structure and called its phenomenological reduction an "infinite task" (see citations in Dreyfus 1982, p. 20). We can refer to this observation as *the primacy of the concrete*; that is, the world as ordinarily experienced is primarily concrete, historical, contextual, meaningful, and, in mathematical terms, effectively infinite dimensional. Conversely, analysis into abstract, context-free, objective, low-dimensional predicates is a comparatively rare activity that we undertake in extraordinary circumstances (e.g., during "breakdowns" in skillful coping, or during scientific analysis); such analyses are always and necessarily incomplete. In Millikan's terms (sect. 4), our understanding of the world is primarily nondescriptionist, proceeding mostly by reidentification of relevantly similar substances; abstract descriptions are ancillary.

The primacy of the concrete is also supported by developments in neuroscience. Although sensory systems are often explained in terms of abstract "feature detectors," this terminology is inaccurate in a number of respects. Certainly, to a first approximation, neurons in early sensory areas appear to be tuned to simple abstract properties: small segments of edges and lines, patches of color, tones, and so on. However, more detailed investigation reveals that most sensory neurons respond to complex combinations of stimulus features. For example, visual cells that respond to oriented edges may also respond to color, motion, and stereo disparity (Pribram 1991, pp. 79-81). Moreover, it is not uncommon to find neurons in *visual* cortex that are tuned to *acoustic* frequencies (Pribram 1991, p. 81, citing Bridgeman 1982; Pribram et al. 1967). Conversely, it has been reported recently (Calvert et al. 1997) that our understanding of face-to-face communication is aided by the response of *auditory* neurons to *visual* stimuli. Finally, it is worth noting that top-down signals in sensory systems can alter the receptive fields of sensory neurons, that is, their response is context-sensitive (Pribram 1991, pp. 257-58). Thus, instead of considering a sensory neuron to be a context-free detector, it is more accurate to view its response as an interaction between a complex combination of activities in the sensory recep-

tors and activity in nonsensory areas (representing context, expectations, etc.).

Much of the persistence of talk about feature detectors in neuroscience can be attributed to the same descriptionist assumptions that pervade philosophy and cognitive science. If we believe that "the only game in town" is the assembly of atomic, context-free features into abstract descriptions, then that is what we will look for in the brain, and to a large extent that is what we will find.

One unfortunate consequence of this descriptivist bias is the "binding problem," which afflicts theories of neural-net knowledge representation: How are context-free features bound together to represent objects (so that, e.g., perception of a red square and a green circle is different from perception of a red circle and a green square)? But the brain does not have to solve a binding problem, because neurons respond to complex combinations of features, that is, to features that are already bound (e.g., there are neurons that respond to the co-occurrence of redness and aspects of circularity but not to the co-occurrence of greenness and circularity, to which other neurons respond). Hence, the joint activity of a population of neurons can represent a unique complex macroscopic constellation of microproperties. In effect, the activity of each neuron represents a small bundle of conjoined microproperties, and the joint activity of a group of neurons represents a co-occurrence of a large number of overlapping bundles.

Millikan's analogy between perceptual tracking and conceptual tracking (sects. 3 and 4) reveals an important idea underlying both: invariance under transformation. Invariance is well known from the psychology of perception (e.g., size and color constancy, invariance of melody under change of pitch). Of course invariance is a central concept in mathematics, but we must be careful in applying mathematical concepts to psychology, since psychological invariants, in particular, are always approximate and limited in range (MacLennan 1994).

Invariants typically arise because various aspects of a stimulus vary coherently, and because of this coherent variation we can have knowledge of the variation of unperceived aspects. For example, when we view a rotating die, we know what the back side is doing and can predict its reappearance. In the conceptual tracking of substances we are interested in aspects that are approximately invariant over successive encounters with the substance (sect. 2). These are the aspects that cohere in the concept and about which it provides information.

Although some invariants are "wired" into the nervous system, others, including many involved in conceptual tracking, are abstracted from the coherent variation (i.e., covariation or contravariation) of multiple aspects of the stimulus. Synapses extract this information by responding to correlated activity between neurons in such a way as to strengthen strong correlations (positive or negative) and to "damp out" weak correlations (Singer 1995). [See also: Phillips & Singer: "In Search of Common Foundations for Cortical Computation" *BBS* 20(4) 1997.] Therefore, after learning, variations in certain aspects of a stimulus will lead to neural activity that mimics or primes the response to variation in other aspects that have been correlated with them in the past. Invariants become a means of generating expectations and filling in missing information. (In this way we can also simultaneously track Fido, dog, fur, and bone; cf. sect. 5.)

The "damping out" of weak correlations causes uncorrelated aspects to be eliminated from the representation, in effect projecting the concrete stimulus from the high-dimensional space of sensory-receptor activity, in which it is given, into a lower-dimensional subspace. The extreme cases – in which a stimulus is projected into a very low-dimensional subspace – produce something approximating a context-free feature detector, but such abstract features are comparatively rare and secondary to the processing of concrete microcorrelations, upon which reidentification depends.

Descriptionist theories make context-free features the elementary constituents of substance concepts, but Millikan's thesis and neural network theory together show how approximately context-

free features are secondary derivatives of concrete substances. Thus, Millikan's two continua, along which the richness of real kinds can differ (sect. 2), can be understood in this way: the multiplicity of supported inferences results from the number of synaptic connections, and the "reliability" of the inferences results from the connections' strength (synaptic efficacy), that is, the number and strength of the correlations.

The primacy of the concrete is also apparent in the context sensitivity of features. That is, the projection into lower-dimensional subspaces is dependent on some behavioral context; different features are salient, depending on whether the animal is hunting prey, seeking a mate, avoiding a predator, and so on. Meaning and relevance are primary; abstractions and features may follow as a consequence. Context-sensitive projections of this kind can be produced by using the neural representations of behavioral contexts to activate or deactivate selectively various sorts of micro-correlations; that is, from the complex combinations of microfeatures to which a neuron responds we select those relevant to the problem at hand (see MacLennan, in press, for possible mechanisms).

Millikan's nondescriptionist theory leads us to expect words to begin as context-dependent condensations out of complex clouds of pragmatic intentions. The apparent reduction of their meaning to simple, low-dimensional predicates is secondary and is in part a consequence of descriptionist presuppositions and values. In conventional terminology, the metaphorical, concrete, and context-sensitive connotation is prior to the abstract, context-free denotation (see also Lakoff & Johnson 1980). This observation even applies to such apparently abstract predicates as color terms, and part of our difficulty in understanding the use of such terms is a consequence of descriptionist assumptions. For example, ancient Greek *chlōros*, nominally translated "green," is applied to many things that are not green in color, such as dew, tears, and blood (Gage 1993, p. 272n7; Zajonc 1993, p. 15). This usage is explicable when we realize that *chlōros*, like the English word "green," may refer to things that are fresh, living, or moist (e.g., green wood, green riders). In addition, many color terms began as univalent material-substance concepts (e.g., names for minerals or dye stuffs) but appear to be polyvalent when supposed to refer to optical color (Gage 1993, pp. 34-35). So, some Medieval scarlets are black, blue, green, or white in color (since scarlet referred primarily to a kind of fabric) (Gage 1993, p. 80). The historical reduction of color to a one-dimensional predicate – wavelength – is partly a consequence of the scientific understanding of light, which began with Newton (and so offended Goethe); but we must not let this blind us to the fact that colors are primarily substances emerging from their complex meaning in our lives.

Whatever happened to meaning?

Jean M. Mandler

Department of Cognitive Science, University of California, San Diego,
La Jolla, CA 92093-0515. jmandler@ucsd.edu

Abstract: Even in infancy, concept formation has to do with creating meaning, not with tracking substances. Preverbal infants can identify a substance such as a dog, but their first concept of this substance is not dog but animal. It is difficult to account for such global concepts by the perceptual processes involved in object identification, yet these concepts are the foundation on which later concepts are built.

Psychologists and philosophers often talk past each other, because they use the same terms in different ways. So it is quite possible that I have misunderstood Millikan. But as best as I can tell, she uses the term "concept" in the way others use the term "symbol," namely, as a pointer to something else. I have always assumed that symbols, which can be arbitrary and themselves empty of content, point to concepts, but for Millikan similarly "empty" concepts point to "natural units in nature." Unfortunately, "natural unit"

goes undefined and Millikan's choice of examples for such units (cat, milk, and Mama) is unduly limited and biases her argument. The first concepts that infants form are not cat or milk, but animal and food (Mandler 1997). Animal and food fit Millikan's definition of substances as categories from which one can derive expectations about new members, but they do not have the properties Millikan needs for her argument to work.

Dogs, milk, and Mama are all categories whose instances are either highly similar perceptually or identical. Millikan uses this characteristic to claim that you do not need a defining description for them; you just need to be able to recognize an exemplar as the same thing or kind of thing you saw before. You can do that because these objects look alike, and our perceptual systems are designed to recognize them from different viewpoints and on different occasions. However, this identification ability requires perceptual similarity to work, at least for instances that appear at different times. In addition, it does not require meaning to see that this is the same squiggly pattern you saw yesterday. You do not need a concept just to identify and reidentify. To locate the core of human conceptual ability in the perceptual ability to identify instances puts the burden of concept formation on mostly unspecified but low-level perceptual processes. Millikan suggests that size and shape constancies can do some of this work, but these are *only* useful for instances of dog, milk, or Mama. Such properties will be of no use in identifying instances of animal, food, or people. Yet when it comes to meaning, as opposed to identification, it is at the latter level that infants operate. Since animals do not look alike, how is this conceptual tracking accomplished? Millikan notes that one must be able to identify a substance under diverse circumstances in order to learn its currently hidden properties (are these the true conceptual cores, then?). It is not clear how one could learn about hidden properties or conceptualize at the level of animal without some description. Millikan suggests that ways of conceptually-tracking substances emerge from insight into the ontological principles that ground them, but again she reverts to perceptual descriptions (built-in responsiveness to faces, or sensitivity to correlations among properties). In either case, since infants conceptualize at the level of animal, insight into ontological principles is unlikely to be required.

I appreciate Millikan's emphasis on developmental data in uncovering how concepts are formed. Unfortunately, her description of early concept formation and the acquisition of nouns is not quite accurate. Recent data show that 3-month-olds can indeed learn to tell dogs from cats (Quinn & Elimas 1997), and so can do the kinds of perceptual identification that Millikan describes. But they do not make conceptual distinctions among these different objects; they treat them all as the same kind of thing. The earliest conceptual distinctions infants make is at the level of animal and vehicle, not at the level of dog and cat (Mandler 1997). Since animals do not look alike in the same way that dogs do, it would seem that the only way that infants can form concepts for them is by creating a reasonably abstract description. I believe that this description is derived from perceptual information in the first place (Mandler 1992), but in itself it is not the kind of perceptual information such as shape, size, and parts (e.g., legs and fur), that we typically use to identify objects. Instead, the information that preverbal infants use in setting up concepts of kinds involves the roles that objects take in events. For example, the earliest concept of animal appears to be that it is an object that moves itself and interacts with other objects from a distance. This is not a bad first conceptual description, and it appears to be the kind of description that babies use to limit their inductive inferences (Mandler & McDonough 1996).

This description, however, is not greatly useful for tracking or identifying individual objects from one occasion to another. Moreover, it is this description that is used to comprehend the first words; hence the initial extension of many nouns is broader than that of adults. Contrary to Millikan's claims, the extensions of children's words change dramatically with development. Thus, a 2-year-old's use of the word "dog" does not necessarily point to a

natural unit in nature. It takes many months for nouns to narrow down to the extensions of adults. Conceptual foundations are laid down before children begin to speak, and early language as an avenue into understanding conceptual structure must be explored with great caution.

Too much substance, not enough cognition

Vincent C. Müller and Stephanie Kelter

Center for Cognitive Science, University of Hamburg, D-22527 Hamburg, Germany. vmueller@informatik.uni-hamburg.de; kelter@informatik.uni-hamburg.de www.informatik.uni-hamburg.de/grkl_eng.html

Abstract: Millikan's account of substance concepts is based on a notion of "substance" expanded from realist notions of individuals and natural kinds. Her metaphysical notion, based on "inductive potential," is shown to be too puristic and needs to incorporate cognizing subjects. This could preserve the realist/nondescriptionist insight that the extension of substances is determined by the world.

Millikan presents her account as a nondescriptionist alternative to standard psychological views of concepts and categorization, but the starting point of her investigations is actually quite a different one. Rather than analyzing a subject's concept formation and application, she constructs a metaphysical view of "substances" and goes on to explain how these determine concepts. Her notion of "substance," however, has serious problems that stem from the attempt to define it in purely metaphysical terms, independently of cognizing subjects.

Millikan defines substances as "natural units in nature" (sect. 1, para. 4) that have "rich inductive potential" (sect. 2, para. 9), which is "*not coincidental but grounded*," and for which there is an "explanation or cause" (sect. 2, para. 20). These causes lie in the nature of the things; neither our conceptual abilities nor our knowledge are parts of it. This is essentially the realist view of substances, for which natural kinds and individuals are paradigmatic, where "belonging to a substance" is determined by nature alone. In contrast to, and as an expansion of the classical realist view (and in contradiction to Millikan's earlier self, cf. 1984, p. 278), Millikan also considers artifacts (chair), social kinds (school teacher), mixtures (milk), and so on as natural units in nature (sect. 1, para. 3; sect. 2, para. 17).

Millikan's substances are specified as metaphysical entities, wholly independent of concepts and conceptual processing. We believe that this will not do for her purposes for the following reasons.

First, on her approach, concept formation runs like this: one encounters an object and identifies the substance, possibly referring indexically, "*that kind of thing*" (cf. Kripke 1972, p. 122). From then on, one just has to "track" that substance, to "reidentify" it. But starting off with Mama, what determines which of the many substances instantiated by her I am supposed to track? Am I to track Mama, mothers, women, redheads, tall people, that smell, or what? (Remember, we did not capture the original object by a description.) What to track is not determined by the original object but rather by the tracking organism, and it will occasionally require decisions as to which of the substances is to be followed up. The problem may possibly be avoided with well-organized and nonoverlapping natural kinds and individuals, as originally envisaged by Putnam (1973; 1975) and Kripke (1972), but not with Millikan's nonnatural substances. They will be overlapping in many dimensions, and it will not even be clear prior to cognition and recognition which substances there are to choose from. The choices will depend on the perceptual and cognitive setup of the subject, and its goals.

Second, even Millikan's criterion for substances, "grounded rich inductive potential" turns out to depend on cognitive factors. If it is to make learning from one encounter to another possible (sect. 2, paras. 3, 10, 21; sect. 3, para. 4), "inductive potential" does not

make sense in the abstract, without a subject in a situation. An omniscient subject cannot "learn from one encounter something about what to expect on other encounters" (sect. 2, para. 10). Even if inductive potential were considered in the abstract, based on clusters of correlated attributes, the problem resurfaces: it remains highly questionable whether "attributes" are simply given in the world. Psychological research suggests that attributes are "constructed" during concept formation, depending either on their usefulness in specifying and distinguishing categories (e.g., Schyns & Murphy 1994) or on the particular selection of objects that are being compared (e.g., Medin et al. 1993).

Third, all things are embedded in numerous correlations and, in the abstract, all such correlations are "grounded" somehow, not merely accidental; Millikan has no way to delineate "the right kind of groundedness." For instance, in the case of complex kinds – such as *bearded woman* – that are said to be merely nominal (sect. 2, para. 12), there *is* inductive potential for some subjects which goes beyond that of its components, so it should be a substance. (The mere composedness of the expression "bearded woman" in English cannot be to the point here.)

Fourth, it is not that something has inductive potential or not, but rather, inductive potential is a gradual affair. Millikan effectively admits this (sect. 2, para. 6), but it is hard to see how she can live with it, since (1) drawing a metaphysical difference between substances and nonsubstances based on "more or less" potential is odd, and (2) taking kinds to be "more or less substantial" would make tracking "more or less" successful.

Millikan's attempt to expand criteria that were made for natural substances just runs into the same difficulty as the traditional psychological approach it is supposed to supplement: conceptual structures and substances cannot be specified independently of one another. And yet Millikan remains in a better position: the substances she needs for her project do not actually require such purely metaphysical criteria. The ability to track substances through time and changes as well as the "capacity to recognize what is objectively the same substance again as the same" (sect. 4, para. 13) can be preserved even if substances are the result of human choice out of the very many classifications the world allows, a choice dependent on both the makeup of the world and the makeup of the choosers. This would preserve the realist insight that once a particular choice has been made, the extension of the substance and concept is determined by the world, not by what subjects believe (thus retaining Millikan's distinction of ontological fact and fallible cognitive processing). It would also fit even better with her suggestions concerning the cognitive mechanisms of conceptual development.

Extensional assumptions in theories of meaning and concepts

Gregory L. Murphy

Department of Psychology, Beckman Institute, University of Illinois, Urbana, IL 61801. glmurphy@uiuc.edu

Abstract: The problems that Millikan addresses in theories of concepts arise from an extensional view of concepts and word meaning. If instead one assumes that concepts are psychological entities intended to explain human behavior and thought, many of these problems dissolve.

In spite of Millikan's final sentence, the arguments for a non-descriptive theory of meaning strike me as being exactly an argument over what "the meaning" is. In essence, Putnam's arguments rely on the assumption that meaning is extension. His examples are those of deciding what is in the extension of "gold" or "tiger." The result is that questions of extension become the sine qua non of concepts and meaning, so that other issues are seen as being not quite about concepts. For example, Millikan points out that having an understanding of a domain "need make no differ-

ence to the *extensions* of one's substance concepts" (sect. 5, para. 13) and so presumably is not really part of the concept. As Putnam famously concluded, and Millikan agrees, this leads to the paradox that what people know does not count as meaning, because people cannot provide criteria for the correct extension of many terms. It therefore becomes necessary for someone like Millikan to make a helpful attempt to explain how people (who cannot specify the extensions) nonetheless appear to have concepts.

It is possible, of course, to take a different approach to meaning (from a Gricean rather than a Fregean perspective), in which linguistic meaning is a property of a language system that is a product of social convention, which therefore must be representable by the members of that social community. We have internalized the conventions of English, which allow us to decide when to use the word "mouse" as opposed to "cat" and "cheese." In contrast, because Putnam thinks that our ability to use these words is not knowledge of their meaning (since we cannot correctly specify their extension), he must develop other theoretical entities to explain what it is that people *do* know.

If one assumes that language users do know the meanings of words and that those (descriptive) meanings are responsible for much language use, the descriptive properties of language can be admitted without shame. For example, instead of insisting that "to call a thing 'gold' or 'mouse' is not to describe it" (sect. 1, para. 4), we can now account for the descriptive use of such terms in predicates:

- (1) – Is that expensive?
– It's solid gold!

Since these nonessential, descriptive properties of gold are used in identifying gold and in determining word use, it seems most natural to include them as part of the concept, rather than to narrowly define the term "concept" so that it excludes the mental representations that actually account for our behavior. After all, it is these actual tokens of behavior – sentences, inferences, word uses, and so on – that we are trying to explain by creating the theoretical construct of a concept.

Millikan's analysis of reidentification suffers from self-imposed extensional handcuffs. On her view, it is the "real extent" (sect. 4, para. 8) of the kind that determines one's concept, not one's internal disposition to apply the concept. There is a real question, however, about how many concepts can be defined extensionally. That is, which concepts have a "real extent" that can be defined separately from their individual or societal definition? For example, it is very difficult to justify any "real extent" of artifacts, which are categorized on the basis of a variety of functional and physical characteristics in a probabilistic manner (Malt & Johnson 1992). Social categories are also difficult to characterize as having some external reality apart from the social conventions that assign their descriptions. Even many apparent natural kinds are not: people's categories of "tree" or "fish" do not correspond to any recognized biological class but to entities that fit certain descriptions to some degree of goodness (e.g., "tree" = tall, woody plant, with leaves and bark). These categories certainly have extensions, but they are the cart and not the horse. No doubt environmental structure is a major factor in the shaping of our concepts. It does not follow from this that one can rely on the environment to define the category on its own.

The extensional account has specific problems in explaining the concepts of preverbal infants and nonverbal animals. Infancy researchers argue that infants have developing concepts of *containers*, *obstacles*, *supports* and other basic physical categories. But since the infant has a very incomplete grasp of what is actually a container, etc., it is impossible to count on physical reality to provide the "real extent" of the infant's concept. Instead, the concept must be based on the infant's own experience and understanding of containers, i.e., it must be descriptive. Similarly, your dog might have the concept of "door," based on its own activity of going in and out of rooms and houses, but it is unlikely that this concept corresponds to any physical or human artifact category (e.g., it might include pet doors or windows that humans do not

think of as doors). In both cases, the only way to determine an extension for the concept is to reproduce the organism's description, perhaps in the context of the rest of its knowledge and beliefs.

The Putnam-Millikan view has apparent plausibility for adult linguistic concepts, because the conventions of language provide an external criterion for meaning – I cannot mean just anything by "door," if I wish to speak English. But it seems to have no purchase on describing the concepts of nonverbal beings whose concepts do not correspond to preexisting physical or conventional categories. Again, in all these cases one can define an extension, but to use this as the definition of the concept is to get the story backwards. Infants decide what they think a container is, and certain items turn out to fit that notion; the conventions of English decide what criteria to use in calling something a "mouse," and then objects turn out to fit it or not. The exact role of the environment in determining these concepts is an interesting and as yet unanswered question. We cannot ask this question with an open mind about the complexity of the answer if we assume that the concept is defined by the "real extent" of the kind.

Beyond substance concepts in cognitive development

Katherine Nelson

Department of Developmental Psychology, City University of New York Graduate School, New York, NY 10036. knelson@email.gc.cuny.edu

Abstract: Millikan's theory of substance concepts has advantages for psychological theories, including those in cognitive development. However, the disadvantage is that it cannot be generalized even to some of the most common concepts that children acquire in the early years of life. For a general theory we must get beyond substances.

Millikan's target article goes some way toward illuminating the long-standing observation that children have difficulty forming artificial concepts in the laboratory but succeed spectacularly in the real world, as evidenced by their acquisition of upwards of seven new vocabulary items per day between the ages of 2 and 6 years old (Nelson 1974a). If I have understood the proposal correctly, there are essentially two parts to the concept: an identification scheme that is causally related to the substance that exists in the real world, and a "box" that holds information about the substance, knowledge gathered through practical experience. The proposal eliminates the concept-formation problem considered as logical classification in terms of necessary and sufficient features, such as "white triangle." It thus disposes of the tradition of research viewing young children as deficient in conceptual structure. Children are competent at forming concepts causally related to things in the world but are not skilled at logical classification.

An advantage of the theory is its connection to evolutionary as well as developmental processes based on experience and practices in the world: "the concept has been tuned to its present accuracy by causal interaction . . . during the evolutionary history of the species or through the learning history of the individual" (sect. 5, para. 14). "Tuning" derives from Gibsonian theory; it denies the concern embedded in most developmental investigations for "correct mappings" between the child's words or concepts and "true" concepts or meanings (backed in recent writing by Quine's [1960] much quoted "gavagai" parable). However, the proposal (like many others in philosophy and psychology) suffers in not going beyond substance concepts to encompass others that children find easy to learn and use, including most of those comprising the vocabulary explosion of the preschool years (Nelson 1995).

Millikan's form of realism apparently prevents her from adequately addressing those concepts that derive not from perception alone but from capacities made possible by human language. She allows that many substance concepts are learned through language

and may be changed through language use, but both the process and the concepts considered are limited in scope. Although infants have substance concepts prior to language, acquiring language changes those concepts in significant ways, specifically from global to basic level (Mandler 1997). Restricting the (preverbal) concept of "dog" from all mammals to (verbal) dogs represents a more fundamental change in the concept than adding new identification features or adding more information to the conception of dog. There is an ambiguity in the word-concept relation discussed by Millikan that ends in the disclaimer "best not to fall into a verbal dispute over what gets to count as 'knowing the meaning.'" This problem becomes more acute when one considers whether this position can be extended to types of concepts other than substances.

At various places Millikan suggests that her proposal can be applied to events, cultural artifacts, social categories, and musical compositions, but she does not tell us how. Like virtually all the authors in the psychological literature that she cites, Millikan recognizes that a theory of concepts must cover more than substances but chooses to start with those and worry about further extensions later. This will not be so easy, however, even for actions, states, and properties (Merriman & Tomasello 1995), much less for abstractions such as "justice" or "electron." Millikan (sect. 2, para. 1) relies on the common assertion that "the bulk of a child's earliest words are concrete nouns." This probably seems harmless, but inasmuch as fewer than half of a child's first vocabulary words are common nouns (Nelson 1995), it leaves much to be explained. Moreover, many of these first nouns are not names of concrete substances but of events (bath), social roles (doctor), localities (park), times (morning), and other constructs that are difficult to identify by pointing.

Thus, an important group of words and concepts lies in an abstract realm between that of "justice" and "truth" and that of "cat" and "cup." These kinds include among others social category terms such as "doctor," "teacher," and "brother" – concepts that cannot be "causally determined by the substance it denotes." For example, children's first concept of "doctor" may point to the person they visit periodically for checkups, but this is a concept of Dr. B, not a concept of "doctor." The latter concept depends on additional experience with the constellation of words as they are publicly used in experienced worlds (Wittgenstein 1953). Consider, for example, "help," a word learned and used very early by young children. Surely the child has a concept of "help," but what could it point to? In place of a real substance there is a constellation of experiences in the world including the term "help" (no less real than Mama herself) that lead the child to an implicit understanding of the term.

The basic assumptions of Millikan's theory may hold for these other kinds of concepts in that they have identifying schemes – they serve practical functions in human lives, are worth acquiring knowledge about, and so on. They are not arbitrary logical constructions, but neither are they causal in the same way that the realist theory presumes. Rather, they depend on both social arrangements and linguistic symbols for their construction. It is not simply a matter of learning words and what they point to, but of how the word-concept fits into the matrix of a social-cultural world.

In sum, the task of constructing a psychological theory covering concepts and words basic to human lives (including children's) requires getting beyond substances and beginning with those common concepts that we use words to express although or because we cannot identify them by pointing.

Room for concept development?

Josef Perner

Department of Psychology, University of Salzburg, Hellbrunnerstrasse 34, A-5020 Salzburg, Germany. josef.perner@sbg.ac.at

Abstract: Millikan's externalist account of concept acquisition cannot completely avoid the distinction between central (defining) and peripheral (characteristic) features, because some knowledge is required to achieve reference and to decide what kind of information to record about the identified substances. However, the emphasis on external reference may provide the requisite principled way to make this distinction.

People who study cognitive development tend to have the strong intuition that children's thinking differs not just in that they have different ideas, but more fundamentally, that some of their elements of thought (i.e., their concepts) differ from ours. Moreover, there is also the intuition that conceptual change is not erratic, switching among unrelated conceptual systems, but that it increasingly approximates the adult system. Although the older child's concepts differ, they do have ancestors in the younger child's thoughts (protoconcepts). This intuition found a natural ally in the classical view of necessary and sufficient conditions, and Keil (1989) spoke of defining and merely characteristic features of concepts to mark the distinction between those aspects that matter for having a concept and those that do not.

With the fall of the classical view it became clear that, strictly speaking, there are no "defining" features. This does not imply that there is no distinction between, say, central features that matter for concept possession and peripheral (merely characteristic) features, but a principled way of making this distinction has been lost. The lack of such principles remains a critical problem for the now-dominant "mid-twentieth-century doctrine that the 'meaning' of a . . . concept is a matter of its connections with other . . . concepts" (sect. 4, para. 7). Failure to find a solution threatens rampant meaning holism and conceptual instability in the face of theory change and, thus, provides one strong motivation for externalism, the path that Millikan follows, where the meaning of concepts is fixed by direct reference to external substances. I now have two questions: (1) Can externalism circumvent the need for drawing the distinction between central and peripheral features? (2) Will it provide us with a theoretical foundation for this distinction?

At first glance, circumvention seems possible, because the distinction does not seem necessary for determining possession of a concept. The distinction belongs to the issue of understanding what the grounds of a substance are (sect. 2, last para.), which is the realm of "conceptions" (Woodfield 1991). If successful, this move may solve the problem of meaning holism, but it will not provide a basis for the developmental psychologist's intuition. In any case, when one considers Millikan's suggestion about concept possession in detail, the success of this move appears less likely. There are two conditions for having a concept of a substance: (1) one must have the means to refer to the correct substance, and (2) one must know what kind of information to collect about that substance. Let me illustrate the issues that these conditions raise with children's acquisition of the mental concepts "belief" and "pretence."

Woodfield (1996) in his endorsement of Wellman (1990) takes the very extreme position that the use of the linguistic term "think" suffices for possession of the concept "belief." Is this enough to achieve reference? In general, reference can only be achieved in a discontinuous world. Only if there are discernible objects in the visual realm (and not a uniform Ganzfeld) can my pointing finger or referential description make clear what I have in mind. However, not all the things we have concepts for are equally distinct. Some are more closely related than others and form a hierarchy of distinguishability. "Belief" and "pretence" are more similar concepts than either of them is to "desire." The former are maps by which we steer, the latter specifies where we want to go.

My claim about 3-year-olds would be that, by and large (even though there is some evidence that they sometimes use "think" to

mean “desire”; Wellman & Bartsch 1988), they have the conceptual ability to distinguish desires from beliefs and pretence, but they cannot yet distinguish belief from pretence. Although they can distinguish false from true, they are not yet sensitive to the possibility that something false can be evaluated as true (Perner 1995). If this is right, then these children are in danger of encoding information about false beliefs as information about pretence, and from their use of “think” it is not clear whether they are referring to beliefs, thoughts, or pretence. Since we have no verbal label for such a concept, I coined the term “prelief” (Perner et al. 1994).

Children’s use of “think” is obviously close on target; it refers to prelief that distinguish beliefs from desires and that seems to be enough so that “think” (meaning prelief) is used correctly on the majority of occasions. However, one cannot conclude from such mostly correct use that children mean belief and not just prelief. Moreover, since prelief lumps belief and pretence together in comparison to other mental states, it is a natural protoconcept of belief and pretence.

I would like to agree with Millikan (and Woodfield) that to have a concept of a substance the child does not need to fully “understand what the grounds of that substance are.” However, for the child to have the same concepts as an adult, the child needs to have enough knowledge resources to focus reference (at least in principle, above chance) to the same fineness of grain as the adult conceptual system. Any excess pieces of knowledge, allowing finer distinctions (e.g., different types of beliefs for which we do not have concepts) are mere conceptions.

I therefore suspect that the answer to my first question will have to be negative: the externalist account of concepts will not be able to circumvent completely the distinction between central and peripheral features, because, as I have tried to illustrate, some features need to be understood for focusing reference. Moreover, if I had space to turn to Millikan’s second criterion of deciding which kind of information to record about identified substances, the need for such a distinction would turn out to be even more pressing.

The answer to my second question may turn out to be positive: the need to identify substances may yield a criterion to determine which features are central and which peripheral. Knowledge of features necessary to focus reference and decide which information to record is central; other knowledge is merely peripheral.

Can mere phonemes be components of Millikan’s substance concepts?

Niko Scharer

Department of Philosophy, University of Toronto, Toronto, Canada,
M5S 1A1. nscharer@chass.utoronto.ca

Abstract: In presenting her attractive theory of concepts, Millikan makes an unwarranted assumption about the role of language in concept acquisition. The phoneme string, rather than the “word” as a semantic entity, may suffice to play the crucial role in the acquisition of substance concepts. Hence Millikan may underestimate the degree of similarity between language and other media of perception.

Language plays a crucial role in Millikan’s attractive, nondescriptionist theory of concepts. She is certainly right in observing that adults acquire information by hearing and believing what others say, and that children gain this ability by learning language (sect. 6, para. 5). Yet there is a gap between how words are used by preschoolers learning to identify cats by hearing and calling animals “kitty” and how they are used by mature language users who can be “handed” a concept like African dormice (sect. 6, para. 8) and for whom the sound of rain can be something like “Hey, guys, it’s raining!” (sect. 6, para. 4). Preschoolers need not believe anything about “kitties,” at least not in the sense that adults must understand and believe the utterances that are the source of their information. But perhaps Millikan is making an unwarranted

assumption about young children’s early use of words. Unlike adults, preschoolers need not be relating to “words” as semantic entities when acquiring concepts. I will argue that on Millikan’s own theory, the phoneme strings which adults treat as words might, in young children, serve their function in concept acquisition without being treated semantically.

As Millikan observes (sect. 1, para. 5), many animals and preverbal humans have substance concepts (i.e., they can reidentify substances). Indeed, many animals respond to verbal commands. Yet we do not assume that animals respond to commands because these are expressed in *words*, but only because commands are audible and distinguishable from other sounds. Likewise newborns respond to their mother’s voice only because the sounds are distinguishable and familiar. Virtually all “words” that children encounter are mere sounds when first heard.

On Millikan’s account, the early development of substance concept involves all the senses as children learn different ways of identifying substances. Yet for many substances the sounds that a child hears when seeing, touching, or smelling a substance are in fact the sounds of words, not the sounds made by the substance seen, touched, or smelled. Consider a preschooler’s experience with cats. The sounds that most regularly accompany perceptions of cats are the phoneme strings “cat,” “kitty,” and “meow.” These are probably heard more regularly than the purring and meowing sounds that cats themselves make. In addition, such phoneme strings are generally produced when the child’s attention is already focused on the cat. By contrast, many of the cat’s own sounds occur independently of the child’s attention.

Using Millikan’s account of concept development, the sounds “cat,” “kitty,” and “meow” are ways of identifying the substance cat. Yet the ability to use these sounds in reidentifying objects may be part of the child’s substance concept in just the same way that using the cat’s smell, shape, or purring to reidentify cats would be, and not as a consequence of the fact that those phoneme strings are *words* or labels applied by competent language users. We should not assume that just because adults treat “kitty” as a *word* for cat, the child does not treat it as a mere *sound* component of cat, albeit a sound made by other individuals. We can accept Millikan’s contention that language is a medium of perception for competent language users, but we need not assume that it functions in the same way during concept development. The relevant medium of perception for the infant could be mere sound, even if when falling on *our* ears those sounds function as words. Indeed, this reflection may work to Millikan’s advantage, since long before the “noun explosion,” hearing phoneme strings might help the child learn to categorize objects into the groups that common words come to pick out.

There are clearly important questions to address concerning the difference between hearing a phoneme string and hearing a word, philosophically and neurophysiologically. My point is simply that we should not assume that the child is using *words* rather than mere *sounds* in acquiring concepts. Early experience with language might be better characterized as experience with phoneme strings.

Although Millikan argues that conceptually tracking substances through language is more like tracking substances through other perceptual means than is generally recognized, she continues to give language special status. She calls words “handles to hang onto,” as if they were attached to substance concepts in the way handles are to cups. That may well be true about words as semantic entities, but Millikan’s own account implies that for very young children the corresponding phoneme strings may be attached to objects in much the same way that color and shape are, as integrated parts of the cup itself.

On Millikan’s account, having a substance concept is having the ability to conceptually track substances, often through different perceptual means. If young children developing substance concepts learn to track substances conceptually using phoneme strings (uttered by others) as well as visual, tactile, and other perceptual cues, then the resulting concept will integrate these

disparate sources of information. This implies that the role “words” play in this early stage is much like that of colors, shapes, textures, and tastes. The child grows to expect objects in the world to be colored, shaped, textured, flavored – and *named*. In the early substance concept a name may play the same role as any other property. Although Millikan and others are right that words can be used to stake out new categories that will later be explored (sect. 6, para. 9), this is also true of other properties. For example, I might ask what spice I tasted in a dish, or I might try to identify a bird I merely heard or saw. To the young child, the name of a substance may be as much a property of that substance as its other persistent features – objects have colors, textures, tastes, and *names*.

In conclusion, Millikan’s nondescriptionist theory of concepts may imply that the role of language in concept acquisition is even more similar to that of other media of perception than Millikan recognizes. Although Millikan emphasizes the similarity between language and other perceptual media, she apparently does not appreciate how deeply her own theory assimilates them.

Explanatory force, antidescriptionism, and the common structure of substance concepts

Jürgen Schröder

Grünwaldstr. 12, 69126 Heidelberg, Germany.
jschroel@urz-mail.urz.uni-heidelberg.de

Abstract: Millikan’s proposal of a common structure of substance concepts does not explain certain conspicuous findings in the psychological literature such as typicality effects, the context sensitivity of these effects, and slips of the tongue. Moreover, it is unclear how antidescriptionism could be relevant to psychological theorizing. Finally, it does not seem to be true that concepts of individuals, stuff, and real kinds have a common structure in older children and in adults.

Explanatory force. A satisfactory account of concepts (including substance concepts) has to be able to account for a variety of empirical findings of psychologists that have also been taken to have implications for the nature of concepts. There are, for example, the findings of Rosch and others (e.g., Rosch 1975a; Rosch & Mervis 1975), concerning the typicality of exemplars of a category. It is not clear how a theory that regards concepts as structured capacities but not as structured representations can explain these findings. Prototype theory, however, which assumes that concepts consist of features that are more or less probable of exemplars of a category, explains typicality effects by how well different exemplars match these features.

True, on Millikan’s account concepts are not representations themselves but capacities to represent, and these capacities are structured insofar as they are realized by subcapacities. But we consider the subcapacities for reidentifying a substance, an explanation of typicality effects does not suggest itself. According to prototype theory it is the matching process between features of the exemplar and features of the concept that is crucial for the explanation of typicality. Not only is there nothing in Millikan’s account that could be analogous to this process, but even if there were it would be unclear on what it would operate: it cannot operate on features, because there are supposed to be none; and if one term of the analogue process were the reidentifying capacities (corresponding to the features of the concept), it would still be unclear what the other term was (the one corresponding to the features of the exemplar). So there seem to be no conceptual resources in Millikan’s theory to deal with typicality effects. Likewise, if the theory cannot explain typicality judgments, then *a fortiori* it cannot explain the context dependence of typicality (Roth & Shoben 1983) or in general the instability of typicality judgments (Barsalou 1987).

Another class of effects that the assumption of structured concepts helps explain are slips of the tongue, for example,

replacing words with semantically similar words (Fromkin 1971; McNamara & Miller 1989). Here it is assumed that the process that selects the words operates on the basis of a comparison between features of the concept representations and features of the word representations. It is a partial malfunction of this process that explains why the wrongly selected words are semantically similar to the correct ones. Again, there seem to be no resources in Millikan’s theory to account for such phenomena.

The relevance of antidescriptionism. Because psychologists are in the business of explaining their data, and because Millikan attempts to give them a new framework by making “substantial preliminary assumptions” about the nature of concepts (sect. 1, para. 1), it is legitimate to ask what the various ingredients of the new framework will buy the psychologist. We have already seen that with respect to several explanatory tasks the ingredient of structured capacities is unsatisfactory. Now we ask whether the property of being nondescriptionist might be of any use to the psychologist. One reason that antidescriptionism may be a good thing is that it ensures that different people can have the same concepts. If concepts are individuated semantically, the meaning of concepts depends in the first place on their extensions; and if the extensions of children’s and adults’ substance concepts are the same, because they are determined by the substances themselves, then children and adults or people in different cultures have the same concepts, because they interact with the same natural kinds.

Apart from this, however, does antidescriptionism shed any new light on the explanatory tasks of the cognitive psychologist? Descriptionism and its antithesis are both views about the relation of representations (natural language words) to the world, since they are views about what determines the reference of a term. According to descriptionism, in the determination of reference the mind is primary. Whatever satisfies the properties associated with a term is its referent. According to the rival view, the world is primary: whatever *actually* governed the application of a term is its referent. Because the contrast between descriptionism and antidescriptionism concerns the relation between representations and the world, and the task of the psychologist is to devise explanations of certain behavioral regularities, it is difficult to see how this contrast could have any bearing on this kind of task. Moreover, the explanatory work is being done by the assumptions about the representations themselves, about the “syntactic” or formal aspect of them (e.g., whether or not they have internal structure). But this aspect is irrelevant in the debate between descriptionism and antidescriptionism. (For an assessment of descriptionism and its rival in the context of proper names, see Schröder 1994.) On the other hand, considerations that deal exclusively with the meaning of representations likewise seem to be irrelevant to the concerns of the cognitive psychologist (except for the shareability issue).

From common to different structures? Perhaps Millikan’s most important proposal is that the structure of concepts for individuals, real kinds, and stuff is the same. This structure consists in the subcapacities that make up a substance concept: (1) the capacity for reidentification, (2) the capacity to know what can be learned about a certain kind of substance, and (3) “the capacity to store away information gathered about it such that it is always represented again with what one understands to be another representation with the same semantic value” (sect. 3, para. 6). Although this structure may be common to all substance concepts, there are important differences in the second of the subcapacities. Knowledge of what carries over in the case of individuals is not the same as knowledge of what carries over in the case of kinds. If I know that the color of my son’s eyes will not change from one encounter to the next, I know equally that the color of the eyes of humans does change from one to another. Likewise for other properties. Furthermore, it *could* be that very small children represent Mama, milk, and mice without the distinction that the first is an individual, the second a stuff, and the third a kind, but when they grow older they are able to make these distinctions, so that the concept of “more” will only be combined with concepts of stuff but not with concepts of individuals. But do these concepts still have

the same structure? Does not the common structure (if “common structure” includes the absence of these distinctions) belong only to the substance concepts of small children? And does not the differentiation in the second subcapacity engender a structural difference?

More me? Substance concepts and self concepts

Carol Slater

Department of Psychology, Alma College, Alma, MI 48801.
cslater@alma.edu

Abstract: User intentions invoked to account for the distinctive way in which public-language natural-kind terms gather their extensions are inapplicable in the case of Millikan’s substance concepts. I suggest that theoretical justification is preferable and available and raise exploratory questions about the applicability of the notion of substance concepts to the genesis of self concepts.

Millikan admires the work of developmental cognitive psychologists, a regard that is undoubtedly mutual. Neo-Vygotskians will applaud Millikan’s advocacy of language as a tool kit for thought; neo-Piagetians will approve of baby scientists busy discovering the nature of the world; Gibsonians and situated action theorists will second the motion that research tasks be vetted for ecological validity, and developmental psychologists of all persuasions will surely already have endorsed Millikan’s call for a “wider developmental psychology,” one that investigates how “normal (ideal) relations between the environment and . . . cognitive structures . . . are put in place in normal (ideal) conditions” (Millikan 1993a, p. 169). Development is, after all, a deeply normative notion; like Millikan, developmental psychologists are concerned with how we take the shape we should take when all goes well rather than with documenting the way things typically turn out. There are, nevertheless, a couple of places in the target article where psychologists might want to solicit Millikan’s assurance that they are on the right track.

The first of these has to do with broadening the historical/causal approach to reference to cover (at least some) mental concepts. The extension of a substance concept, Millikan tells us, is (non-descriptively) determined “not by one’s fallible dispositions to recognize portions of its extent, but by the real extent of the substance that has governed the development of these dispositions” (sect. 4, para. 9). Even a psychologist who is familiar with the Kripke/Putnam historical/causal account might be uncertain about what is supposed to make this the case for such concepts. That a particular public language term “X” gathers its extension in the manner characteristic of natural kind terms has variously been attributed to the intentions of language users, to their pragmatic competences, or (perhaps) to a natural kind “marker” associated with X in the user’s lexicon: it is, for example, proposed that it is Xs being launched (transmitted, deployed) by users with the intention of referring to a natural kind with the same microstructure as certain ostended instances that gives X a causal/historical semantics (or presemantics).

This, however, does not seem a promising line to take with regard to mental concepts, not the least because Millikan has explicitly rejected speaker intentions as a basis for construing linguistic meaning (Millikan 1984). Millikan would, I think, prefer and have available a straightforwardly theoretical justification. A possible parallel might be found in Richard Boyd’s (1979) suggestion that linguistic natural kind terms refer to whatever “natural substance” they afford “epistemic access” because this provides the best theoretical explanation of how language performs its function in a communal knowledge-gathering enterprise.

Psychologists might also want to hear more from Millikan about how – or even whether – her account of substance concepts could be applied to what is often called a *self* concept. In some ways it

appears eminently applicable, but there are also some odd stretches. To begin with, do Baby’s capacities for purely perceptual tracking deliver “more me” in the same way that they deliver “more Mama”? And if tracking oneself over short-time spans is ever accomplished purely perceptually, would the temporary implicit intensions iterated for the occasion – inner “me”s? – be indexical because perceptual? It would seem so. But could mental “me’s” be indexical any more than mental “I,” given that the referent of neither mental term actually varies with context? (see Millikan 1993b).

It also seems distinctly odd to think of Baby as having or developing a capacity to carry over information about herself from one encounter to another (Encounter with herself? Are we ever out of contact with ourselves?). Still, different contexts surely make available to us different information about ourselves just as they do about other substances. (Even grown-ups have been known to surprise themselves on occasion.) And a 5-month-old’s recognition that the swinging legs she sees on a video screen are the swinging legs she feels (or, perhaps, the legs she feels herself swinging?) clearly evidences a capacity to coordinate information about herself obtained from different sources (Watson 1985, cited in Vasta et al. 1995), as does performance on a standard self-recognition task (e.g., touching on one’s face the rouge spot one sees in a mirror).

Moreover, as in the case of other substance concepts, children have to learn what of their own properties can be “projected” from one time to the next: that one (at least normally) stays male or female is a notoriously late discovery. Finally, here as elsewhere it would appear that adequate concept-development requires a cooperative world. It is at least suggestive that poorly attached children are reported to be delayed in self-knowledge (Pipp et al. 1992, cited in Vasta et al. 1995). Perhaps sensitively contingent response by caretakers contributes simultaneously to adequate concepts of both Mama and me.

A psychologist who wants to get into serious conversation about any of this will I think, have to give at least as much careful consideration to crafting “self concept” into a useful theoretical term as Millikan has given to developing the notion of a substance concept. Clearly, more research needs to be done.

Semantic realism, rigid designation, and dynamic semantics

Alice G. B. ter Meulen

Department of Philosophy, Indiana University, Bloomington, IN 47405
atm@phil.indiana.edu

Abstract: Semantic realism fits Millikan’s account of kind terms in its focus on information-theoretic abilities and strategic ways of gathering information in human communication. Instead of the traditional logical necessity, we should interpret rigid designation in a dynamic semantics as a legislative act to constrain possible ways in which our belief may change.

Millikan’s account of the role that kind terms play in our information gathering strikes a deeply congenial chord with me, a working natural language semanticist. Semantic realism – that is, the ecological program of a theory of meaning and interpretation that locates information in the world and relates it to our actions – embodies an anti-essentialism that sharpens its controversy with Kripkean possible world semantics. It may still be possible to consider the origin, biological or otherwise, of an object, an attribute it cannot lose without losing its metaphysical identity. But as Millikan aptly points out (sect. 2, para. 8), there is no semantic need to impose such a Kripkean-Aristotelian conception of rigid designators on the ways we refer to kinds or individuals. The modality at stake in the axiom of rigid designation (i.e., if $a = x$, then it is logically necessary that $a = x$) should be interpreted as an *epistemic* one, not as a metaphysical claim about the ways in which the world could or could not be different from how it is, given the logical conventions of fixing the reference of free variables.

Epistemic modalities are interpreted in a dynamic semantics as quantifying over ways in which a current belief state may change. When we establish a referent and assert of it that it is (an) XYZ, we create a context that commits anyone who intends to be a cooperative conversational partner to use XYZ subsequently with the intention to co-refer and to share our belief that in some way, perhaps unknown to both of us, the targets of our referential acts are considered similar. Initiating a chain of reference is hence a legislative information-theoretic act, constraining the ways in which the given context may change.

This is how Millikan's "rich inductive potential" is guaranteed to accumulate information about the same kind or individual, which provides kind terms with their flexible, if fallible, and adaptive explanatory power. Anyone violating the referential constraint forces himself out of our context, is considered uncooperative, and must hence be dismissed by us as someone belonging to another way of interpreting the world, another "worldview." Reverberating in this pragmatic conception of the semantics of kind terms and proper names is the classical Kantian schema, the *a priori* structuring of the content of our possible experiences, making causal claims assertible and providing a richly structured context of justification and explanation.

Proper names, kind terms, and indexicals give us the same ability to form beliefs inductively with the guarantee that the information we gather is about the same object, no matter how it changes. This allows us to regard any observed changes in an object as a change of properties and not a violation of its identity. Even name changes can now be considered a change of property, for a name is merely a label of one way of accumulating information about an object and relabeling is a simple information-theoretic housekeeping act. Even if the two names are retained, we can go smoothly back and forth between the two ways of referring to what we know is the same object, as long as the information is preserved that the two names co-refer. But major problems arise in our "information housekeeping," however, when we somehow find out that two expressions whose co-reference we had taken for granted turn out to refer to two distinct objects. Sorting out what in the information we stored is about which object is a comprehensive overhaul of our belief state that is as inevitable as our need for providing sound explanations about the two objects.

It is a major gain of semantic realism, as Millikan points out (sect. 6, para. 2), that it considers different ways of gathering information to contribute equally to the formation of beliefs. Perception is just another way of forming a belief, which has much in common with the interpretation of what someone communicates to you in language. There is the same reason for caution: seeing is believing, not necessarily knowing! Similarly, understanding what someone tells you does not logically entail its truth,

nor does it entail your attitude of holding it to be true. But in daily practice, we take perceptions under what we consider to be circumstances that favor veridicality and minimize distortion to be a source of reliable information, just as we normally cruise on the trusting assumption that someone who tells you something must have good reason to consider it true.

A white thing

J. van Brakel

*Institute of Philosophy, University of Leuven, 3000 Leuven, Belgium.
pop00127@cc5.kuleuven.ac.be*

Abstract: I have no problem with Millikan's saying that Mama, milk, and mouse are substances, but I do not see why this list cannot be extended with white, red cows, things, *vovetas*, *lhenxa*, GRUE, and so on. In the right circumstances, given the right training, the characteristics of substances that Millikan provides work equally well for each of them.

Millikan says that "no knowledge whatever carries over about nonsubstance kinds, such as *the red square*," and a white thing "is not on the scale with substances, for there is nothing to be learned about it" (sect. 2, para. 8). I cannot follow this. Millikan's appeal to Quine suggests that in addition to more gold, more milk, and more Mama, there would also be more white and more red (cf. Quine 1960, pp. 90–105; 1969, pp. 35f). This is supported when she says that "red square" is a nonsubstance, because knowledge that applies to it "applies to one or another of the analytical parts of these complexes taken separately" (sect. 2, para. 8). Hence, red and square are, presumably, substances. (However, a square is an equilateral rectangle, so it is also *not* a substance.) Following this line or reasoning, "white thing" would not be a substance, because it "applies to one or another of the analytical parts of these complexes taken separately." But if all substance concepts are subject concepts (as Millikan says they are), they are all "things." This might lead to the conclusion that there is only one substance left, namely, "thing." However, I am not sure how Millikan uses the words "thing," "subject concept," "object," or "individual." Would she agree with many (any?) of the quotations in Table 1? Is "thing" an artifact of English grammar (perhaps similar to classifiers in Japanese)? But if "thing" is neither a substance nor a nonsubstance, why is white or red not a substance? In addition to the Quinean "more gold," why would "more yellow" not qualify as a substance? Or "more white," by analogy with "more milk"?

Millikan saying that a "white thing" is not on the scale with substances, because nothing can be learned about it, reminds one of Mill (1843, p. 122): "White things are not distinguished by any common properties, except whiteness: or if they are, it is only by

Table 1 (van Brakel). *What is a thing, object, individual?*

Source	View expressed
Chomsky (1995, p. 30)	What is a thing, and if so, what thing it is depends on specific configurations of human interests, intentions, goals, and actions – an observation as old as Aristotle.
Hardcastle (1994, p. 590)	The world may be, but probably is, radically different from the way we perceive it. At the very extreme, we may have an "object" proto-theory built into our perceptual system that forces us merely to <i>interpret</i> the world as being filled with things.
Lowe (1989, p. 11)	Individual are only recognizable as <i>individuals of a sort</i> , whereas sorts are only intelligible as <i>sorts of individuals</i> .
Quine (1992, p. 6)	The very notion of an object at all, concrete or abstract, is a human contribution, a feature of our inherited apparatus for organizing the amorphous welter of neural input.
Sellers (1963, p. 9)	To ask what are the basic objects of a (given) framework is to ask not for a <i>list</i> , but a <i>classification</i> .

Table 2 (van Brakel). *Which are the real substances?*

Substance	Description
Milk, juice, mouse, gold	Examples used in target article.
Mama, shoe, house	Examples used in target article.
<i>Vovetas</i>	Tsistsistas [Cheyenne] word, the reference of which includes most vultures (<i>Cathartidae</i>), the common nighthawk (<i>Chordeiled minor</i>), swarms of green darners (<i>Anax junius</i> , a dragonfly), swarms of red skimmers (<i>Libellula saturata</i>), and tornados (meteorological events), which, among other things, are perceptually similar in displaying the same kind of typical whirling movements (van Brakel 1991).
Red square, red flag, red cow	"Red square" is an example of a nonsubstance in the target article.
White	Will always reflect most of the incident light; there is no transparent white; the reflection of white objects (as contrasted with the luminance) is the same throughout changes in illumination (Westphal 1987, p. 12–39).
Thing, object, individual	Cf. Table 1.
<i>Lhenxa</i>	Kwakw'ala (Kwakiutl) word, the reference of which includes most green and yellows (Saunders & van Brakel 1996); both a yellow lemon and a green apple are <i>lhenxa</i> .
GRUE	Applies to all things examined before <i>t</i> just in case they are green but to other things just in case they are blue (Goodman 1972, p. 381).
UV-white, non-UV-white	Pigeons can be trained to sort systematically a pile of feathers that are equally white to human observation into two piles: the ones that reflect UV-light and the ones that do not (van Brakel 1992).
Prey-TWS, mate-TWS	Jumping spiders discriminate between prey and mate using four classes of photoreceptors (Nuboer 1986).

such as are in some way connected with whiteness." It also reminds one of Hacking (1991b, p. 115): "No one in the great tradition of natural kinds has seriously regarded the colours as natural kinds." On this reading "thing" is not a substance and "white" is excluded, "for there is nothing to be learned about it" (sect. 2, para. 8). However, this is not true. Various things can be learned about white (cf. Table 2), and it does not seem to fail any of Millikan's criteria for substances listed in Table 3.

Millikan takes Rosch's (1973; 1975) ideas about basic level concepts for granted. (This would also suggest that colors are substances, because colors are Rosch's prime examples of basic level concepts.) Millikan even goes so far as to refer to "those intermediate level categories such as *shoe* and *mouse* and *house* that children in all cultures learn first" (sect. 1, para. 3). Perhaps nowadays this is true for children of all cultures in the USA, but (and I feel a little embarrassed in pointing this out), it is not true for all children in all times and places.

The underlying problem is of course the question of what the right primitives are: the substances in terms of which all other nonsubstances can be analytically defined. The way of staging this problem also makes it a nonsensical problem. There is no space here to argue in general for this, so let me stick with examples. Why can't red square or red flag or red cow be primitive? Intuitive appeal seems to be made to the obviousness of what is primitive and what is complex. White, red, green, yellow, blue, black would then be primitive (as on Rosch's account); *lhenxa* or GRUE would not be; mouse, and perhaps vulture, would be primitive, and *vovetas* not (see Table 2 for descriptions of the "weird" categories). The only reason I know of to support what is primitive and what not is an appeal to the superiority of twentieth-century English, and the way the primitives of this language constrain psychological and linguistic research (Saunders & van Brakel 1997). Conceptual tracking abilities and the other characteristics listed in Table 3 work equally well for *lhenxa* and yellow.

In response to Goodman's GRUE, Chomsky has said that "every language learner (in fact, every mouse, chimpanzee, etc.) uses green rather than grue as a basis for generalization," as quoted in Goodman (1972, p. 78). But Goodman is surely right to say that speakers accustomed to projecting "grue" rather than "green"

would be equally confident that animals use grue rather than green as a basis for generalization. This may sound counterintuitive, but then using *lhenxa*, or *vovetas* will also sound counterintuitive to most readers of this commentary, whereas using *mouse* or *house* may seem counterintuitive to some nonreaders of this commentary.

Table 3 (van Brakel). *Characteristics of (the concept of) a substance*

Location in target article	Quote
Sect. 2, penultimate para.	A "substance" is something about which one can learn from one encounter things to apply on other occasions <i>where this possibility is not coincidental but grounded</i> . That is, there is an explanation or cause of the samenesses.
Sect. 3, para. 1	The "concept" of a substance . . . is the capacity to represent the substance in thought for the purpose of information gathering and storage, inference, and ultimately guidance of action.
Sect. 4, para. 7	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.
Sect. 5, para. 12	The cognitive systems are designed by evolution and tuned by experience to find real world substances, not random logically possible ones.
Sect. 2, para. 7	It is not a matter of logic that these things will not vary from meeting to meeting.

Consider also the jumping spider (*Salticidae*). It can differentiate prey and mate, because four classes of photoreceptors in its frontal eyes respond differently to prey and mate (Nuboer 1986). Let us call this discrimination capacity tetrachromatic wavelength sensitivity (TWS). Does the spider have the concepts of prey-TWS and mate-TWS? Probably not, because this is not a form of “practical knowledge [collected] over time of how to relate to specific stuffs, individuals, and real kinds” (sect. 1, para. 5). But then are prey-TWS and mate-TWS substances for human beings who have over time collected the practical knowledge for using the spider’s behavior as an indicator of the presence of prey-TWS (or spider-prey) and mate-TWS (or spider-mate)? Or are they nonsubstances, or what?

And what about pigeons? They are born with the potential capacity to discriminate UV-white, and non-UV-white. Assume they do not use this capacity automatically, but are trained by their parents to use this capacity (e.g., by learning to avoid some but not all white-feathered birds). Do they have the concepts of the substances UV-white and non-UV-white? On the other hand, it is sometimes said that human babies are born with the capacity to recognize a human face (on first sight, so to say). Are they born with the concept of a human face? Or, if they are born with no concepts, what is the sequence of learning events? Millikan does not say which concepts underlie the infant’s capacity to recognize its mother. Is it the concept “mother” in the sense of “whoever cares for me”? Should the “whoever cares for me” show his/her face to the infant (preferably?, necessarily?); should the accompanying smells, sounds, and touchings be within a certain range? The only possibility of a primitive substance in this context might perhaps be something like “the genetic mother” (not necessarily a human being anymore). But “genetic” is not a concept that can be innate or developed prior to language by infants.

I have no problem with Millikan following Quine and saying that at some point there is no distinction between “more Mama, more milk, and more mouse,” but I do not see why this list cannot be extended by adding “more white, more red cows, more things, more *voetas*, more *thenxa*, more GRUE,” and even “more UV-white, more prey-TWS,” and so on. In the right circumstances, given the right training, all characteristics listed in Table 3 work equally well for all of them.

I may have misunderstood much of what Millikan proposes, and therefore I may have quoted her out of context. Perhaps the best reply Millikan could give would be to provide a table with typical examples of substances on the left and typical nonsubstances on the right, including all examples from my Table 1, and perhaps also the following candidates – human face, stone, hawk, washing, quadruplet, electricity, being an electron or a positron, dead person, multiple sclerosis, God, *Kwoth*, gene, science, muddle. Let us take for granted that GRUE is, for whatever reason, not a substance. But what then about the others?

Words are invitations to learn about categories

Sandra Waxman^a and William Thompson^b

^aDepartment of Psychology, Northwestern University, Evanston, IL 60201
s-waxman@nwu.edu; ^bDepartment of Linguistics, Northwestern University,
Evanston, IL 60201 wkt@nwu.edu

Abstract: Evidence from language acquisition suggests that words are powerful mechanisms in the acquisition of substance concepts. Infants initially approach language with the general expectation that words refer to real kinds, regardless of grammatical cues to the contrary.

Millikan’s insightful target article brings to the foreground several fundamental issues regarding human concepts. Our goal is to marshal recent evidence from infants and young children to amplify the potential of Millikan’s nondescriptive account and to

elaborate upon the powerful contributions of naming in the acquisition of substance concepts. We focus primarily on categories of objects (real kinds, in Millikan’s terminology).

Under the general rubric of “substance concepts,” Millikan includes three kinds of kinds (real kinds, stuffs, and individuals). It is interesting that in English, each of these kinds of kinds are realized as particular subclasses of the grammatical category noun (count nouns, mass nouns, and proper nouns, respectively).

It appears, however, that among these three kinds of kinds, Millikan’s real kinds, or what Locke termed “sortals,” enjoy a special conceptual and linguistic status. Only sortal concepts (such as “table”) provide criteria of individuation and criteria of identity. Concepts of stuffs (such as “gold”) and concepts of individuals (such as “Mama”) do not (Lowe 1989; Macnamara 1982). Sortals also enjoy a special status in language acquisition. Millikan notes that substance terms constitute a major proportion of the early lexicon. Even more to the point, these early substance terms primarily denote sortals, corresponding to concepts denoted by count nouns in the adult language. Moreover, infant word learners seem to be biased toward interpreting most nominal terms (and adjectives as well) as denoting the sortal concept that includes the named object. Indeed, they appear to favor such interpretations over any other alternatives, including concepts of stuff or individuals (Hall 1993; Soja et al. 1991; Waxman & Hall 1993).

This is related to recent experimental evidence suggesting that at the onset of acquisition, infants may harbor what is at first a general expectation that a novel word (independent of grammatical form), applied ostensibly to an individual object, will refer to commonalities among objects and will therefore support the establishment of sortal concepts (Waxman & Markow 1995). More finely tuned linkages between specific grammatical forms and specific types of meaning (e.g., that counts nouns denote object categories, mass nouns denote stuff) emerge later. These more precise mappings are malleable; they are shaped by language-specific experience. This developmental finding fits well with the cross-linguistic observation that the lines of demarcation among distinct grammatical forms are drawn at different points in different languages (Choi & Bowerman 1991; Imai & Gentner 1993; Waxman et al. 1997).

Another appealing aspect of Millikan’s nondescriptive account is that it can accommodate potentially large changes in the “meaning” underlying a category or a name without requiring that there be radical changes in extension. This is important, because adults typically introduce novel words ostensibly (“Look at the zebra”). Adults also tend to correct errors in extension (“That’s not a horse, it is a zebra”). Although tutorials like these may help to delimit the boundaries of extension, they provide no evidence of the underlying intension or meaning. Children must therefore arrive at (most of) the deeper characteristics of substance concepts without the benefit of explicit tutorials (Waxman, in press). We suspect that their ability to do so is, at least in part, a consequence of naming.

This brings us to our final point. Millikan suggests that language serves as just another source of evidence about objects and object categories. Our view is quite different. We suggest that words serve as invitations to form (or, in the spirit of Millikan’s realist approach, to “learn about”) categories. Novel words direct infants’ attention toward commonalities (and differences) among objects and in this way promote object categorization. Providing a common name for a set of disparate objects (e.g., animals) promotes comparison among objects, allowing infants to notice deeper and subtler commonalities among them. This powerful influence of naming is especially clear in the early acquisition of object categories at nonbasic levels (Waxman & Markow 1995). We, therefore, see naming as central to the enterprise of discovering the important nonobvious commonalities that characterize our most powerful and inclusive categories of objects. Millikan’s assertion that language is key in acquisition of substance concepts rests comfortably with this position.

Concepts are not beliefs, but having concepts is having beliefs

Fei Xu, Joshua B. Tenenbaum, and Cristina M. Sorrentino

Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139. fei@psyche.mit.edu; jbt@psyche.mit.edu; cristina@psyche.mit.edu

Abstract: We applaud Millikan's psychologically plausible version of the causal theory of reference. Her proposal offers a significant clarification of the much-debated relation between concepts and beliefs, and suggests positive directions for future empirical studies of conceptual development. However, Millikan's revision of the causal theory may leave us with no generally satisfying account of concept individuation in the mind.

On Millikan's view, substance concepts are pointers to natural units of the world and "having a concept" means being able to use such a pointer to reidentify reliably instances of a substance. Treating concept possession as an ability of the mind helps to clarify the question of how beliefs are related to concepts. Beliefs are not constitutive of concepts, although *having* certain beliefs may constitute having certain concepts if those beliefs form the basis of the ability to reidentify. For example, knowledge that Mama has brown eyes, or that milk is white, along with many other beliefs about Mama and milk, forms the basis for the ability to reidentify Mama or milk on disparate occasions.

The failure to make the distinction between what a concept is and what it is to have a concept is one of the major sources of philosophical discontent with the cognitive science of concepts (Fodor 1995). Making this distinction allows us to stop arguing about whether children and nonhuman animals have the same concepts as human adults. To the extent that children and nonhuman animals have the same reidentification ability, they have the same concepts, albeit different "conceptions" (i.e., different beliefs). Nonetheless, although beliefs are not constitutive of concepts, we cannot study concept possession without studying beliefs. Millikan's view thus calls for a reinterpretation of existing research on infant concept possession.

The psychological investigation of infant concept possession has taken one of two forms. Some have argued that infants as young as 3 to 4 months have bona fide representations of substances such as HORSE or CAT (e.g., Quinn & Eimas 1993). That is, infants at this age will habituate to members of a particular category, such as horses, if shown several pictures of horses. These infants will then show a visual preference for an exemplar of a new category, such as a cat, relative to a new exemplar from the old category, namely, another horse. Others have argued that pre-linguistic infants younger than about 12 months do not conceptualize cups and balls as sortal concepts, which provide criteria for individuation and identity (Xu & Carey 1996); sortal concepts are considered necessary for acquiring count nouns such as "cup" and "ball" (Wiggins 1980).

Inspired by Millikan's new view of concepts, we suggest that these two empirical tests of concept possession are too weak and too strong, respectively. Having a concept is the ability to reidentify members of the substance on different occasions so that, crucially, beliefs acquired on different occasions are about the same set of entities. The infant habituation studies suggest that infants have a perceptual system that is similar to that of adults, so that infants carve up the world in more or less the same way adults do. However, these studies do not address the question of whether infants would acquire beliefs based on these categories. In fact, Mandler and McDonough (1996) have suggested that infants start by basing their inductions on more global categories such as ANIMAL and VEHICLE.

Although the infant habituation studies are too weak as a test of concept possession, the individuation and identity studies are too strong. These studies suggest that cups and balls are not conceptualized as individuated entities qua CUP and BALL by young infants. However, as Millikan explicitly says, it is not necessary for the child to conceptualize mice as an individuated entities to have

the concept MOUSE. That is, it is not necessary for the infant to have the sortal concept CUP to gather information about cups. All that is needed is the ability to reidentify cups or cupness. Reidentification without individuation is possible for infants with respect to count nouns, just as reidentification without individuation is possible with mass nouns. Adults can accumulate beliefs about milk without individuating milk every time they encounter it. Millikan's analysis of concept possession calls for new empirical research aimed at showing that prelinguistic infants can acquire beliefs about the same set of entities on different occasions.

We are encouraged by these possibilities for new empirical work. However, if beliefs are an integral part of the reidentification procedure, this raises problems for concept individuation in the mind. Millikan asserts that a concept is "a concept of A, rather than B, not because the thinker will always succeed in reidentifying A, never confusing it with B," but because of two factors: (1) "A is what the thinker has been conceptually . . . tracking and picking up information about", and (2) "the concept has been tuned to its present accuracy by causal interaction with members of A's specific domain" (sect. 5, para. 14). The second of these factors is the standard causal theory, in which reference is the result of a unidirectional causal link from the world to the mind. The first factor, in which the mind now plays an essential role in concept individuation, is new in Millikan's proposal and necessary under her definition of "substance," which includes many concepts that can apply to any one experience. Whether thinkers in the presence of Fido are accumulating beliefs about Fido, as opposed to dogs, fur, or bone depends only on whether they are "conceptually tracking" Fido.

But how can we know whether individuals are conceptually tracking substance A, unless we observe that they are by-and-large successfully reidentifying A? Millikan denies that consistent reidentification is the criterion for concept individuation, and instead appeals to causal links to the world filtered through a conceptual tracking mechanism. Yet she provides no diagnostic means for successful tracking other than reidentification. The problem of individuating concepts has been replaced by the problem of individuating conceptual tracks.

In conclusion, Millikan's theory clarifies the relation between concepts and beliefs in a way that suggests useful avenues for empirical studies of concept possession in infants. However, without an empirically viable criterion for concept individuation, the attempt to bridge the psychology and philosophy of concepts remains incomplete.

Author's Response

Words, concepts, and entities: With enemies like these, I don't need friends

Ruth Garrett Millikan

Department of Philosophy, University of Connecticut, Storrs, CT 06269-2054
millikan@uconnvm.uconn.edu

Abstract: A number of clarifications of the target article and some corrections are made. I clarify which concepts the thesis was intended to be about, what "descriptionism" means, the difference between "concepts" and "conceptions," and why extensions are not determined by conceptions. I clarify the meaning of "substances," how one knows what inductions to project over them, the connection with "basic level categories," how it is determined what substance a given substance concept is of, how equivocation in concepts occurs, and the role of language in the conception of substances. Finally, I clarify exactly why I said that concepts of individuals, real kinds, and stuffs have "a common structure,"

showing that, rightly understood, this view is not in conflict with data on infant concepts.

Writing a paper for people in other fields and trying to interpret their reactions is a deep study in communication. There is a quip in Italian: "With friends like you, I don't need enemies." A number of excellent commentaries evoke the converse of this thought (e.g., **Boyer, Müller & Kelter**, much of **Keil**, last paragraph of **Waxman & Thompson**), treating as objections claims or suggestions I had tried to make myself. Then there is warmly proclaimed agreement that I do not completely understand (**ter Meulen**). And there are the very many objections based on misunderstandings of what I intended, whatever I managed to say. Indeed, disappointing perhaps to those spoiling for a fight, there is not all that much here with which I fully disagree. And happily, in the midst of all this communicative confusion, very many of the right substantive questions seem to me to have been raised, offering an invaluable opportunity to clarify and extend my thought as well as my exposition. I am grateful to the commentators and for the *BBS* medium.

I will begin with clarifications that concern the general project of the target article, and some of the terminology. Then I will take up more specific issues.

R1. The overall program and some comments on terminology

In *Philosophical Investigations*, Wittgenstein (1953) said: "Think of the tools in a toolbox: there is a hammer, pliers, a saw, a screwdriver, a glue pot, nails and screws – The functions of words are as diverse as the functions of these objects" (para. 11). There is a use of the word "concept" – indeed, Wittgenstein helped to found this tradition – that equates a concept with whatever one has to learn to use a certain word correctly. So one can talk about the concept *or* and the concepts *of*, *hurrah*, and *the* and also the concepts *because*, *necessarily*, *ouch*, *good*, *true*, *two*, *exists*, and *is*. One can do that, but then remember Wittgenstein's warning. There will be little or nothing to be said in common about any two of these various "concepts." One must not expect a theory of how the tape measure works to double as a theory of how the glue works.

I have proposed a thesis about the nature of one and only one kind of concept, namely, concepts of what I have called "substances." It is legitimate to ask me to be more exact about what I mean by a "substance," as **van Brakel, Carlson, Nelson**, and **Mandler** (my "natural units in nature" are substances) have done. But it is not legitimate to ask me about any other kind of concepts (**Hampton, Hauser & Fitch, Nelson, Perner**). (**Allen** has other work of mine in mind in which I make more general claims about empirical concepts. I am sorry there is not room to explore those issues here.) By no means am I claiming that all concepts are substance concepts. I do not claim, for example, that no concepts are classifying concepts (**Gauker**). But I will try later to explain *some* relations I think substances and substance concepts have to *some* other ontological categories and *some* other psychological abilities. The world and thought are both vastly complicated. There is no Gordian knot waiting to be cut.

Equally important, the theory of substance concepts that I have proposed is not in the first instance about *words* for substances (**Franks & Braisby, Mandler, Murphy**,

Schröder). Rather, it belongs to the theory of cognition, in exactly the same way that theories of perception do. Thus **Allen** rightly asks about the substance concepts that languageless animals have, and **Cangelosi & Parisi** would even supply computers with substance concepts. A reasonable comparison might be made between the proposal made here and Marr's (1982) proposed first level of analysis in the theory of vision. I have attempted a task analysis for substance concepts, a description of what their function is, and why we need to have them. Marr claimed (rightly or wrongly) that the task of vision is to construct representations of three-dimensional objects starting from retinal images. I claim that the task of substance concepts is to enable us to reidentify substances in such a way that we can accumulate practical skills and theoretical knowledge about them and use what we have learned. I identify a substance concept with such an ability. (Incidentally, an ability is not a process, though of course abilities are usually implemented through processes – **Komatsu**.) Perhaps this ability will be best understood as equaling the ability to token and process or use a mental word appropriately such that it *constitutes* the thought of a substance (re: **Gauker, Komatsu, Mandler**) or perhaps an image of representation that is not as harsh as mental words is more suitable (Millikan 1997). What kind of entity does **Mandler** take a "concept" to be? It is because of its function that a mental representation is a representation at all, and a representation of this rather than that. It is concerning the functions of mental representations of substances, hence what it is that constitutes them as representations of substances, that I have made my initial claims.

This claim is on (something like) Marr's first level of analysis. Filling in the higher levels of analysis is primarily a job for psychologists: 1. How do we manage to perform this conceptual task; 2. and what are the details of the development of this sort of skill? I take it that all the traditional work in experimental psychology on substance concepts has been addressed implicitly to these higher levels, but work on these levels should be interpreted and given direction in light of an understanding of the function of substance concepts. I have made some tentative suggestions about how to do this.

An important one of these suggestions (as **ter Meulen** notes) is that multiple means are typically used conjointly and alternatively for identifying any given substance. This gives rise to the distinction between a "concept" and a "conception." The conception one has of a substance is the ways one knows to identify that substance plus the disposition to project certain kinds of invariances rather than others from one's experiences with it. I take it that what psychologists have typically studied is "conceptions" in this sense – the conceptions that people have of substances – and this is exactly what they should be studying. The concern that I am advocating the general abandonment of traditional ways of studying concepts (conceptions) is not warranted (**Hampton, Keil, Murphy, Schröder**). Nor is there a "paradox that what people know does not count as meaning" (Murphy). Without conceptions, no substances would be conceived of; "conception" is one of the things that "meaning" means. On the other hand, insofar as it has traditionally been assumed that for each thing that might be conceived or meant there corresponds but *one* possible conception (in my defined sense), or that for each univocal word in a language there corresponds just *one* conception, I

would have to disagree. There is no such thing as “the” conception of a substance nor is there “the” conception that corresponds to a public language term for a substance. Different people competently speaking the same language may have quite different – indeed, nonoverlapping – conceptions corresponding to the same substance term, and a single person may have quite different conceptions corresponding to the same substance at different times. (I am certainly not suggesting that there is no such thing as conceptual change, change in conceptions – **Keil**). This is why I say we should not be seduced into disputes about “what gets to count as ‘knowing the meaning’” of a certain substance term (**Murphy, Nelson**). And this is one of the reasons why the content of concepts “varies according to contextual and pragmatic constraints” (**Franks & Braisby**; cf. **Schröder**).

An embarrassment in terminology results from this divergence from a more traditional position, however. What I am calling a “conception” is in many ways much like what tradition has called a “concept.” But then tradition speaks of “the,” not “a concept dog,” and I think this is wrong if what is meant is a conception. I reserve the term “concept” then for what we do have only one of per person per substance, and only one of per word for a substance, namely, for abilities to recognize these substances and to know something of their potential for inductive use (**Murphy**). Or, because these abilities are what lend thoughts of substances their referential content (their representational values), we can also think of substance concepts as corresponding to mental representations of substances, say, to mental words for substances but to words *qua meaningful* (**Gauker, Komatsu**).

“Descriptionsm” is another word that has caused trouble (**Hampton, Hauser & Fitch, Keil, Komatsu, Mandler**). Some commentators seem to have assimilated my central claim against descriptionism with my quite independent claim that it is possible to identify substances without using mental descriptions of them, without using prior concepts of properties, and/or with my claim that recognizing a substance, even such a substance as mouse(kind), is not mentally describing something. Perhaps because of this assimilation, some thought that I was also claiming that substances are never identified by knowing descriptions of them. I am not certain exactly to which of these assimilations to attribute which problems seen by which of the above commentators, but let me clarify my position on each of these points (in reverse order, however).

Gopnik, and also **Xu et al.** have my position exactly right that early tracking mechanisms tend eventually to be replaced, certainly to be supplemented, by others (e.g., I no longer recognize my mother by smell) and may eventually be determined almost entirely by beliefs about the ontological structure of the world and, more mundanely, by beliefs about what *properties* tend to be diagnostic of what substances. (**Hauser & Fitch** and **Murphy**). Conceptual tracking is not equated with perceptual tracking (**Bloom?**); “reidentification” is not, in general, “cognitively impenetrable” (**Komatsu**); and certainly it is not claimed that infants recognize instances of kinds because they “look alike” (**Mandler**). It is just that some fundamental kinds of conceptual tracking begin with tracking perceptually, especially by recognizing object and property constancy. Indeed, the disposition to make an explicit inference, for example, from “the stuff has gone green” to “there’s copper

in it” (Quine 1960) is a paradigm conception of copper that helps to effect the conceptual tracking of that substance (recall my reference to the tool bag of tricks used by the chemist).

That recognizing a substance as such is not describing or classifying something follows from the claim that a substance does not equal merely a set of properties, nor is the concept of a substance shorthand for a set of concepts of properties. But “does not equal” does not of course imply “has no connection with.” Recognizing Mama by smell is certainly not classifying her nor is it conceiving of her as whatever bears that smell. It is more accurate to imagine it as a tokening of the mental term “Mama” in response to a smell. The thought is of Mama, not of smells, but it arises in response to a smell. Similarly, recognizing copper by the fact that the stuff has gone green is not conceiving of it as being just a green-turning thing. One tokens a mental term for copper in response to the knowledge that has gone green. What makes it a mental term for copper is, roughly, that it serves as a repository for incoming information about copper and that its tokenings are controlled by previous experience with copper, including explicit knowledge previously gained about copper.

In claiming that substances can be identified without necessarily using prior mental descriptions (concepts of properties; **Allen, Hampton, Hauser & Fitch**) I assumed that a property concept would involve more than the property’s causing a difference to one’s generalized response dispositions. In support, see **MacLennan**’s very instructive commentary on the neurological primacy of the concrete. (**Müller & Kelter** will like this, too, re: how “attributes are ‘constructed.’”) I had in mind that concepts of properties would involve representations of properties, which would imply the capacity to recombine these properties in thought with other subjects. Certainly a mere response to a presented property, such as a discriminating reflex response, requires no concepts (**Murphy**).

But none of these claims was what I had in mind in rejecting “descriptionism.” The descriptionist holds that the conception one has of a substance determines its extension. That is, the methods one uses for reidentifying, that is, for determining applications of a substance concept, determines what the concept is a concept of. I am fully in charge of the extensions of my substance concepts; whatever I am disposed to apply them to is what they are concepts of. I called this view “descriptionism,” because the extant views of substance concepts in the psychological literature uniformly take the conceptions we have of substances to be governed by descriptions. But because I spent considerable time arguing that certain of our most basic conceptions of substances are not governed by descriptions, this was inaccurate and misleading terminology. For I intended, equally, that conceptions of substances based merely on abilities to track them perceptually do not determine the extensions of those substances either. This is my “externalism.”

Franks & Braisby are mistaken in thinking that I use Putnam-Kripke style “counterfactual” arguments to defend my externalism. As **ter Meulen** remarks, I eschew possible world semantics and the Kripke interpretation of rigid designation. And as **Gopnik** remarks, I am not advocating an “internalist psychological essentialism.” That is, it is not my claim that substance concepts have the extensions they do because that is how people intend or believe them to

refer, or because that is the way people proceed with the use of words for substances. **Slater** has asked exactly the right question here, wanting to know, exactly why (given that it is not in accordance with the thinker's intentions) the reference of substance terms is determined as I say.

The answer is that it is not the purposes of individuals, but the biological functions of their inborn concept-tuning mechanisms that connects their substance concepts with certain extensions. No one supposes that the function of vision is determined by the intentions of the individuals who happen to have eyes. Similarly, the function of substance concepts is not determined by the intentions or dispositions of individuals who happen to have them. I have proposed a theory telling what the function of substance concepts is. It is their job to make contact with substances as these are objectively defined in nature. Only insofar as they manage to do this can they help us proceed with successful inductions. One could call whatever a certain conception happens to "corral" part of its "extension." 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 the conception to capture. A parallel would be to label whatever a frog happens to snap up with its tongue reflex – say, a beebee – as one of its "prey," and then be forced to coin another term to designate the things its reflex snap was designed to capture.

Franks & Braisby are right, however, that under stress, words can vacillate between continuing to be names for substances and taking on a more classificatory function (see also **Carlson**; more on this soon). They are also right, of course, to ask about the old chestnuts, empty names and informative identity statements. I refer them to Millikan (1984, Ch. 12; 1993, Ch. 14; 1997).

Gopnik and **Schröder** suggest that the interest displayed here in biological function or evolutionary purpose rather than current dispositions of the thinker may be irrelevant to psychological explanation. This opens a well-known can of worms about what it is psychology's job to explain. I discuss the issues fully in Millikan (1993, especially Chs. 7, 8, and 9), arguing for a widely ecological psychology that understands itself as a branch of biology, where the central concept is function understood as determined by natural selection.

Another confusing term seems to have been "pointing," which was taken by a number of commentators far more literally than I intended. "Pointing" is a metaphor having nothing to do with ostension (**Franks & Braisby, Nelson**), and not implying inarticulateness or lack of descriptive "content" in the conception that does the pointing (**Keil**). Nor are substance concepts literally "indexical" – that was a rather unfortunate metaphor of Putnam's. The idea is just that the relation of the conception to the extension of the concept is not logical but, in the broad sense, historical. It is a causal-order relation, a relation in the actual space-time world. (If I understand them correctly, **Xu et al.** have asked me to explain precisely what relation. I will get to that.)

R2. What are substances? The ontology

The ontology is supported by arguments in (Millikan 1984, Chs. 14–17). Here I make only assertions.

There is no single set of ontological "elements," no unique way of carving the ontology of the world, but a variety of basic patterns to be discovered there. The cate-

gory of substances, as I have defined it, is at root an epistemological one, and **Ghiselin** is quite right that it cuts straight across many more familiar distinctions in ontology. All that is required for an entity to be a substance is that it be such that it can be encountered on different occasions and such that it will remain invariant in certain respects over these encounters, not by accident but for a cause or reason, that is, in accordance with some kind of natural necessity. Beethoven's Fifth has many properties that are more or less the same from performance to performance (you can recognize it and know what is coming next). Places have properties, many of which remain the same over time. Dinnertime and siesta time have pretty definite properties in many cultures. War has certain properties that remain the same over the ages. Squares and cubes of material are things one can learn to recognize and about which one can learn a number of stable things such as how they fit together, how they balance, that their sides, angles, and diagonals are equal, and so forth. As **Cangelosi & Parisi** remark (correcting me), white gets dirty easily and, I now add, shows up easily in dim light, stays cool in sunlight but also tends to blind us, and so on. (Note the naturalness of noun forms here: "a square," "a cube," "white gets. . .") (It is not clear that individual events or processes can be understood as substances – **Ghiselin, Nelson** – although one can of course encounter the same event from various perspectives, for example, filtered through the medium of records made in various humans' memories or other recording media.) What makes a substance a substance is that it can be appropriated by cognition for the grounded – not accidental – running of inductions, or projecting of invariants. This will be possible in different cases for very different reasons, the result of very different sorts of causes, which is, of course, exactly what interests me about substances. It is their variety, considered from other ontological perspectives, that makes it easy to overlook their similarity relative to the project of cognition.

Ghiselin objects to my treating biological species as kinds rather than individuals, thus "concealing" the fundamentally different causes of the properties held in common by members of species and members of natural kinds: "laws of nature on the one hand and history on the other." I say bravo to the distinction, which I have generalized in Millikan (forthcoming), where I discuss "historical kinds" as indeed special, important, and neglected. The majority of our everyday substance kinds are historical kinds, their members being alike, typically, not because of some common inner essence, but in part because some form of copying has been going on in what is, relevantly, the same ongoing historical environment: Beethoven's Fifth, architectural kinds, living species, social kinds, professional kinds, the most common artifact kinds, automobile models, and so forth (**Murphy, Nelson**; for Nelson, that is how doctors constitute a real kind. They are an actual-world group, not a set of possible properties in a set of possible worlds. That is why their attitudes and practices can be studied empirically). But that a distinction is important in many contexts does not show that a similarity is not important in others. Natural kinds, historical kinds, stuffs, and individuals are very different indeed, but they also have something in common.

Substances vary greatly both in the number of inductions they support, and in the reliability of these inductions. (The latter gives rise, I suppose, to typicality effects. It seems

natural that people should work with a stereotype taken from knowledge of the most stable properties of substances when asked to describe the substance, in making guesses about category membership, when asked to make inferences about unobserved members, and so forth – **Schröder**.) Thus there is no sharp boundary between what is and what is not a substance. Rather, some things are, as it were, better substances than others, some are worth understanding as substances, others are too marginal or uninteresting (**Cangelosi & Parisi, Müller & Kelter, Nelson, van Brakel**). What holds a substance together, making it more than a mere set of similar items or encounters, is that the uniformity of its properties over encounters is not a coincidence. There is a reason why different encounters with the substance yield results as uniform as they do. But many substances have vague boundaries, indeed, some shade at the edges, in one or more dimensions, into other substances. Then the concepts (and words) that denote them may be equally vague. White shades into black and water shades into mud; these substances have natural paradigms, not natural boundaries. It does not follow that you cannot learn stable things about each. When is it really a war? Who is really a member of the working class? In the latter cases, the principle or principles that cause or tend to cause the members to be alike catch up some members more squarely than others.

But the interesting question is not so much one of pinpointing the substances, but rather what it is for something, rightly or wrongly, productively or unproductively, to be understood as a substance. Red sulphur is not just sulphur that is red, but an allotrope of sulphur with its own suite of properties different from other forms of sulphur (cf. **Müller & Kelter** regarding bearded women). One might have merely a classifying concept of red sulphur, however, capturing exactly the same extension, understanding it as just sulphur that is also red. Alternatively, one might have a concept of red sulphur as a substance, without knowing that it is always red or that it is a form of sulphur. One might use quite different means to identify this same substance. Concepts that classify are analytical concepts. They are composed of conjunctions (or disjunctions or other functions) of prior concepts. Substance concepts are synthetic. They may rest wholly or in part on prior concepts used in the process of identifying, but they are not equivalent to any mere function of prior concepts. The 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. (For **Hauser & Fitch**: this is one reason ants do not have substance concepts of dead ants. But mother vervets may well have substance concepts of their infants.) It is as if a substance concept made an inarticulate claim that there is some substance out there that it is hooked into.

Identifying and classifying are different things to do and they have different purposes (as discussed in the target article). Still, concepts can act partly as identifiers and partly as classifiers, or they can vacillate between these two functions. Any substance concept or term can be used for purposes of classification, and where substance boundaries are vague in nature, the purposes of classification may be served by drawing artificial boundaries around the extensions of these substances. For certain purposes, what counts as war and what counts as the working class may be quite sharply but artificially defined. Also, when confidence

is lost in the reality of a substance or in the univocity of a substance term, it may begin to be used strictly as a classifier (**Carlson, Franks & Braisby**). This is one reason not to attempt a list of words for substances (**Nelson, van Brakel**).

R3. Substance templates

A substance concept anticipates the validity of certain kinds of inductions. But how can one know ahead of time what kinds of inductions may hold? How does the child know, for example, to expect different constancies in a new uncle and a new piece of furniture? This must be done by having a grasp of more general categories within which substances can fall, the member substances having determinables in common. Determinables are not (determinate) properties like *red* or *square*, but rather disjunctions of contrary properties like *colored* (equals red or blue or green or . . .), and *shaped* (equals square or triangular or circular or . . .). I will call categories of substances that correspond in this way to sets of determinables “substance templates.” It is possible that *physical object* is a pure substance template. To be a physical object in the broadest sense, a thing need have no particular determinate properties, but it has to have some mass, some position and velocity at each time, some extension, some charge. With rare exceptions, however, concepts of substance templates are not pure. They capture substances that bring substance templates with them.

Animal and *vehicle* are such substances (**Mandler**). There is very little to be learned about either of these as such. What is most interesting about animals, for example, is that they divide into species, and that roughly the same sorts of questions can be asked about each of these species, and answered once and for all after one or a few observations. The main interest of the category *animal* is as a substance template. Because *animal* is not something there is much to find out about, there is also not much to say about it, and it is not surprising that the word “animal” enters the child’s vocabulary rather late. But because recognizing the substance template *animal* is crucial to learning about the various species of animals, it is equally unsurprising that animals might be recognized as such very early. Indeed, as **Boyer** suggests, the ability to track animals conceptually may have a strong boost from endogenous factors. As **Xu et al.** suggest, however, there is no reason to suppose that the infant’s differential response to animals indicates a substance concept of animal – no reason to suppose the infant is busy collecting information about the character of animals as such. (For Boyer, the understanding that “is not necessary to having the concept of a substance” is an explicit theory about what holds the substance together.)

A grasp of rough substance templates is a prerequisite to having genuine substance concepts. It is the requirement for substance concepts that one have an idea what to use them for. Paradigmatic substances are those that fall squarely under rich substance templates, such as animal, mineral, vegetable, and vehicle (for real kinds), and person (for individuals). Within each of these categories it is easy to find many substances, for each of which much the same questions can be asked: How big does it grow? How does it move? What organs are inside? Or, what is its melting point? Does it burn? How hard is it? How dense is it? Does

it corrode? Does it conduct? Or, what is it like inside? Where does it grow? Is it edible? – and so forth. I intended agreement here with **Boyer** that we may have built into us ways of conceptually tracking in a variety of different substance template domains, a boost toward conceptual tracking in each as well as a boost toward knowing why we should bother tracking. (That a particular method of conceptual tracking is always used for some particular purpose is what **Müller & Kelter** seem to have missed, plus the claim that boosts of the kind Boyer suggests must be built in.)

Real kinds nearly always bring with them substance templates covering their individual members. Thus the ability to identify cats is easily applied to discovering what sorts of questions can be asked about individual cats. What color is this cat (it will not change as with chameleons)? Is it tame or wild (not applicable to flies)? And does it have feline leukemia or a loud purr (not applicable to dogs)? For **Slater**: If “me” conceived as a substance begins with perceiving my body, then there is no particular problem about how a self-concept begins. It rides on grasp of the substance template for persons, who are tracked in the first instance by their bodies. The peculiarity is only how I recognize certain “inner” properties, because I have such a peculiar perspective on them. (I think I was wrong in Millikan [1984] that perceptual representations are indexical.)

R4. Basic level categories?

Making reference to the notion “basic level categories” in the target article was a mistake (though **van Brakel** may be less embarrassed to note that I gave references on the cross-cultural claims). Certainly, I did not mean that they are the only substances (**Carlson**), or that they are fundamentally different from other substances in some way (**Komatsu**). Nor was I thinking of basic-level categories as defined by “where perceptual similarity among exemplars is high” (Mandler & McDonough 1996), but had in mind more what Komatsu says about them. I was trying to make contact with current psychological terminology and theory, but Komatsu’s question shows how I failed. I hereby disclaim any opinion about why certain substance categories tend to be learned first cross-culturally. There may well be different reasons for different categories. Also, as noted above, there is reason to think that the most important substance template categories a child knows have no cause to manifest themselves in early speech. But more important, Komatsu is right that the change from thinking of all categories as classifiers to recognizing that many name substances radically challenges the more traditional framework in which many have theorized about “basic-level” categories, the framework that posits a “horizontal” and a “vertical” level of kind distinctions.

This traditional framework assumes a hierarchical structure among categories, so that they form a logical tree. This framework, the doctrine of “real definition,” or of natural ordering by genus and differentia, originated with Aristotle (which may be the best reason to believe it). Tree structure is what a good classification system must have, but it is not the structure of the logical space of substances nor of most of any subspaces of this space.

Consider stuffs on the one hand and people on the other. Clearly there is no way to hang these on the same logical

tree. They are neither beside one another (horizontal) under some higher substance, nor is one included in the other (vertical), nor is there some more inclusive substance covering them both. (Aristotle would have said they are both subsumed under *substance* and under *Being*, but *substance* is not a substance and neither is *Being*.) When we look within domains rather than across them, matters are no tidier. Susan is a mother and a professor and a diabetic. Each of these is a rough substance category, but there is no logical tree on which they all hang. Heated modern debates among biologists about principles of classification (phenetics, cladistics, evolutionary classification) reflect exactly this: there is no way to organize the substances that are of interest to the zoologist or botanist into a single hierarchy above the level of species – and in certain subdomains, even that level is problematic. The demand for biological taxonomy to settle on a single hierarchy is of course quite rational. A good classification system is needed for information storage and retrieval among the various biologists. The actual systems of classification used by biologists are compromises between good classification and respect for natural substance boundaries (compare Mayr 1981). In the natural domain of substances there is a confusing crisscrossing, every which way. On the other hand, the existence of real kinds that bring with them substance templates for their members does indeed impose a degree of hierarchy and order on the domain of substances.

Komatsu speaks of “sacrificing inductive richness” and of “variation in number of inferences supported” as one moves up a classification hierarchy to more inclusive categories. Indeed, categories lower down have all the properties of those above plus more, so there is more that is true of them, but how does this make them more “inductively rich”? Two different dimensions of induction are relevant here. There is the question of how many inductions, if one knew to venture them, would yield correct conclusions; and the question of how many one knows to venture. The more interesting question of inductive potential concerns how many determinables we *know* we can find stable values for, not how many stable properties the substance actually has. Hence the good substances are the ones for which there are rich, known, substance templates, for example, the chemical elements and compounds, the various living species, and also individual members of these species and most ordinary individual physical objects. These are things we know how to learn about without wasting time on dozens of observations verifying the stability of each trait. If one were to recognize only the lowest level substances, say, only the individual animals or the species, although it is true that these have the greatest number of properties, learning about these properties would be a hopelessly inefficient process. One would have to start all over with each individual object or species, exploring its individual features, with no contribution from prior knowledge of higher substances, either about its properties or its relevant determinables. The question of which level of categories are inductively the most fertile does not appear to be a well-formed one.

R5. Concept individuation (**Xu et al.**) and focusing reference (**Perner**)

There are multitudes of crisscrossing substances, very many more, surely, than those for which we have ideas. The ones that are picked up by thought and by language are those

that have properties of interest to us (**Livingston, Müller & Kelter, Mandler** on the “meaning” of concepts), but that they are interesting does nothing, of course, to make their status as substances less than fully objective (Livingston is lucid on this). The need for conceptual choice from among the multiplicity of substances does introduce another problem, however. If the substance I am thinking of is not distinguished by my having a disposition to track it correctly, if I can make mistakes in tracking what exactly determines or “individuates” what substance my concept is a concept of (**Xu et al.**)?

Many believe that what a substance concept is a concept of is determined by what fits the features or properties one represents it as having. Any more direct route from the mind to the substance would be mysterious. But what determines what features or properties are the ones one is representing? Surely no one is infallible at recognizing properties either, so how can prior thoughts of properties help us out here? A standard reply is that we recognize properties infallibly “in normal conditions.” How, then, do we define “normal conditions,” such that they are appropriately different for seeing the shapes of big things like mountains and small things like fleas, appropriately different for hearing loud sounds and soft, and different for seeing colors, tasting foods, and so forth (consider how tea tasters prepare themselves)? We must take care that “normal conditions” do not turn out to be just the conditions under which one perceives each of these various properties correctly, for that would be marching in place. On the other hand, if there is some noncircular way of defining “normal conditions” for perceptions of various properties, why is it that we cannot use the same technique to define “normal conditions” for the tracking of substances? The two problems are parallel.

Now biologists are usually concerned, first, with understanding normal function. They may be interested in disease or other abnormal functions, too, but these are defined relative to normal function. I take it that normal function itself, in this context, is best defined relative to a history of natural selection (Millikan 1984, Chs. 1 and 2; 1993, Chs. 1 and 2), but one can supply a favorite theory of normal function if there is one and it will serve the argument just as well. My suggestion is that cognitive psychologists, too, are, or should be, interested in normal function. For the most part, however, biological items have to be under certain conditions to perform normally. My preference is also to define normal conditions relative to selectionist history, as conditions under which that function was performed historically such that it has been selected for (Millikan 1984; 1993), but if a better definition exists, I have no objection. The point is that if we can give a definition of normal biological function and normal conditions for the performance of this or that function, we can also apply it to the performance of psychological functions, such as developing substance concepts and applying them.

Grant, then, that there is a normal way (or ways) to develop substance concepts (perhaps different for different substance domains). That is, assume that normal developmental psychology is a viable field. There will be a normal way or ways that children or adults first recognize the manifestations of a substance impinging on their perceptual organs, a normal way that they attempt to track that substance, and normal conditions for their success in tracking and in building conceptions adequate to that substance.

There will also be normal conditions for applying the concept so built, these being described, in part, relative to the conditions under which the concept was built. When everything goes exactly right, there will be no question what the concept is a concept of, even if there is a disposition to apply it incorrectly under conditions abnormal (specifically) for it. The problems arise when things do not go exactly right, when they deviate from the ideal.

Biological items, in general, are defined relative to an ideal. A diseased, damaged or malformed heart is a heart nonetheless because of the relations it has to hearts that perform normally. Once again, I leave the reader to define that relation, or to adopt the one I describe in Millikan (1984, Ch. 1). The important point here is that, having described how normal hearts are structured and how they function, it is of no interest to biologists how far away from that ideal a thing has to be before one stops calling it a “heart.” There are no exact borders of the substance *heart* in nature, and the biologist is concerned with nature. Similarly, I suggest, to press the question, in sufficiently abnormal cases, “But please, really, what is the referent of this person’s substance concept?” is useless.

On the other hand, there may be common and interesting abnormalities, divergences from the ideal, that are well worth studying. An obvious one is a substance concept that hovers between two or more substances, each of which has played a part in the normal development of a concept, but unfortunately got mixed together. Indeed, it is likely that normal development of many kinds of concepts involves a process of differentiating between substances originally confused together – **Perner** calls it a process of “focusing reference.” It is tempting to interpret much of the history of science as an attempt to focus reference, for example, distinguishing weight from mass and oxygen from other oxidizers. Where referential concepts have unfocused references they are equivocal. For example, if I should have twins confused together in my mind, thinking there is only one person out there and not two, my concept would be equivocal. I see no reason to suppose, however, that disambiguating my concept would require that I focus on *one specific* set of features defining each twin (Perner).

For **Xu et al.** and **Allen**: Does what my concept is about have to be any more “empirically viable” than whether I am really remembering something? Answers to both questions are objective but rooted in the past. For Allen: Assuming that baby vervet monkeys are designed by natural selection to develop, specifically, certain predator calls, there is no ambiguity in the baby vervet’s inept call. It is an immature signal for predators.

R6. Words and the depth of perception

Learning words for substances is in part a matter of focusing reference. Substances are tracked through words and in other ways. If a subject of information that arrives through language, tracked by a word manifesting it, is then merged under the same concept with some different subject of information arriving through other perceptual media, there will be equivocation in the resulting concept. We say in such cases that the person does not know the meaning of the word. The case would be exactly similar, however, had the subject mixed a person known only through phone calls with someone else known from glimpses at the beach. We could just as well say, using the same sense of “meaning,”

that we did not know the meaning of the voice over the phone. Thus, **Ghiselin's** child who calls the whole genus *Felis* by the name "kitty" does not yet know the meaning of "kitty," but it also has an equivocal concept. The child's word "kitty" hovers between referring to felines generally and house cats specifically. The child will be putting all information gleaned through language and specific to house cats in the same bin as information gleaned about tigers and lions at the zoo. I think that what **Keil** intended to say was that one can intend to refer to something that is completely different from that to which one actually refers. The paradigm case, I believe, would be one in which one's concept is very strongly and clearly channeled outside the linguistic medium to something other than what one's word generally carries information about in the public language. Still, there is a bit of ambiguity in such word usage, so long as the speaker is ready to confuse information gathered via the word with information about an outside referent – so long, that is, as the speaker did not just misspeak.

Ghiselin's child's conception of "kitty" is equivocal because part of it is channeled through the method of tracking that is understanding language. Because it is possible for a conception to be channeled completely through this method of tracking, it is possible to have a substance concept through nothing but a word plus a grasp of its substance template and enough relevant grammar. **Perner** finds this unintuitive, and I sympathize, but my point is that filling out of the concept into a more and more adequate one happens in degrees. There is no special thing that gets added at some later point that suddenly makes it into a "real concept." It can be filled out more; it can get better and better. But there is no magic moment when it has attained some essence required for true concepthood. That, as **Perner** notes, is what caused the analytic/synthetic distinction to die.

I did not mean to suggest that "language serves as just another source of evidence about objects and object categories" (**Waxman & Thompson**). I take it that new words serve in huge numbers as seed crystals around which fuller concepts are then quickly formed. That is why Helen Keller was, as she later described it, pretty much unable to *think* until Annie Sullivan taught her some language. And it is why there can be such differences between the concepts available in cultures not historically related.

The view of language proposed is quite different from Kripke's "causal chain of reference borrowing" (**Franks & Braisby**) and quite different from the view that a child or anyone else takes out loans on concepts knowing there are experts out there to pay up (**Keil**). both these images have someone out there who "really" has the concept, whereas the rest of us do not. But even if we soften this to just the claim that some people out there have (or had – consider our concept of Socrates) the concept in a way that was different because it was focused without any reliance on public language, the image is still wrong. Thus, **Gauker** wants to know how the "very content of a person's thought" could depend on "aspects of usage that that person has not personally detected," and **Hauser & Fitch** have me supposing that "language provides a vehicle for concept exchange." No. Let me try it again.

Children come into the world without any knowledge of how minds work or what goes on inside people when they speak. (Indeed, we ourselves seem to be a bit short on such knowledge). For children, language serves simply as an-

other medium through which to perceive the world, just as they perceive the world through their eyes without knowing anything about light, and through their sense of touch and smell without knowing anything about physical forces or chemicals. (**Keil** asserts that language does not transmit structural isomorphisms. Wittgenstein (1922) argued that it does. Compositional semantics assumes that it does. What is **Keil's** argument?)

How can that be, one might say, because Mamma's words are right here and the dog she talks about is way over there? Well, how does one perceive oneself in the mirror? What's funny about language, I have said, is that it does not show one's relation to the things one perceives through it. How can that be, one may ask, because a lot of what is said may not be true. Well, in water, oars look bent and the reflections of the trees show them moving in ripples. **Gendler** objects that "I am dying" uttered by a laughing 8-year-old does not present the "appearance" that someone is dying. Of course one is not tempted to believe it, nor is one tempted to believe that the oars are bent in the water.

But the main barrier lies here. The philosophical tradition, and the psychological tradition following after, resolutely hold that for each of the physical senses there is just one layer of the world that it perceives directly; all other layers are known only through inference. This is the premise I am denying. There is no single "given" layer of perception (again, see **MacLennan**). What one perceives depends on where one's mind is focused.

So my answer to **Scharer** is that it is neither words nor phoneme strings nor sounds that the young child perceives when learning language. It is, in the first instance, the world. (Nor do children learn many words by ostension. They learn them by hearing complete sentences containing them – e.g., Gleitman 1990; Grimshaw 1994; Pinker 1994). For **Mandler** (1996): I do not believe there is any such thing as a layer of "perceptual properties," and certainly there is no such thing as a "perceptual category." For **Gendler**: The opposite side of the coin of there being no perceptual level of givenness is that all levels of perception are "theory laden," that is, fallible – though "theory" is not analyzed here in traditional twentieth-century empiricist terms. There is no necessary involvement of deductive inference norms or inference dispositions in this kind of "theory."

R7. A common structure for concepts of substances?

As I interpret **Bloom**, **Keil**, and **Schröder**, all require a clearer statement of what "common structure" I intend.

First, let me be clear that my thesis is not Quine's (1960). Quine thinks that language both inspires and creates the principles that "individuate" the various substances. I think neither of these things. I used Quine's words "more Mama" in describing children's thoughts of their mothers. I should have just said "Mama again," for whatever Quine had in mind, I certainly did not intend to imply that a child cannot differentiate between individual objects and stuffs, and certainly did not mean that the child thinks individuals *are* stuffs! Begin with this comparison. The child differentiates between dogs and cats, but it does not follow that the child's concepts of these do not have a common structure. The child differentiates between animals and vehicles, but it does not follow that the child's concepts of things in these

two domains do not have a common structure. Similarly, children differentiate among individuals, stuffs, and real kinds, but it does not follow that their concepts of things in these domains do not have a common structure. What counts as a common structure obviously depends on one's interests.

So what was the point of my title? The fact is that tradition has implicitly claimed there is nothing common to the structures of concepts of individuals, kinds, and stuffs (let alone of "here's Beethoven's Fifth again" and "here's white again.") But my claim is that the most important fact about the structure of each of these concepts (insofar as it operates as a substance concept) is something they all have in common. Each contains some means or other of tracking its appointed substance and a grasp of how to project some of the invariants defining this substance to new encounters with it. This is the most important fact about the structure of these concepts because it defines their function, explains the reason we have them. Moreover, I am claiming, against a long tradition at least in philosophy (this being an interdisciplinary journal), that an "explicit understanding of the ontological principles that ground substance tracking" definitely is not "needed for a substance concept" (Keil). (Ghiselin: Notice that it was not until very recently that anyone had an explicit understanding of the ontological principles that ground tracking of the various biological species, and we still do not have many details in the case, for example, of asexual animals and easily hybridized plants.) Therefore I claimed that the concepts of individuals, of real kinds, and of stuffs do not have to differ from one another in the way many have assumed is essential.

We can contrast these similarities among substance concepts with some very important differences (Bloom). There are two important dimensions to any substance concept, first, a method of tracking and second, a projection of invariants. Children have, perhaps, a concept of Mama and also a concept of women. They use different methods to track these, and project different invariants over encounters with them. Tracking Mama is one of the means of tracking women: If it is Mama again, it is a woman again. But the concepts are entirely separate, not at all confused together. Similarly, knowing what to expect of a connected physical object and knowing to expect something different of a pile of sand shows that children are capable of distinguishing between the domains of application of certain substance templates. Their methods of conceptual tracking are surely different for objects and stuffs. Similarly, their methods of tracking *cat* may allow them to generalize from the cat now seen on the left to a cat now seen on the right, whereas their methods of tracking individuals, hence Tabby, do not. These are real differences. But they do not erase the sameness in which I was interested.

Xu et al. reflect on experiments suggesting that infants less than a year old do not have "sortal concepts," for example, of *cup* and *ball*, proposing that this may not show a lack of concepts of these as substances. This seems right to me, but I would rather express it by saying they do not yet have the idea of a certain substance template covering a much wider range of objects than cups and balls. The idea that certain concepts bring with them criteria for individuation and identity (see also Waxman & Thompson) should be suspect to a realist, for it suggests (as Wiggins [1980] intended it to) that the mind – not the world – decides such things.

References

Letters *a* and *r* appearing before authors' initials refer to target article and response respectively.

- Abbot, V., Black, J. B. & Smith, E. E. (1985) The representation of scripts in memory. *Journal of Memory and Language* 24:179–99. [aRGM]
- Allen, C. & Hauser, M. D. (1991) Concept attribution in nonhuman animals: Theoretical and methodological problems in ascribing complex mental processes. *Philosophy of Science* 58:221–40. [MDH]
- Allen, C. & Sidel, E. (1997) The evolution of reference. In: *The evolution of mind*, ed. D. D. Cummins & C. Allen. Oxford University Press. [CA]
- Angelin, J. M. (1977) Classifiers. *Language* 53:285–311. [aRGM]
- Atran, S. (1989) Basic conceptual domains. *Mind and Language* 4(1/2): 7–16. [aRGM]
- Baillargeon, R., Spelke, L. & Wasserman, S. (1985) Object permanence in five-month-old infants. *Cognition* 20:191–208. [PBI]
- Barsalou, L.W. (1987) The instability of graded structure: Implications for the nature of concepts. In: *Concepts and conceptual development*, ed. U. Neisser. Cambridge University Press. [aRGM, JS]
- Bartlett, F. C. (1932) *Remembering*. Cambridge University Press. [aRGM]
- Bealer, G. (1987) The philosophical limits of scientific essentialism. In: *Metaphysics*, ed. J. Tomberlin, vol. 1 of *Philosophical perspectives*. Atascadero: Ridgeview. [BF]
- Billman, D. (1992) Modeling category learning and use. In: *Percepts, concepts, and categories; The representation and processing of information*, ed. B. Burns. North-Holland. [aRGM]
- Bloom, P. (1990) Syntactic distinctions in child language. *Journal of Child Language* 17:343–55. [PBI]
- (1996) Possible individuals in language and cognition. *Current Directions in Psychological Science* 5:90–94. [PBI]
- Bloom, P. & Kelemen, D. (1995) Syntactic cues in the acquisition of collective nouns. *Cognition* 56:1–30. [PBI]
- Boyd, R. (1979) Metaphor and theory change: What are metaphors a metaphor for? In: *Metaphor and thought*, ed. A. Ortony. Cambridge University Press. [CS]
- (1989) What realism implies and what it does not. *Dialectica* 43:5–29. [aRGM]
- (1991) Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies* 61:127–48. [aRGM]
- Braisby, N. R., Franks, B. & Hampton, J. A. (1996) Essentialism, word use, and concepts. *Cognition* 59:247–74. [BF]
- Bridgeman, B. (1982) Multiplexing in single cells of the alert monkey's visual cortex during brightness discrimination. *Neuropsychologia* 20:33–42. [BJM]
- Burge, T. (1979) Individualism and the mental. In: *Studies in metaphysics: Midwest studies in philosophy IV*, ed. P. French, T. Uehling, & H. Wettstein. University of Minnesota. [aRGM]
- (1982) Other bodies. In: *Thought and object*, ed. A. Woodfield. Clarendon. [aRGM]
- (1986) Individualism and psychology. *Philosophical Review* 95:3–45. [aRGM]
- Burns, B. ed. (1992) *Percepts, concepts, and categories; The representation and processing of information*. North-Holland. [aRGM]
- Byrnes, J. P. & Gelman, S. A. (1991) Perspectives on thought and language: Traditional and contemporary views. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- Calvert, G. A., Bullmore, E. T., Brammer, M. J., Campbell, R., Williams, S. C., McGuire, P. K., Woodruff, P. W., Iversen, S. D. & David, A. S. (1997) Activation of auditory cortex during silent lipreading. *Science* 276:593–96. [BJM]
- Cangelosi, A., Denaro, D. & Parisi, D. (1997) Language as an aid to categorization. Technical Report SNVA-9707, Institute of Psychology, National Research Council, Rome. [AC]
- Cangelosi, A. & Parisi, D. (submitted) The emergence of a "language" in an evolving population of neural networks. Paper presented at the Special Session on Modeling Language Evolution, 18th Cognitive Science Conference, San Diego, July 1996, and at the 1996 Annual Meeting of the Language Origins Society, Baltimore, July. [AC]
- Carey, S. (1985) *Conceptual change in children*. The MIT Press. [aRGM]
- Carey, S. & Gelman, R., eds. (1993) *The epigenesis of mind: Essays on biology and cognition*. Erlbaum. [aRGM]
- Carlson, G. (1980) *Reference to kinds in English*. Garland. [GC]
- Cheney, D. L. & Seyfarth, R. M. (1980) Vocal recognition in free-ranging vervet monkeys. *Animal Behaviour* 28:362–67. [MDH]
- (1990) *How monkeys see the world: Inside the mind of another species*. University of Chicago Press. [CA, MDH]

- Chiang, W.-C. & Wynn, K. (1997) Eight-month-olds reasoning about collections. Poster presented at the meeting of the Society for Research in Child Development, Washington, DC, April 4. [PBI]
- Choi, S., & Bowerman, M. (1991) Learning to express motion events in English and Korean: The influence of language-specific lexicalization patterns. *Cognition* 41:83–121. [SW]
- 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*. Center for the Study of Language and Information. [aRGM]
- Chomsky, N. (1995) Language and nature. *Mind* 104:1–61. [aRGM, JvB]
- Clark, E.V. (1991) Acquisitional principles in lexical development. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- Dickens, C. (1854/1955) *Hard times*, Oxford University Press. [TSG]
- Dodwell, P., Humphrey, K. & Muir, D. (1987) Shape and pattern perception. In: *Handbook of infant perception*, vol. 2 of *From perception to cognition*, ed. P. Salapatek & L. Cohen. Harcourt Brace Jovanovich. [aRGM]
- Dreyfus, H. L. (1979) *What computers can't do*. Harper & Row. [BJM]
- (1982) Introduction. In: *Husserl, intentionality and cognitive science*, ed. H. L. Dreyfus with H. Hall. MIT Press. [BJM]
- (1991) *Being-in-the-world: A commentary on Heidegger's Being and Time, Division I*. MIT Press. [BJM]
- Dromi, E. (1987) *Early lexical development*. Cambridge University Press. [aRGM]
- Evans, C. & Marler, P. (1995) Language and animal communication: Parallels and contrasts. In: *Comparative approaches to cognitive science*, ed. J.-A. Meyer & H. L. Roitblat. MIT Press. [CA]
- Fodor, J. A. (1972) Some reflections on L. S. Vygotsky's *Thought and language*. *Cognition* 1:83–95. [aRGM]
- (1994) Concepts: A potboiler. *Cognition* 50:95–114. [BF, FX]
- Franks, B. & Braisby, N. R. (submitted) On the relation between word use and concepts. Perspectives. [BF]
- Frege, G. (1892/1966) On sense and reference. In: *Translations from the philosophical writings of Gottlob Frege*, ed. P. Geach & M. Black. Blackwell. [BF]
- Fromkin, V. A. (1971) The nonanomalous nature of anomalous utterances. *Language* 47:27–52. [JS]
- Fumerton, R. (1989) Russelling causal theories. In: *Rereading Russell: Essays in Bertrand Russell's metaphysics and epistemology*, ed. C. W. Savage & C. A. Anderson, vol. 12 of *Minnesota studies in the philosophy of science*. University of Minnesota Press. [aRGM]
- Gage, J. (1993) *Color and culture: Practice and meaning from antiquity to abstraction*. Little, Brown. [BJM]
- Gallistel, C. R. (1990) *The organization of learning*. MIT Press. [PBI]
- Gallistel, C., Brown, A., Carey, S., Gelman, R. & Keil, F. (1993) Lessons from animal learning for the study of cognitive development. In: *The epigenesis of mind: Essays on biology and cognition*, ed. S. Carey & R. Gelman. Erlbaum. [aRGM]
- Gelman, S. A. (1988a) Children's expectations concerning natural kind categories. *Human Development* 38:213–44. [aRGM]
- (1988b) The development of induction within natural kind and artifact categories. *Cognitive Psychology* 20:65–95. [aRGM]
- Gelman, S. A. & Byrnes, J. P., eds. (1991) *Perspectives on language and thought*. Cambridge University Press. [aRGM]
- Gelman, S. A. & Coley, J. D. (1991) Language and categorization: The acquisition of natural kind terms. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- Gelman, S. A. & Markman, E. M. (1987) Young children's inductions from natural kinds: The role of categories and appearances. *Child Development* 58:1532–41. [aRGM]
- Gelman, S. A. & Wellman, H. M. (1991) Insides and essences: Early understanding of the non-obvious. *Cognition* 38:213–44. [AG]
- Gentner, D. (1982) Why nouns are learned before verbs: Linguistic relativity versus natural partitioning. In: *Language development*, ed. S. Kuczaj, vol. 2 of *Language, thought and culture*. Erlbaum. [aRGM]
- Ghiselin, M. (1974) A radical solution to the species problem. *Systematic Zoology* 23:536–44. [aRGM]
- (1981) Categories, life and thinking. *Behavioral and Brain Sciences* 4:269–83. [aRGM, MTG]
- (1997) *Metaphysics and the origin of species*. State University of New York Press. [MTG]
- Gilbert, D. (1993) The assent of man: Mental representation and the control of belief. In: *Handbook of mental control*, ed. D. M. Wegner & J. W. Pennebaker. Prentice-Hall. [aRGM, TSG]
- Gleitman, L. R. (1990) The structural sources of verb meanings. *Language Acquisition* 1:1–55. [rRGM]
- Goodman, N. (1972) *Problems and projects*. Bobbs-Merrill. [JvB]
- Gopnik, A. & Meltzoff, A. (1993) Words and thoughts in infancy: The specificity hypothesis and the development of categorization and naming. In: *Advances in infancy research*, ed. C. Rovee-Collier & L. Lipsitt. Ablex. [aRGM]
- (1996) *Words, thoughts and theories*. MIT Press. [aRGM]
- Gopnik, A. & Wellman, H. M. (1994) The theory theory. In: *Mapping the mind*, ed. Hirschfeld & S. A. Gelman. Cambridge University Press. [TSG]
- Grimshaw, J. (1991) Lexical reconciliation. In: *The acquisition of the lexicon*, ed. L. Gleitman & B. Landau. MIT Press. [rRGM]
- Hacking, I. A. (1991a) On Boyd. *Philosophical Studies* 61:149–54. [aRGM]
- (1991b) A tradition of natural kinds. *Philosophical Studies* 61:109–26. [JvB]
- Hall, D. G. (1993) Basic-level individuals. *Cognition* 48:199–221. [SW]
- Hall, G. (1993) Assumptions about word-meaning: Individuation and basic-level kinds. *Child Development* 64:1550–70. [PBo]
- Hardcastle, V. G. (1994) The image of observables. *The British Journal for the Philosophy of Science* 4:585–97. [JvB]
- Hauser, M. D. & Carey, S. (in press) Building a cognitive creature from a set of primitives: Evolutionary and developmental insights. In: *The evolution of mind*, ed. D. Cummins & C. Allen. Oxford University Press. [MDH]
- Hull, D. L. (1978) A matter of individuality. *Philosophy of Science* 45:335–360. (Reprinted in Sober, E., ed. (1994) *Conceptual issues in evolutionary biology*. MIT Press.) [aRGM]
- Huntley-Fenner, G. & Carey, S. (1995) Individuation of objects and portions of nonsolid substances. Poster presented at the meeting of the Society for Research in Child Development, Indianapolis, IN, March 30. [PBI]
- Imai, M. & Gentner, D. (1993) Linguistic relativity vs. universal ontology: Cross-linguistic studies of the object/substance distinction. *Proceedings of the Chicago Linguistic Society*. [SW]
- Ingram, D. (1989) *First language acquisition: Method, description and explanation*. Cambridge University Press. [aRGM]
- John, O. (1985) 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. [aRGM]
- Johnson, M. H., Dziurawiec, S., Ellis, H. D. & Morton, J. (1991) Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition* 40:1–21. [aRGM]
- Johnson-Laird, P. N. (1983) *Mental models*. Harvard University Press. [aRGM]
- Kant, I. (1787/1929) *Critique of pure reason*. Trans. N. Kemp Smith. Macmillan. [TSG]
- Karmiloff-Smith, A. (1992) *Beyond modularity*. MIT Press. [MDH]
- Katz, J. (1972) *Semantic theory*. Harper & Row. [aRGM]
- Katz, J. & Fodor, J. A. (1963) The structure of a semantic theory. *Language* 39:170–210. [aRGM]
- Katz, N., Baker, E. & Macnamara, J. (1974) What's in a name? A study of how children learn common and proper names. *Child Development* 45:469–73. [PBI]
- Keil, F. C. (1979) *Semantic and conceptual development: An ontological perspective*. Harvard University Press. [aRGM]
- (1983) Semantic inferences and the acquisition of word meaning. In: *Concept development and the development of word meaning*, ed. T. B. Seiler & W. Wannemacher. Springer-Verlag. [aRGM]
- (1987) Conceptual development and category structure. In: *Concepts and conceptual development: Ecological and intellectual factors in categorization*, ed. U. Neisser. Cambridge University Press. [aRGM]
- (1989) *Concept, kinds and cognitive development*. MIT Press. [aRGM, AG, LKK, JP]
- (1991) Theories, concepts, and the acquisition of word meaning. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- Kelley, D. (1984) A theory of abstraction. *Cognition and Brain Theory* 7:329–57. [KRL]
- Komatsu, L. K. (1992) Recent views of conceptual structure. *Psychological Bulletin* 112.3:500–26. [aRGM]
- Kornblith, H. (1993) *Inductive inference and its natural ground*. MIT Press. [aRGM]
- Kripke, S. (1972; 1980) *Naming and necessity*. Blackwell; Harvard University Press. [aRGM, GC, BF, AG, VCM]
- (1977) Speaker's reference and semantic reference. In: *Contemporary perspectives in the philosophy of language*, ed. P. A. French, D. Uehling, & H. Wettstein. University of Minnesota Press. [BF]
- Kuhn, T. S. (1962) *The structure of scientific revolutions*. University of Chicago Press. [TSG]
- Lakoff, G. (1987) *Women, fire, and dangerous things*. Chicago University Press. [aRGM]
- Lakoff, G. & Johnson, M. (1980) *Metaphors we live by*. University of Chicago Press. [BJM]
- Lakatos, I. (1970) Falsification and the methodology of scientific research

- programs. In: *Criticism and the growth of knowledge*, ed. I. Lakatos & A. Musgrave. Cambridge University Press. [TSG]
- Livingston, K. R. & Andrews, J. A. (1995) On the interaction of prior knowledge and stimulus structure in category learning. *Quarterly Journal of Experimental Psychology* 48A:208–36. [KRL]
- Livingston, K. R., Andrews, J. A. & Harnad, S. (in press) Categorical perception effects induced by category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. [KRL]
- Lowe, E. J. (1989) Kinds of being: A study of individuation, identity and the logic of sortal terms. Blackwell. [JvB, SW]
- MacFarlane, A. (1977) *The psychology of childbirth*. Harvard University Press. [aRGM]
- MacLennan, B. J. (1994) Continuous computation and the emergence of the discrete. In: *Origins: Brain and self-organization*, ed. K. H. Pribram. Erlbaum. [BJM]
- (in press) Mixing memory and desire: Want and will in neural modeling. In: *The brain and values: Proceedings of the Fifth Appalachian Conference on Behavioral Neurodynamics*, ed. J. King & K. H. Pribram. Erlbaum. [BJM]
- Macnamara, J. (1986) A border dispute: *The place of logic in psychology*. MIT Press. [SW]
- Malt, B. C. (1995) Category coherence in cross-cultural perspective. *Cognitive Psychology* 29:85–148. [KRL]
- Malt, B. C. & Johnson, E. C. (1992) Do artifact concepts have cores? *Journal of Memory and Language* 31:195–217. [GLM]
- Mandler, J. M. (1992) How to build a baby II: Conceptual primitives. *Psychological Review* 99:587–604. [JMM]
- (1997) Development of categorization: Perceptual and conceptual categories. In: *Infant development: Recent advances*, ed. G. Bremner, A. Slater & C. Butterworth. Erlbaum; Psychological Press. [JMM, KN]
- Mandler, J. M. & McDonough, L. (1996) Drinking and driving don't mix: Inductive generalization in infancy. *Cognition* 59:307–35. [rRGM, JMM, FX]
- Markman, E. (1989) *Categorization and naming in children*. MIT Press. [aRGM]
- (1991) The whole-object, taxonomic and mutual exclusivity assumptions as initial constraints on word meaning. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- Marler, P. (1993) The instinct to learn. In: *The epigenesis of mind: Essays on biology and cognition*, ed. S. Carey & R. Gelman. Erlbaum. [aRGM]
- Marr, D. (1982) *Vision*. Freeman. [rRGM]
- Mayr, E. (1981) Biological classification: Toward a synthesis of opposing methodologies. *Science* 510–16. (Reprinted in Sober, E., ed. (1994) *Conceptual issues in evolutionary biology*, 2nd ed. MIT Press.) [rRGM]
- McNamara, T. P. & Miller, D. L. (1989) Attributes of theories of meaning. *Psychological Bulletin* 106:355–76. [JS]
- Medin, D. L. (1989) Concepts and conceptual structure. *American Psychologist* 44(12):1469–81. [aRGM]
- Medin, D. L., Goldstone, R. L. & Gentner, D. (1993) Respects for similarity. *Psychological Review* 100:254–78. [KRL, VCM]
- Medin, D. L., Lynch, E. B., Coley, J. D. & Atran, S. (1997) Categorization and reasoning among tree experts: Do all roads lead to Rome? *Cognitive Psychology* 32:49–96. [KRL]
- Medin, D. L. & Ortony, A. (1989) Psychological essentialism. In: *Similarity and analogical reasoning*, ed. S. Vosniadou & A. Ortony. Cambridge University Press. [aRGM]
- Medin, D. L. & Schaffer, M. M. (1978) Context theory of classification learning. *Psychological Review* 85:207–38. [aRGM]
- Medin, D. L. & Smith, E. E. (1981) Strategies and classification learning. *Journal of Experimental Psychology: Human Learning and Memory* 7:241–53. [aRGM]
- (1984) Concepts and concept formation. *Annual Review of Psychology* 35:113–38. [aRGM]
- Merriman, W. & M. Tomasello, eds. (1995) *Beyond names for things*. Erlbaum. [KN]
- Mervis, C. B. (1987) Child-basic categories and early lexical development. In: *Concepts and conceptual development*, ed. U. Neisser. Cambridge University Press. [aRGM]
- Mervis, C. B. Crisafi, M. A. (1982) Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development* 53:258–66. [aRGM]
- Meulen, A. ter (1981) An intensional logic for mass terms. *Philosophical Studies* 40:105–25. [AGBtm]
- (1994) Demonstrations, indications and experiments. *The Monist* 77:239–56. [AGBtm]
- (1995) Semantic constraints on type-shifting anaphora. In: *The generic book*, ed. G. Carlson & J. Pelletier. University of Chicago Press. [AGBtm]
- Mill, J. S. (1843) A system of logic. In: *Collected works of John Stuart Mill*, vol. 7, ed. J. M. Robertson. University of Toronto Press. [JvB]
- Miller, G. A. & Johnson-Laird, P. N. (1976) *Language and perception*. Cambridge University Press. [LKK]
- Millikan, R. (1984) *Language, thought and other biological categories*. MIT Press. [arRGM, VCM, CS]
- (1991) Perceptual content and Fregean myth. *Mind* 100.4:439–59. [aRGM]
- (1993a) *White Queen psychology and other essays for Alice*. MIT Press. [arRGM, CG, CS]
- (1993b) On mentalese orthography. In: *Dennett and his critics*, ed. B. Dahlbom. Blackwell. [aRGM]
- (1994) On unclear and indistinct ideas. In: *Philosophical perspectives*, vol. VIII, ed. James Tomberlin. Ridgeview Publishing. [aRGM]
- (1997) Images of identity. *Mind* 106:423. [rRGM]
- (forthcoming) Historical kinds and the special sciences. *Philosophical Studies* (The Oberlin Colloquium, 1997). [rRGM]
- Minsky, M. (1975) A framework for representing knowledge. In: *The psychology of computer vision*, ed. P. H. Winston. McGraw-Hill. [aRGM]
- Murphy, G. L. & Medin, D. L. (1985) The role of theories in conceptual coherence. *Psychological Review* 92:289–316. [aRGM]
- Neisser, U. (1975) *Cognition and reality*. W.H. Freeman. [aRGM]
- Neisser, U., ed. (1987) *Concepts and conceptual development*. Cambridge University Press. [aRGM]
- Nelson, C. & Horowitz, F. (1987) Visual motion perception in infancy: A review and synthesis. In: *Handbook of infant perception*, ed. P. Salapatek & L. B. Cohen. Harcourt Brace Jovanovich. [aRGM]
- Nelson, K. (1974a) Concept, word, and sentence: Interrelations in acquisition and development. *Psychological Review* 81:267–85. [KN]
- (1974b) Variations in children's concepts by age and category. *Child Development* 45:577–84. [aRGM]
- (1991) The matter of time: Interdependencies between language and thought in development. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- (1995) The dual category problem in lexical acquisition. In: *Beyond names for things*, ed. W. Merriman & M. Tomasello. Erlbaum. [KN]
- Nuboe, J. F. W. (1986) A comparative view on colour vision. *Netherlands Journal of Zoology* 36:344–88. [JvB]
- Perner, J. (1995) The many faces of belief: Reflections on Fodor's and the child's theory of mind. *Cognition* 57:241–69. [JP]
- Perner, J., Baker, S. & Hutton, D. (1994) Prebelief: The conceptual origins of belief and pretence. In: *Children's early understanding of mind: Origins and development*, ed. C. Lewis & P. Mitchell. Erlbaum. [JP]
- Piaget, J. (1926) *The language and thought of the child*. Harcourt Brace. [aRGM]
- Pinker, S. (1994) How could a child use verb syntax to learn verb semantics? In: *The acquisition of the lexicon*, ed. L. Gleitman & B. Landau. MIT Press. [rRGM]
- (1997) *How the mind works*. W. Morrow. [PBI]
- Pipp, S., Easterbrooks, M. A. & Harmon, R. L. (1992) The relation between attachment and knowledge of self and mother in one- to three-year-old infants. *Child Development* 63:738–50. [CS]
- Pribram, K. H. (1991) *Brain and perception: Holonomy and structure in figural processing*. Erlbaum. [BJM]
- Pribram, K. H., Spinelli, D. N. & Kamback, M. C. (1967) Electrooculographic correlates of stimulus response and reinforcement. *Science* 157:94–96. [BJM]
- Putnam, H. (1973) Explanation and reference. In: *Philosophical papers*, vol. II. Cambridge University Press. [VCM]
- (1975a) The meaning of "meaning." In: *Language, mind and knowledge*, ed. Keith Gunderson, vol. 7 of *Minnesota studies in the philosophy of science*. University of Minnesota Press. [aRGM, VCM]
- (1975b) *Mind, language, and reality*, vol. II of *Philosophical papers*. Cambridge University Press. [GC, BF, AG]
- Quine, W. V. (1960) *Word and object*. MIT Press. [arRGM, PBI, KN, JvB]
- (1969) *Ontological relativity and other essays*. Columbia University Press. [JvB]
- (1992) Structure and nature. *The Journal of Philosophy* 89:5–9. [JvB]
- Quinn, P. C. & Eimas, P. D. (1993) Evidence for representations of perceptually similar natural categories by 3- and 4-month-old infants. *Perception* 22:463–75. [FX]
- (1997) Perceptual organization and categorization in young infants. In: *Advances in infancy research*, vol. 11, ed. C. Rovee-Collier & L. P. Lipsitt. Ablex. [JMM]
- Rifkin, A. (1985) Evidence for a basic-level in event taxonomies. *Memory and Cognition* 13:538–56. [aRGM]
- Rosch, E. (1973) Natural categories. *Cognitive Psychology* 4:328–50. [aRGM, JvB]
- (1975a) Cognitive representations of semantic categories. *Journal of Experimental Psychology* 104:192–233. [JS]

- (1975b) Universals and cultural specifics in human categorization. In: *Cross cultural perspectives on learning*, ed. R. Brislin, S. Bochner & W. Honner. Halsted. [aRGM]
- (1978) Principles of categorization. In: *Cognition and categorization*, ed. E. Rosch & B. B. Lloyd. Erlbaum. [aRGM, LKK]
- Rosch, E. & Mervis, C. B. (1975) Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology* 7: 573–605. [aRGM, JS]
- 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. [aRGM]
- Roth, E. M. & Shoben, E. J. (1983) The effect of context on the structure of categories. *Cognitive Psychology* 15:346–78. [JS]
- Rumelhardt, D. E. (1980) Schemata: The building blocks of cognition. In: *Theoretical issues in reading comprehension*, ed. R. J. Spiro, B. C. Bruce & W. F. Brewer. Erlbaum. [aRGM]
- Saunders, B. A. C. & van Brakel, J. (1996) The phantom objectivity of colour: With reference to the works of Franz Boas on the Kwakiutl. In: *Translation of sensitive texts*, ed. K. Simms. Rodopi. [JvB]
- (1997) Are there nontrivial constraints on colour categorisation? *Behavioral and Brain Sciences* 20:167–232. [JvB]
- Schacter, D. L. & Cooper, L. A. (1993) Implicit and explicit memory for novel visual objects: Structure and function. *Journal of Experimental Psychology: Learning, Memory and Cognition* 19:995–1009. [PBo]
- Schank, R. C. & Abelson, R. P. (1977) *Scripts, plans, goals and understanding*. Erlbaum. [aRGM]
- Schröder, J. (1994) On the reference of proper names. *Conceptus* 27:219–37. [JS]
- Schyns, P. G. & Murphy, G. L. (1994) The ontogeny of part representation of object concepts. In: *The psychology of learning and motivation*, vol. 31, ed. D. L. Medin. Academic Press. [VCM]
- Sellars, W. (1963) *Science, perception and reality*. Routledge and Kegan Paul. [JvB]
- Shepard, R. (1976) Perceptual illusion of rotation of three-dimensional objects. *Science* 191:952–54. [aRGM]
- (1983) Path-guided apparent motion. *Science* 220:632–34. [aRGM]
- Singer, W. (1995) Development and plasticity of cortical processing architectures. *Science* 270:758–64. [BJM]
- 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. [aRGM]
- Smith, E. E. (1995) Concepts and categorization. In: *Thinking*, 2nd ed., ed. E. E. Smith & D. N. Osherson, vol. 3 of *An invitation to cognitive science*. MIT Press. [LKK]
- Smith, E. E. & Medin, D. L. (1981) *Categories and concepts*. Harvard University Press. [aRGM, LKK]
- Soja, N. N., Carey, S. & Spelke, E. S. (1991) Ontological categories guide young children's inductions of word meaning: Object terms and substance terms. *Cognition* 38:179–211. [PBI, SW]
- Spelke, E. (1989) The origins of physical knowledge. In: *Thought without language*, ed. L. Weiskrantz. Oxford University Press. [aRGM]
- Spelke, E. S. (1993) Physical knowledge in infancy: Reflections of Piaget's theory. In: *The epigenesis of mind: Essays on biology and cognition*, ed. S. Carey & R. Gelman. Erlbaum. [aRGM]
- (1994) Initial knowledge: Six suggestions. *Cognition* 50:431–45. [PBI]
- Tversky, A. (1977) Features of similarity. *Psychological Review* 84:327–52. [KRL]
- Tversky, B. & Hemenway, K. (1984) Objects, parts and categories. *Journal of Experimental Psychology: General* 113:169–93. [aRGM]
- Uller, C. (1997) Origins of numerical concepts: A comparative study of human infants and nonhuman primates. Ph.D. dissertation, MIT. [MDH]
- Uller, C., Xu, F., Carey, S. & Hauser, M. D. (in press) Is language needed for constructing sortal concepts? A study with nonhuman primates. *Proceedings of the 21st annual Boston University Conference on Language Acquisition*. [MDH]
- van Brakel, J. (1991) Meaning, prototypes and the future of cognitive science. *Minds and Machines* 1:233–257. [JvB]
- (1992) The ethnocentricity of colour. *Behavioral and Brain Sciences* 15:53. [JvB]
- Vasta, R., Haith, M. M. & Miller, S. A. (1995) *Child psychology: The modern science*, 2nd ed. Wiley. [CS]
- Ward, T. B. & Becker, A. H. (1992) Intentional and incidental learning. In: *Percepts, concepts, and categories; The representation and processing of information*, ed. B. Burns. North-Holland. [aRGM]
- Watson, J. S. (1985) Contingency perception in early social development. In: *Social perception in infants*, ed. T. M. Field & N. A. Fox. Ablex. [CS]
- Waxman, S. R. (1991) Semantic and conceptual organization in preschoolers. In: *Perspectives on language and thought*, ed. S. A. Gelman & J. P. Byrnes. Cambridge University Press. [aRGM]
- (in press) The dubbing ceremony revisited: Object naming and categorization in infancy and early childhood. In: *Folkbiology*, ed. D. L. Medin & S. Atran. MIT Press. [SW]
- Waxman, S. R. & Hall, D. G. (1993) Assumptions about word meaning: Individuation and basic level kinds. *Child Development* 64(5):1550–70. [SW]
- Waxman, S. R. & Markow, D. B. (1995) Words as invitations to form categories: Evidence from 12-month-old infants. *Cognitive Psychology* 29:257–302. [SW]
- Waxman, S. R., Senghas, A. & Benveniste, S. (1997) A cross-linguistic study of the noun-category bias. Evidence from the French- and Spanish-speaking preschool children. *Cognitive Psychology* 43:183–218. [SW]
- Wellman, H. K. & Bartsch, K. (1988) Young children's reasoning about beliefs. *Cognition* 30:239–70. [JP]
- Wellman, H. M. (1990) *The child's theory of mind*. MIT Press. [JP]
- Westphal, J. (1987) Colour: *Some philosophical problems from Wittgenstein*. Blackwell. [JvB]
- Wettstein, H. (1988) Cognitive significance without cognitive content. *Mind* 97:1–28. [aRGM]
- Wiggins, D. (1980) *Sameness and substance*. Blackwell. [rRGM, FX]
- Wilson, E. O. (1970) *Insect societies*. Harvard University Press. [MDH]
- Winograd, T. (1975) Frame representations and the declarative-procedural controversy. In: *Representation and understanding: Studies in cognitive science*, ed. D. G. Bobrow & A. Collins. Academic Press. [aRGM]
- Wittgenstein, L. (1922) *Tractatus logico-philosophicus*. Harvard University Press. [rRGM]
- (1953) *Philosophical investigations*. Blackwell. [rRGM, KN]
- Woodfield, A. (1991) Conceptions. *Mind* 4:547–72. [JP]
- (1996) Which theoretical concepts do children use? *Philosophical Papers* 25:1–20. [JP]
- Wynn, K. (1992) Addition and subtraction by human infants. *Nature* 358:749–50. [PBI]
- Xu, F. & Carey, S. (1996) Infant metaphysics: The case of numerical identity. *Cognitive Psychology* 30:111–53. [aRGM, PBI, MDH, FX]
- Zajonc, A. (1993) *Catching the light: The entwined history of light and mind*. Bantam. [BJM]