

JESSE PRINZ, ed.

Oxford University Press (UK), Handbook of Philosophy of Psychology

Thoughts of Real Kinds

An enormous number of experiments have been done by psychologists on "categorization," the way that we divide various things into kinds. Most "categories" that have been studied correspond to simple common nouns such as "bird" and "chair," "fruit," "furniture" and "house," although there have also been studies of how people learn artificial category words defined by the experimenter. The usual assumption behind these experiments has been that these words correspond merely to classes, either exact, or perhaps built on paradigms or exemplars. I will argue that most of these common nouns stand instead for what I will call "real kinds," real kinds being natural units that are not defined, as ordinary classes are, merely by the fact that the members possess common or overlapping properties.

Three entwined but distinguishable assumptions were incorporated in the classical view of categorization as classification, the first having already been questioned by some researchers. By the end I hope to show that these three assumptions all fail for the same reason, and what should be put in their place.

(1) Each kind word in the public language corresponds to a class defined by a set of properties. Since a word, say, *bird*, corresponds to a class, then to judge that an item is a bird is just a shorthand way to attribute a set of properties to it. Kind terms are just complex property terms. They correspond to logical predicates of judgment.

According to the classical view, exhibiting these defining properties was necessary and

sufficient to determine membership in the kind. In the last quarter century, however, a prevalent view has been that many or most common kind terms are defined instead by sufficient resemblance to certain paradigm cases or to certain exemplars. The members may resemble these central instances in somewhat different ways and to different degrees, having properties that merely overlap closely rather than a single set in common. Call classes determined in this manner "blurry classes." But blurry classes are classes none the less. Membership in such a kind would still depend only on having some qualifying set of properties. Kind terms of this sort would still be just shorthand for complex, partly disjunctive, sets of properties.

(2) The properties determining a kind are "represented in the mind" of the person who uses the kind word correctly. Thus to think of a kind requires prior capacities both to think of properties and, presumably, to think of individuals bearing those properties.

To judge that one was encountering a certain kind, such as mankind again or dogkind again involved the capacity to judge concerning an encountered individual that it had certain properties. One had to have the prior capacity to make judgments of the form, say, "Object x is square" or "This object is heavy" in order to employ or understand a kind term. This feature of the classical approach to the study of categories has not been challenged, I think, in the literature. And although talk of "categories" often blurs into talk of "concepts" in the psychological literature, no general analysis of what it is to have a concept or the capacity to think either of an individual or of a property has been attempted.

(3) Association of the relevant properties with the kind word constitutes grasp of

the public meaning of the kind word. For two people to understand a kind word as having the same meaning is for them to represent the kind as having the same properties or blurry set of properties.

If people use different properties in deciding when to classify examples under a given kind term then, although this term may be coextensive in their idiolects (as "featherless biped" may be coextensive with "human") it will not have the same meaning in their idiolects.

In the 1980s a somewhat different view of categories began to emerge in some of the psychological literature, a view inherited from the philosophical tradition after Putnam (1975) and Kripke (1972) and somewhat closer to the one that I will defend. The view concerned certain kinds that were labeled "natural kinds." Natural kinds were taken to be units found in nature rather than assembled by language conventions or by human dispositions to classify. Natural kinds had essences, deep "theoretical" characteristics, from which their more superficial and readily observable properties flowed. But according to this view, these essential properties were not typically represented in the mind of a person thinking of the natural kind but only uncovered through the investigations of empirical science. Various experiments involving both adults and children suggested that people do take the unity of many kinds to rest on "hidden essences" that are unknown to them (e.g., Medin and Ortony 1989) --a finding that fits well with the fact that people quite systematically fail to give definitions exempt from counter-examples when asked for definitions of common kind terms. But this position left wide open the question, what does thinking of a natural kind amount to? This question was sometimes acknowledged with a brief nod toward what philosophers call "externalism" about mental content. Externalism is

the view that what one is thinking of is not determined solely by what is within one's mind, by one's (narrow) psychological state, but in part by some sort of causal-order relation between one's thought and its object. Exactly what constitutes this external relation was not discussed in the psychological literature, however, nor was there any agreement among philosophers.

The resulting psychological theory remained inadequate, I believe, in several ways. Although it denied that the extension of a kind term is determined merely by certain properties represented in the mind of the user, it did not say what about a mind does determine that it is thinking of a certain natural kind. Also, the kinds with which the theory was concerned were still kinds that depended for their unity merely on common properties, though inner or hidden ones, whereas I will argue that most of the kinds that have been studied by psychologists, such as "bird," "chair," "fruit" and "house," are not examples of natural kinds in this narrow sense. Further, remnants of classical theses (2) and (3) above tended to cling to the new approach. Knowing the meaning of a kind word, though no longer assimilated to knowing defining properties of the kind, was usually assumed to require representing in one's mind at least a few definite properties shared by typical kind members, which properties had to be associated with the term in order to understand its meaning. For example, Putnam claimed that it is part of the meaning of "tiger" that stereotypical tigers are animals, large and striped and if someone doesn't know this they don't properly understand the word "tiger."¹

¹ Barsalau and his colleagues (1987), not working within an externalist tradition, have claimed that different people designate different sets of properties as being typical of

Let me move toward a new way of thinking about kind terms by noting two facts about child development that don't fit easily either with classical views or with externalist views that retain assumptions (2) and (3).

First is the finding that children learn names for individuals and for many concrete kinds and stuffs ("milk," "juice") well before they acquire names for even the most obvious and common properties, such as shapes and sizes and colors. The bulk of a child's earliest words are concrete nouns, including names of individuals, names of concrete kinds and stuffs, dozens of these being acquired in a rush between about eighteen months and two years. Adjectives come later and more slowly, and abstract nouns later still (Markman 1991). And there is 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 appear to be less fundamental than those expressed with simple concrete nouns. This is puzzling on the view that a prior capacity to think of various properties characterizing a kind is required in order to recognize or to think of the kind. Notice that this finding is also puzzling on the traditional philosophical view that thinking of an individual requires thinking of properties unique to that individual. I will return to this latter point soon.

The second finding that doesn't fit is that children acquire an average of at least five

a kind, indeed, that what properties a single person will designate as typical changes depending on context. But this does not speak to the question whether there are certain key properties that everyone always has to have in mind when understanding a public-language kind term.

to nine new words each day between the ages of two and six (Waxman 1991, Clark 1991, Byrnes and Gelman 1991). Indeed, Chomsky says they acquire "about a word an hour from ages two to eight with lexical items typically acquired on a single exposure..." (Chomsky 1995, p. 15). These words are mostly acquired in linguistic context and without accompanying explanations of their meanings. This is puzzling on the view that learning the meaning of a public-language kind word involves acquiring certain information about properties of the kind that are either defining or considered typical by all competent users. One wonders how the relevant information on meaning is imparted to the child, apparently effortlessly and on the fly.

If the proposal I will make is correct, these two findings from child development are exactly as would be expected. But on this proposal, assumptions (1) through (3) all fail. Full details of the proposal are given in my On Clear and Confused Ideas (Millikan 2000, hereafter OCCI). Here I will try to explain enough of the theory to make clear what it puts in place of each of the three classical assumptions, how it fills in some gaps for externalist theories on thought content, and how it accounts for the above-mentioned facts on child language. The bulk of the work required to defend the theory cannot be accomplished in this short a compass, indeed, I can sketch only the tip of the iceberg in any detail. But I will try to make this exposed part very clear, and then reference where an examination of various submerged parts can be found.²

² Portions of the remainder this essay have been revised and adapted from "Why (Most) Concepts aren't Categories," in H. Cohen and C. Lefebvre eds., Categorization in Cognitive Science.

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

On the other hand, in the case of species, blood relatedness is bound to be accompanied by considerable overlap in properties. If the species reproduces asexually, the reason is that progeny are clones. If the species reproduces sexually, then each of the genes in the gene pool has to fit in well enough with a random selection of other genes to help produce a viable individual often enough to remain in the gene pool. No single gene that changes the animal in very extreme ways can do this. The result is what is called "homeostasis" in the gene pool. Thus the various individuals within a species mostly resemble one another in a great variety of ways, but do not all resemble one another in any particular ways. What pulls them together as a group is not just that they have common or overlapping properties, but that they tend to have common and overlapping properties for a

good reason. One individual is like the next for a good reason. There is a good explanation of why one is likely to be like the next. Various kinds of inductions drawn over the members of a species are likely to be sound owing to certain kinds of causal connections among these members.

On this analysis of what pulls the members of a species together, species are not in the first instance classes. Classes are defined by the members having certain common properties. Blurry classes may be defined by the members having overlapping properties or by their having many properties in common with a paradigm. But the members of a class do not need to be like one another for any reason. They may be like one another quite by accident. Species names are not mere names for classes. They are names for real kinds, kinds gathered together by real connections.

Now let me explain, first, why this point is important and, second, let me generalize it.

The point is important because it explains why it is possible to study a species as such, that is, efficiently to gather stable information about it. If there is a reason why one dog is likely to be like the next in a good number of respects, then there is a reason why studying one dog is likely to yield probable knowledge about the next dog. In fact, of course, dogs are something it is possible to learn a great deal about. Consider how much may be learned about dogs by a student at veterinary college. But this knowledge is only probable knowledge. Whatever one learns about dogs, it won't be analytic or necessary that every individual dog has all of those properties. But mere classes are not things one can learn anything at all about by induction. If there is no reason, given one member of a

class, why the next member is likely to be like it, then if any inductions over the class turn out true conclusions, it can only be by accident. For example, it seems likely that there is no reason why one red triangular object should tend to be like the next in any respect other than redness and triangularity, so it is not likely that discovering, say, that one red triangular object is sweet will be of any use in predicting the taste of the next.

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

The point about dogs is generalized by noticing other kinds of relations that tend to cement a unit together such that there is a reason why one individual from the unit is liable to be like another. Units of this general sort are what I am calling "real kinds." Following Aristotle, I also call them "secondary substances" so as to point up their similarity to Aristotle's "primary substances," that is, to individuals, for this similarity turns out to be important. Recall that what Ghiselin and Hull said about dogs was that they were big, scattered, historical individuals enduring through time. In OCCI chapter 2, I discuss a variety of differently structured real kinds. Here I will mention only a few important ones, but enough, I hope, to make the general principle clear. Most important, I will try to bring out the similarity of real kinds to individuals. For my claim will be that elementary thoughts of real

kinds are similar in structure to thoughts of individuals, neither having to rest on prior capacities to think of properties. To this end I will often speak just of "substances," using the Aristotelian term that covers both real kinds and individuals, for many of the same points need to be made about both.

Substances fall roughly into three basic sorts: "historical kinds," "eternal kinds" and "individuals." Historical kinds are like dogs. They are collections of individuals scattered within a roughly bounded spatio-temporal area that are causally related to one another in such a manner that each is likely to be like the next in a variety of respects. There are two obvious kinds of causes that can make members of an historical kind tend to be alike. First, something akin to reproduction or copying may have been going on, all the various individuals having been produced from one another or from the same models. Second, the various members may have been produced by, in, or in response to, the very same ongoing historical environment, for example, in response to the presence of members of other ongoing historical kinds. A third and ubiquitous causal factor often supporting the first is that some function is served by members of the kind, understanding "function" roughly in the biological sense of an effect raising the probability that it's cause will be reproduced. It is common for these three causes of likeness to be combined.

Consider chairs, for example. Chairs have been designed to fit the physical dimensions and practical and aesthetic preferences of humans, who are much alike in relevant respects for the same reasons dogs are. Moreover, the design of a chair is pretty invariably influenced by the design of previous chairs, typically because these previous chairs have functioned well and were aesthetically pleasing within a cultural setting, relevant

aspects of the cultural setting being reproduced elements as well. For these reasons, chairs form a rough historical kind. There are reasons that have nothing to do with any arbitrary points of definition why one knows roughly what to expect when someone offers to bring a chair.

Renditions of "The Irish Washer Woman" and of "The Rite of Spring" form an historical kinds. They are copied from one another or from scores that are transcribed from earlier renditions or copied from earlier scores. McDonald's restaurants form an historical kind. There are causes of their being so much alike. Professors, doctors, and businessmen form historical kinds, especially well integrated ones when these groups are studied as limited to particular historical cultural contexts. Members of these groups are likely to act similarly in certain ways and to have attitudes in common as a result of similar training handed down from person to person (reproduction or copying), as a result of custom (more copying), as a result of natural human dispositions (compare dog dispositions) or social pressures to conform to role models (copying again) and/or as a result of legal practices handed down from univocal sources. There is a reason why it may be productive to investigate, say, "the attitude of American doctors toward acupuncture." These attitudes are contagious. They spread within a communicating group with common backgrounds, traditions, interests and so forth.

Eternal kinds are often called just "natural kinds." The members of eternal kinds are like one another for a different kind of reason than historical kinds. They are alike because they possess a common inner nature of some sort, such as an inner molecular structure, from which the more superficial or easily observable properties of the kind's instances flow.

The inner structure results in a certain selection of surface properties, or results in given selections of properties under given conditions. Popular examples of this sort of kind are the various chemical elements and compounds. Water molecules, electrons, protons, and so forth, are also members of eternal kinds.

The members of an historical kind need not all have any properties in common. Being a member of the kind merely raises the probability of having certain properties. Not all dogs bark, and not all doctors have typical attitudes towards acupuncture. Also, many historical kinds do not have sharp boundaries. Typical members may fade slowly into borderline cases. Eternal kinds, on the other hand, can be said to have "essences" in a very traditional sense, essences that are not nominal but real, often discovered only through empirical investigation. The members of eternal kinds have many properties in common because they have a few fundamental properties and/or causes in common that account, given natural laws, for the others. Eternal kinds do form classes, all of whose members are alike in a variety of respects. But eternal kinds are also more than mere classes; they are alike not by accident but in accordance with a causal explanation.

The third kind of Aristotelian substances are individuals. Individuals have been taken in modern times to have a very different sort of unity than the unity of a kind, but there is a way in which the cement that holds a single individual together as it endures through time is quite a lot like the cement that holds an historical kind together. Ghiselin and Hull claimed that species are actually individuals, because they are held together not by a traditional essence but through historical causal connections. The other side of this coin is that individuals are rather like species. A species is a "homeostatic system....amazingly

well-buffered to resist change and maintain stability in the face of disturbing influences" (Eldredge and Gould p.114, quoted by Hull, p. 199). Similarly, an individual animal is a "homeostatic system....amazingly well-buffered to resist change and maintain stability in the face of disturbing influences." If a person is tall, brown haired, knowledgeable about electronics and a good piano player today, it is likely, though not certain, that she will have each of these traits also tomorrow. The various members of a species are like one another in part because they are, speaking roughly, copied from one another. An individual physical object tends to have the same physical properties the next day as it had the last because of natural conservation laws which tend, rather than copying, to preserve its properties from day to day. The effect, however, is much the same. The inference that an individual animal or inanimate object will probably have these and those properties tomorrow because it has them today is likely to yield a true conclusion for the same sort of reason that an inference that other members of a species probably have these and those properties because this member has them is likely to yield a true conclusion. Individual objects are things that inductive knowledge can be collected about over time for the same sort of reason that historical kinds and, more broadly, all real kinds, are things that knowledge can be collected about over time.

I have explained why historical kinds, eternal kinds and individuals are similar and why it is possible to gain inductive knowledge about one part of the cemented-together unity each composes from other parts of the unity. None of these Aristotelian substances are merely classes, either focused or blurry. Because real kinds are not units cemented together merely by having some set of common or overlapping properties, the extension of

the concept of a real kind cannot be determined merely by a conjunctive or disjunctive set of properties represented in the mind. What is it to have a thought of a certain real kind, then, if not to have in mind a set of central properties?

First I will argue that there are, in general, many different ways to have thoughts of the same real kind, any one of which will equally enable a person to understand the meaning of a public-language term designating that kind. That is, I will argue that classical assumption (3) is mistaken. There need be nothing in common to the inner psychological states of two persons who understand the same public-language kind term correctly. Then I will argue that assumption (2) is also mistaken. Being able to think of a real kind need not entail the capacity to think of any of its properties. I am going to argue these two points by discussing, first, what it is to have a thought of an Aristotelian primary substance --of an individual. Then I will generalize to real kinds.

What is involved in thinking of an individual? One traditional answer is that to think of an individual is to capture that individual in thought with a description uniquely identifying it. Another is that to think of an individual requires that one know how to identify it one way or another, perhaps by description or perhaps just by being able to recognize it, to differentiate it from other individuals in perception. These views share an important assumption, namely, that there is more than one way to think of the same individual. Indeed, surely there are innumerable ways to think of the same individual. An indefinite number of individuating descriptions apply to every individual. Further, there are, in general, very numerous different ways that the same individual might be recognized by sight, from front, back or side, by face or body stance from different distances, by voice or by other

characteristic sounds, by smell (animals, but also babies, are good at identifying individuals by smell) and so forth. Contrast the ways Helen Keller recognized her friends with the ways they recognized one another. Clearly there is no single or definite set of properties that one must either think of or be able to discriminate in order to have a concept, a thought, of a given individual. Different people can have different kinds of concepts of the same individual by using quite different methods of recognition. Nor do the methods any one person uses to recognize another constitute a definition for the one person of the other. Your mother is not defined for you by the way you recognize her, say, by the look of her face and the sound of her voice. She doesn't have a definition, either an appearance or a set of properties, that make her be who she is. She is not a class that happens to contain only one member.

Similarly, there is no central set of properties, all or some of which one must be able to think of, recognize or discriminate in order to think of the real kind dog, hence in order to learn about dogs, to understand things said about dogs and so forth. The species dog is a unity of which different people can have quite different kinds of conceptions, using quite different methods of identification or recognition. What makes real kinds worth having thoughts of is that so often there is much that can be found out and known about them. J. S. Mill said about real kinds, "a hundred generations have not exhausted the common properties of animals or plants... nor do we suppose them to be exhaustible, but proceed to new observations and experiments, in the full confidence of discovering new properties which were by no means implied in those we previously knew" (from Hacking 1991, p. 118). Typically, very numerous of these properties, taken either alone or in small sets, are

each diagnostic of the substance. That is, each of these properties or property sets will be found only or typically when the substance itself is encountered. At least this may be so within the spatial and temporal locale traversed by the thinker. Mistakes one might have made had one lived in a different place or time, say, when an Aunt Samantha look-alike or a prehistoric grey squirrel look-alike was about, are not relevant to one's actual capacities to recognize these substances (OCCI chapters 4, 13-14). Thus people can have concepts of the very same real kind using very different means. Children and chemists have different ways of recognizing sugar. Some people recognize red pines by their bark, others by their needles. Probably Helen Keller recognized dogs, cats and chickens by smell, exactly as very young babies do their mothers (MacFarlane 1977). Indeed, you and Helen Keller probably have at least somewhat different ways of recognizing every common real kind. Nor does the method or methods that a person uses for recognizing a real kind constitute a definition of what that real kind is, even for that person. A real kind is not just a class that happens to contain so and so many members. Indeed, for many familiar real kinds, one's ability to recognize the kind may be constantly improving, as one learns to recognize it by a wider diversity of diagnostic signs, and under a wider diversity of conditions. None of these recognition methods is any more definitional than any other.

Think why being able to recognize real kinds in alternative ways is useful. If all of the properties typically possessed by members of a real kind were observed on every encounter there would be no gain in learning to recognize the kind itself. One could just deal with the observed properties that happened to be relevant to one's needs at the moment. But in fact different properties are observed on different occasions. That is why "a

hundred generations have not exhausted the common properties of animals or plants," and why it is useful to bother remembering these properties once discovered. Carrying knowledge of a real kind about with us is useful only because most of the kind's properties are hidden from us most of the time. Different properties show themselves on different occasions. For the same reason you need to be able to recognize the kind in a variety of different ways, as wide a variety as possible. Even the most easily observable properties of substances are, of course, distal to our senses, affecting them quite differently under different conditions and when bearing different relations to us. Nor is there anything to make any of the various ways one learns to identify a real kind any more definitional of it than any other.

But something should be said here about fallibility. For it might be supposed that one can't have an ability to recognize dogs unless the diagnostic appearances or diagnostic properties one relies on for recognizing dogs are perfectly correlated with dogs --unless they perfectly define the kind. But, once again, compare our ways of recognizing individuals. The ways we have of recognizing individuals are always fallible. There always are conditions under which you might misidentify even your closest friends or relatives. More generally, one's capacities to recognize anything's objective properties are intrinsically fallible, depending on external intervening or mediating conditions such as lighting conditions, atmospheric conditions, sound absorption and reflectance properties of surrounding objects, obscuring conditions such as intervening objects, masking sounds and odors, and so forth. Nor is there an independent way of ascertaining what these mediating conditions happen to be in a particular case. Having the ability to recognize an

individual cannot mean one is infallible in recognizing them. I have the ability to walk. It is one of my finest abilities. It does not follow that I will never trip and fall when trying to walk. Needed here is a description of the nature of abilities, such as walking or recognizing one's mother or dogs, that does not collapse an ability into a simple disposition. Such an analysis has been given in OCCI chapter 4. But with or without analysis, that all abilities are fallible is surely common sense. Of course an ability to recognize Mama, or dogs, or chairs is never infallible.

Just as none of the ways that a person knows to diagnose a substance need be infallible ways, no particular set of properties used to diagnose a substance are ever definitional of it. It is true that in the case, especially, of eternal kinds, empirical investigation may reveal (with probability) that, in fact, some sets are reliably diagnostic. But it is always logically possible that there is some other substance that has parts of its cemented-together unity that share the very same properties as the properties one is using, with practical success, for diagnosis of a certain substance. The possibility of twin earth water, certainly of twin earth dogs and, indeed, of twin earth Mama, indistinguishable from your mother, is never ruled out by logic alone. How then is it determined that your current thought is of water and not of twin earth water or of your mother and not of twin earth Mama? The analysis I have given of abilities in OCCI chapters 4 and 13-14 implies that the object of an ability is determined not merely by current dispositions but in part by the origins of the mechanisms responsible for those dispositions. Your abilities to recognize water and Mama were acquired on earth through interaction with earth water and with your mother. That is what makes the difference.

Next consider whether having thoughts of properties is needed for having thoughts of real kinds. Again, begin by comparing thoughts of individuals. Does the tiny infant that recognizes his mother by smell have a thought of a particular smell, judge that this perceptual object before him has this smell and conclude that this perceptual object is Mama? If the infant recognizes Mama on account of, or because of, her smell, does it follow that he must think of this smell? If he recognizes Mama on account of the quality of her voice does this mean that he thinks of her voice, or of her, as having a certain auditory property? Put generally, does the capacity to discriminate a certain quality so as to react differently in its presence, say, by identifying some currently apprehended object differently, require having a thought of that quality and of something as having that quality? Do dogs that can recognize their masters then have thoughts of various properties of their masters? How about mice when they recognize one another?

We need to avoid a verbal dispute here. But suppose that we take a thought of an object or of a property to be not just any mental representation of it but a representation of a sort that has somewhat versatile uses. That is, the representation is not dedicated to just one function, but might be used in different contexts to help serve different functions. I do not mean to impose anything as strict as Gareth Evans' "generality constraint" on thoughts or concepts (Evans 1982). But if the capacity to react differentially to a certain smell or auditory quality is confined to exactly one use, encapsulated within a mechanism that does nothing, say, but recognize Mama, I believe it will only cause confusion to call the neural representation of that smell or quality a "thought." Consider for comparison the perceptual mechanisms that are responsible for color constancy and for depth perception. To

recognize the same color as the same under various lighting conditions requires recognition, on some neurological level, of the average illumination of the scene one is observing. Somewhere in the neurological machinery average illumination gets represented. Similarly, perception of depth is achieved in part by sensitivity to ocular disparity. Somewhere in the neurological machinery ocular disparity is represented. But we should distinguish dedicated representations of this sort occurring within encapsulated neurological mechanisms from thoughts. Children who show color constancy and who have good depth perception surely do not have thoughts or concepts of average illumination or of ocular disparity. Nor, presumably, do infants who discriminate between Mama and others by smell and voice quality have thoughts of smells and of voice qualities.

There is no reason why the situation with real kinds should be any different. To recognize dog again, to recognize and learn how to deal with dogs, surely a cat does not need to think about any properties that are diagnostic of dogs. Nor need a human infant have thoughts of cat properties to recognize cats so as to learn what to expect from them or how to interact with them happily. There is nothing surprising, then, in the fact that children have words for individuals and for many real kinds and stuffs (milk, juice) long before they acquire concepts of separable dimensions, such as color, shape, and size.

More radically, I have claimed in OCCl that a very young infant need not have thoughts of any of the individual members of a real kind in order to have thoughts of the kind itself. Just as an adult need not have thoughts of the various temporal slices of an individual in order to grasp that the individual has properties that endure over time, a very young infant, or a dog, need not have concepts of individual portions of milk to recognize

milk again nor concepts of individual cats in order to recognize cat again. There are many properties of milk that endure over most samples of milk, and many properties of cat that endure over most instances of cat. To have a concept of milk or of cat is, in the first instance, merely to be capable of learning from experience how to deal with milk or cats, and to know when one has encountered these secondary substances again. Thus it is that children acquire words for both individuals and real kinds well in advance of words for even the most common properties.

An important question remains. How it is possible for a young child, quite effortlessly, to acquire five to nine words along with their meanings each day, a large proportion of these being words for kinds? In talking about concepts of individuals and real kinds, so far I have discussed only the capacities to recognize these, as it were, in the flesh. I have claimed that we possess concepts of individuals and of real kinds, which involve abilities to recognize them,³ in order to collect information about them, which information we pick up on some occasions and then apply on others. To pick up information about a substance you must be in a position to interact with the substance, or with other things that interact with the substance, other things that are influenced by the substance or that influence it. Natural information is transmitted in the space-time-causal order, and you have to be in the same causal order with whatever the information is information about in order to receive that information.⁴ To apply one's knowledge of an

³ I do not mean that a concept of a substance simply equals an ability to recognize it. On what all else is involved, see OCC1.

⁴ I am using the notion natural information in a way somewhat like the way Dretske

individual also requires that one interact with it in the causal order, and that one can recognize it or its manifestations as encountered in perception. But is this claim really plausible?

Surely you can have a thought of the last dinosaur species on earth to go completely extinct and of the first baby to be born next year, or of any other Aristotelian substance which, although you have never encountered it, you do know for it an identifying description. Surely you can have these thoughts without having the slightest idea how to identify any of these things in the flesh. You can have a concept of molybdenum --you can think about it and ask questions about it-- without being able to identify it in the laboratory. You can have a concept of Socrates without being able to identify him in the flesh, even if you were to be transported back to ancient Athens. Let me tackle the descriptions first, then come back to molybdenum and Socrates, for they will prove more interesting.

The descriptions are handled this way. That your circumstances are such that you never get a chance to use an ability that you have does not take that ability away from you. You won't lose the ability to swim just because they chain you to a post in the Sahara desert for the rest of your life. If you understand the terms in any description and know how to apply them, that is, you know how to recognize the other objects and properties and

uses it in Knowledge and the flow of Information (Dretske 1981), yet not quite that way. For the purposes here the difference may not matter, but a careful description of the kind of information I have in mind is in (Millikan 2004 chapters 3-5) where it is called "local information."

relations mentioned in the description, and if you are right that the description is identifying, then you know a way to identify the substance that the description describes. You would do so by encountering something that you could recognize directly as fitting that description, or by coming across something else that you recognized as carrying information telling what fits that description. There are many cases in which you just aren't at all likely to come across any such information, but that is irrelevant to whether you have a capacity to recognize the substance. I am assuming here a fairly usual reading of the notion of natural information, according to which information about the past and about the future are entirely routine kinds of information (also see footnote 4 above). And I am about to claim that language is a standard medium through which natural information is transmitted, hence a standard medium through which substances are recognized exactly as they are recognized "in the flesh" through other media such as light and sound.

Now consider molybdenum and Socrates. It is an obvious fact that many of our concepts of substances have been acquired without encountering those substances directly but only by hearing about them. Moreover, as Kripke (1972), Putnam (1975) and Burge (1979) have observed, we often have no unique descriptions of these substances in mind either. How then can we be said to know how to recognize them? The answer, I claim, is that speech is just as direct a medium for the perception of objects and events and their properties as is the light reflected off objects, the smells emanating from objects, the sounds emanating from events in the environment, or the mechanical stimulations caused by objects in contact with one's skin. This is a thesis that requires defense, and I have defended it in OCCI chapter 6 and in (Millikan 2004) chapter Nine. Here I advance only the

rough idea, hoping that if you are dubious you will look to these defenses.

The claim is that hearing and immediately believing a sentence about a fact or occurrence is in relevant respects just like, for example, immediately believing what you see. There is experimental evidence that what one is told goes directly into belief unless cognitive work is done to prevent this, just as what one perceives in other ways, through other media, does. Loading the cognitive systems with other tasks, such as having simultaneously to count backwards by threes, has the effect of facilitating belief fixation regarding whatever one hears or reads (Gilbert 1993). Recognizing a linguistic reference to a substance is as much a way of recognizing the substance "in the flesh" as any other way of recognizing it. It is identifying it and recognizing natural information concerning it through one more medium of manifestation. Think of this medium, the speech of another person, as like an instrument that aids perception. The lens of one's eye is, of course, an instrument that aids perception. If one wears corrective lenses or looks through a telescope these are other such instruments. The speech of another person is the output of a somewhat more complicated instruments of this kind. Like a camera, a radio, a cat scan, or a microscope, another person who talks to you picks up information-bearing patterns from the environment, focuses them, translates them into a new medium and beams them at you. Think of living in a language community as like being inundated in one more sea of ambient energy. Like the surrounding light, surrounding people transmit the structure of the environment to you in ways that, barring certain interferences, you have become tuned to interpret. Becoming tuned to interpret the information-bearing patterns that are common in a certain language community is coming to understand the language of that community.

Similarly, a radiologist must learn to interpret the information contained on X-ray images and the auto mechanic must learn to interpret the information contained in the sounds emanating from ailing automobile engines. (Notice, while we are here, that identifying a substance through its manifestations in a language medium is a transparent example of identification that does not go by way of thoughts of properties of the substance.)

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

A second reason that the comparison between ordinary perception and language comprehension is unintuitive is that ordinary perception is so much more reliable than what one hears said, at least under common circumstances. It is hard to fool ordinary

perception. To create strong perceptual illusions requires a good deal of knowledge about the perceptual mechanisms and often quite special equipment, of the kind, for example, that optometrists have in their examination rooms. This is a difference of degree, however, a mere difference in frequency, not a difference in kind. Recalling that film dubbing is currently the rule rather than the exception, what differences are there, for example, among (1) believing what you apparently see when you look through the peephole into an Ames room (2) believing what you see when a film has been dubbed and (3) believing what you hear someone say when it's false? In the modern world, if you want to believe only what's true, you often have to apply heavy filters to other methods of perception as well as to perception through language.

So we can understand how it is possible to recognize a substance through the information that language bears, indeed, how it is possible to come to be able to recognize a substance pretty much merely by learning a word for it. This is how we manage to have a concept of Socrates and, for most of us, how we manage to have a concept of molybdenum. To have a word for a substance is already to have an essential part of an ability to recognize manifestations of it, and thus incoming information about it, within a particular language community. Nor is there some other more basic method of identifying each substance that is somehow ingredient in knowing the public meaning of its name. There are no linguistic laws about the way you have to identify each of the various substances named in a language in order to be speaking and understanding the language correctly. That is why it is possible for small children to learn, as Chomsky put it, "a word an hour" between two and eight years of age.

I have argued that common words for kinds do not, in general, correspond to mere classes. I have argued that understanding a kind word does not necessarily require that one represent it in one's mind by any properties. And I have argued that grasping the meaning of a kind word does not, in general, require that one associate with it the same properties as others in one's language community do. But many important questions about names for kinds have not been addressed in this essay. I have painted certain features in fair detail, but neglected many more. A more complete picture can be found in OCCI.

References

- Byrnes, J.P. and Gelman, S.A. 1991. "Perspectives on thought and language: Traditional and contemporary views." In Gelman and Byrnes 1991, 3-27.
- Barsalou, L. "The instability of graded structure: implications for the nature of concepts" Ulric Neisser ed., Concepts and Conceptual Development, Ecological and Intellectual Factors in Categorization, Cambridge: Cambridge University Press, 101-141.
- Burge, T. 1979. "Individualism and the mental." In P. French, T. Uehling and H. Wettstein (eds.), Studies in Metaphysics: Midwest Studies in Philosophy IV, 73-121. Minneapolis: University of Minnesota Press.
- Chomsky, N. 1995. "Language and nature." Mind 104, 1-61.
- Clark, E.V. 1991. "Acquisitional principles in lexical development." In Gelman and Byrnes 1991, 31-71.
- Dretske, F. 1981 Knowledge and the Flow of Information, Cambridge MA: The MIT Press.
- Eldredge, N. and Gould, S.J. 1972. "Punctuated equilibria: An alternative to phyletic gradualism." In T.J.M. Schopf (ed.), Models in Paleobiology, San Francisco: Freeman, Cooper and Company, 82-155. Quoted in Hull 1978.
- Evans, G. 1982. The Varieties of Reference, Oxford: Clarendon Press.
- Gelman, S.A. and Byrnes, J.P. eds. 1991. Perspectives on Language and Thought, Cambridge: Cambridge University Press.
- Ghiselin, M. 1974. "A radical solution to the species problem." Systematic Zoology 23, 536-44.

- Ghiselin, M. 1981. "Categories, life and thinking." Behavioral and Brain Sciences 4, 269-283.
- Gilbert, D. 1993. "The assent of man: Mental representation and the control of belief." In D.M.Wegner and J.W.Pennebaker (eds.), Handbook of Mental Control, Englewood Cliffs NJ: Prentice-Hall.
- Hacking, I.A. 1991. "On Boyd." Philosophical Studies 61, 149-54.
- Hull, D.L. 1978. "A matter of individuality." Philosophy of Science 45, 335-360. Reprinted in Sober 1994, 193-216.
- Keil, F.C. 1989. Concept, Kinds and Cognitive Development, Cambridge MA: MIT Press.
- Kripke, S. 1972. Naming and Necessity. Cambridge MA: Harvard University Press.
- MacFarlane, A. 1977. The Psychology of Childbirth. Cambridge MA: Harvard University Press.
- Markman, E. 1991. "The whole-object, taxonomic and mutual exclusivity assumptions as initial constraints on word meaning." In Gelman and Byrnes 1991, 72-106.
- Medin, D.L. and Ortony, A. 1989. Psychological essentialism. In S. Vosniadou and A. Ortony (eds.), Similarity and Analogical Reasoning, 179-195, New York: Cambridge University Press.
- Millikan, R.G. 2000 On Clear and Confused Ideas; An Essay on Substance Concepts, Cambridge UK: Cambridge University Press.
- Millikan, R.G. 2004 Varieties of Meaning Cambridge MA: The MIT Press.
- Putnam, H. 1975. "The meaning of 'meaning'." In ed. Keith Gunderson, Language, Mind

and Knowledge, Minnesota Studies in the Philosophy of Science vol. 7.

Minneapolis: University of Minnesota Press.

Sober, E. (ed.) 1994. Conceptual Issues in Evolutionary Biology, Cambridge MA: MIT Press.

Waxman, S.R. 1991. "Semantic and conceptual organization in preschoolers." In Gelman and Byrnes 1991, 107-145.