

## Handbook of Embedded Cognition.

### Embedded Rationality<sup>1</sup>

Philosophers and laymen alike have traditionally assumed that whether you can reason well, make valid inferences, avoid logical mistakes and so forth is entirely a matter of how well the cogs in your head are fashioned and oiled. Partner to this is the assumption that careful reflection is always the method by which we discover whether an inference or reasoning process is correct. In particular, further experience, observation or experiment never bear on the question whether an inference is valid. Validity is best checked with your eyes tightly closed so you can attend solely to the internal relations among your ideas.

There seems no need to defend these assumptions, nor, to my knowledge, has anyone ever tried. They are pure common sense. Occasionally, however, common sense is a repository for obdurate error. My claim will be that that is so in this case. Rather than being an a priori matter, I will argue, good reasoning needs constant empirical support. Clear thinking is possible only as embedded in a cooperative external world. Because this claim flies rather rudely in the face of common sense, I will first introduce it slowly in several different ways and then illustrate it with a variety of examples.

Among the most common informal fallacies in reasoning are fallacies of ambiguity. These are mistakes that hinge on a word or phrase that has one meaning in some or all of the premises of the argument but another meaning in other premises or in the conclusion. A traditional toy example runs as follows:

The police enforce the laws.

---

<sup>1</sup> Copyright Ruth Millikan 2005

The law of gravity is a law.

So the police enforce the law of gravity.

Real examples of this fallacy are not so blatant, of course. But according to tradition, if any such fallacy should occur in one's reasoning, it will always be detectable by careful enough reflection on the meanings of the terms in the premisses and the conclusion. Terms in a language (words, phrases) are importantly unlike terms in one's thought (mental terms, thoughts of things) in that exactly the same word can denote different things in different contexts, but the same thought always denotes the same thing. Thoughts are never ambiguous, never equivocal. By paying careful attention to the thoughts behind the words, refusing to let the words get in the way, a rational person can always avoid fallacies of ambiguity.

That is the first part of the view to be challenged. I will argue that one can fail to know that the denotation of a thought is unstable. It is possible to have an equivocal thought, to possess a single mental term that denotes more than one thing. You can have two people confused together in your mind, taking them for the same person, storing all the information you have about each of them in one and the same bin in your head. Or you can have two properties confused in your mind, as mass and weight were confused together in people's minds before Newton. Just like words, equivocal thoughts do not display their ambiguities on the surface. It is, in general, an a posteriori matter, a matter of experience, whether one's thoughts are equivocal. So it is possible to fall into fallacies of ambiguity that are not discernable by hard thinking alone.

The opposite of fallacies of ambiguity are inferences one fails to make because unaware that two words or phrases denote the same thing. For example,

Woodchucks are mammals

Mammals are warm blooded

Therefore groundhogs are warm blooded

is a perfectly rational inference if you happen to know that woodchucks are the same thing as groundhogs but, of course, you might not know that. Although tradition says that the same thought is always a thought of the same thing, it also accepts, what is obvious, that different thoughts can be of the same thing, and can be so without your knowing it. You can think of Mark Twain and then think of Samuel Clemens without knowing they are the same man, or think of woodchucks and then think of groundhogs without knowing they are the same species. The same mental term always denotes the same thing, but different mental terms can also denote the same thing, and this will not generally be known a priori. But tradition then says that inferences like the woodchuck-groundhog one above are invalid, even if you do already happen to know that woodchucks are groundhogs. They are invalid, that is, unless you explicitly add into your premises that woodchucks are the same thing as groundhogs. The inference is invalid because one does not know of the identity of woodchucks and groundhogs by mere reflection, so one cannot know just by reflection that the conclusion of such an argument follows from the premises.

The above doctrines both flow directly from the premise that what logically follows can always be known a priori. If you can always tell by reflection alone that an argument is valid and does not equivocate, then the same mental term must always denote the same thing. And if you cannot tell by reflection alone that different mental terms denote the same then arguments that turn on different thoughts of the same (woodchuck and groundhog) are invalid. Both these conclusions follow from the premise that the world outside you, the world

known through sensory experience, is in no way involved in your being rational. Being rational is something you do in your head.

But again, I will argue, this is a mistake. Whether a mental term is equivocal is an a posteriori matter, a matter of experience. So if any inferences at all are to count as valid, inferences that are valid in all other ways and that also do not commit fallacies of ambiguity should surely count as valid, even though the fact that they are not prey to ambiguities is known only a posteriori. Otherwise no mediate inferences at all (that contain empirical terms) will ever count as valid. That is, either we define "valid inference" such that all inferences that are valid must be known to be valid a priori but there are no valid inferences, or we define "valid inference" such that there are some valid inferences but they are not known to be valid a priori. The latter seems a more sensible way to speak. We do not want the term "valid inference" to be empty. But the way experience teaches us whether a mental term is equivocal or not is exactly the same way that it teaches us whether one term is equivalent to another. So if one knows from experience that two terms denote the same thing, inferences in which one freely substitutes the one term for the other should be valid so long as no other is fallacy present.

Here is a third way of explaining the proposal I want to make. Closely related to the rationalist view of rationality described above is the common sense view that if you are genuinely thinking of something, you really cannot fail to know what it is you are thinking of. Bertrand Russell put the matter this way:

...it is scarcely conceivable that we can make a judgment or entertain a supposition without knowing what it is we are judging or supposing about...(The problems of Philosophy, p. 58.)

So if you were thinking of two different things merged under one thought, of course you would know it. But Russell drew a more dramatic conclusion. The passage continues:

...We must attach some meaning to the words we use, if we are to speak significantly and not utter mere noise; and the meaning we attach to our words must be something with which we are acquainted...[but] Julius Caesar is not himself before our minds, since we are not acquainted with him. We have in mind some description of Julius Caesar..., some description of him which is composed wholly of particulars and universals with which we are acquainted. (58-59)

By "not acquainted with Julius Caesar" Russell means not that we have never met Julius Caesar but that Julius Caesar is not a sense datum! For, Russell supposes, only sense data, sensory impressions, auditory or visual sensations, etc., and their properties, can literally be directly before one's mind, and only if what one thinks of is, in this way, literally before or within one's conscious mind could one be sure what it was one was thinking of. The immediate result of this sort of Russellian move was the emergence of various verificationisms and phenomenalisms, an era to which we have no wish, I imagine, to return.

A ubiquitous contemporary and parallel move is clearly illustrated by Michael Dummett:

Meaning is transparent in the sense that, if someone attaches a meaning to each of two words, he must know whether these meanings are the same.

(1978, p. 131)

But, of course, you can fail to know that two words in your vocabulary stand for the

same thing. You can fail to know that Mark Twain is Samuel Clemens; you can fail to know that woodchucks are groundhogs. You can have two thoughts of the same thing without knowing it. Still, according to Dummett you cannot attach the same meaning to two words without knowing it. It follows that the meaning you think of when you understand a word must be separable from the thing you think of when you hear the word. It must be separable from the word's referent or denotation. Although "Mark Twain" and "Samuel Clemens" refer to the same man they must have different meanings or express different thoughts. Perhaps they express different descriptions, as Russell suggested, through which one thinks of the same man. Although "woodchuck" and "groundhog" denote the same species they too must have different meanings or express different thoughts, perhaps expressing different descriptions of this species. Accordingly, very numerous philosophers have distinguished what they call "reference" --what the thought is of-- from something else, which has gone under various different names, such as "sense," "mode of presentation," "intension," or sometimes just "meaning" or "the thought."

According to this view, valid arguments can be known to be valid a priori because instead of concerning what thoughts are about --reference, denotation-- correct reasoning concerns only sense, modes of presentation (or whatever). It concerns only what is before the mind, but referents (such as Julius Caesar himself) are never directly before the mind, of course. Sense determines reference: there are no equivocal thoughts; the same thought always has the same referent. It is just that reference does not determine sense: the same referent is not always thought of in the same way. Thus it is that the toy argument above about laws can be known to be invalid a priori and the one about woodchucks and groundhogs can be known to be invalid a priori as well.

Now that meaning is not the same thing as reference makes a lot of sense when one considers words like "moreover" and "hurrah" and "or" that clearly do have meanings yet do not seem to have any referents at all. Similarly for words like "phlogiston" and "Santa Clause." But to claim that words that do have referents and do have the same referents always have different meanings whenever someone might not know their referents were the same is peculiar. I have a daughter who collects nicknames the way fugitives collect aliases. Natasha gets called "Tasha," she gets called "Nat," sometimes "Ta" sometimes "Banana," often "Mouse" or "Mousie," and so forth. There may easily be folks who don't know, say, that Nat Millikan is Ta Millikan. Does it really follow that "Nat Millikan" and "Ta Millikan" have different meanings? Indeed, do "woodchuck" and "groundhog" really have different meanings?

Two different questions, often confused, should be distinguished here. The first is whether a name such as "Mark Twain" or "Tasha Millikan" or "woodchuck" corresponds to some one particular thought or sense or mode of presentation (whatever) that every competent user of that name must have in mind when using that name, or whether, instead, different competent speakers may use different thoughts, different modes of presentation, when understanding a name of this kind. For example, is there a particular description or descriptions of Mark Twain that everyone who uses the name "Samuel Clemens" with comprehension must have in mind, or might different competent people associate entirely disjoint sets of descriptions with "Samuel Clemens"? Is there some definition of "woodchuck" that is different from the definition of "groundhog" that every competent user of "woodchuck" must have in mind? Or perhaps you are not a competent user of the word "woodchuck" unless you know that woodchucks are groundhogs (and vice versa)? The

position that I would argue for, indeed have argued for at some length,<sup>1</sup> is that with few exceptions, words that denote do not correspond to any definition or mode of presentation that all competent users understand in common. But that will be only a secondary theme in this essay, for I want mainly to talk about thought rather than language.

The second question to be distinguished might count as the central theme of this essay. The question is whether when one thinks of a thing such as Mark Twain or woodchucks one does generally think of it under some mode of presentation or, as is often said, in some particular way. When I think to myself Mark Twain was a writer and when I think to myself Samuel Clemens was a writer am I really thinking two different thoughts? When I think to myself Woodchucks look like they have zippers down the front and when I think to myself Groundhogs look like they have zippers down the front am I really thinking two different thoughts? Obviously, if I don't know that Mark Twain is Samuel Clemens or if I don't know that woodchucks are groundhogs then these are indeed different thoughts. They form, as it were, different patterns in my head. But does it follow that they remain different thoughts after I am thoroughly convinced of the relevant identities? Moving closer to home, I cannot really make out how my thoughts Nat is grey-eyed and Ta is grey-eyed differ in any way. Try it yourself with an old friend and the friend's nickname. Perhaps different associations may go with a person's formal name and their nickname, but a whole different way of thinking about them?

The mistake that has been made, I believe, is to confuse different ways of recognizing, or different keys to recognizing what it is that one is receiving information about, with different ways of thinking about that thing. How I recognize that it is my daughter that someone is talking about is not a way of thinking of her. I recognize who is being talked

about when people say "Ta" and when they say "Nat" and when they say "Mousie" and so forth, but surely I have only one way of thinking about this daughter. I have only one term in my mental vocabulary for her. There could be some people, of course, that I have different ways of recognizing without knowing it so that I keep in my mind more than one mental term for them. But surely for my very good friends, I generally have very many different ways of recognizing them, both in direct perception and through a variety of linguistic manifestations such as names and descriptions, but only one way of thinking of them, only one mental term for them. Indeed, the idea that every way that I have of recognizing a thing yields a different way of thinking of it, a different mode of presentation of it, drifts into incoherency as soon as we seriously try to count ways of recognizing. Ways of recognizing a thing typically are as uncountable as portions of water in a pond. "Ways of recognizing" is not a count noun. Or so I will argue.

To have a concept of an individual, of an empirically evidenced natural property, or of a natural kind typically involves a capacity to recognize that same thing, as such, in a great variety of ways. Speaking more exactly, typically it involves a keen, though fallible, ability to channel both natural information of numerous kinds that may impact on one's sensory surfaces and much information contained in the language one hears so that it comes to a single focus in one's mind, being understood to concern one and the same thing. Knowing what you are thinking of is having this capacity with regard to the thing you are thinking of, and Russell was surely right that one cannot "make a judgment or entertain a supposition" about some thing without having this capacity at least to some degree. But Russell was wrong to suppose that knowing what you are thinking about is an all or nothing affair. Many forms of information may be recognized as concerning the same thing without

all forms of information about that thing being recognized. And although it may not be usual, it certainly is possible to have two concepts, two focal points for information, two mental terms, for the same thing. --"without knowing it" obviously goes without saying, for knowing it could only be constituted by merging these two mental terms into one, putting all the incoming information into one folder.<sup>2</sup> That is the thesis I will support, mostly by displaying a variety of illustrative examples.

Good examples with which to begin are thoughts of empirical properties. Many modern theories describe concepts of individuals or kinds as though these thoughts were reducible to thoughts or judgments about complexes of properties and then ignore the question what it is to think of a property. Thoughts of individuals are analyzed in terms of definite descriptions (as in the above quote from Russell) and thoughts of kinds in terms of properties supposed to define them. Soon I will argue against these classical analyses, but supposing they were right, then showing that thoughts of properties are discovered to be univocal only a posteriori would show that thoughts of individuals and kinds were so as well.

Thoughts of properties obviously cannot all be analyzed in terms of thoughts of complexes of more fundamental properties without regress. The objects (the content) of our most basic concepts of perceptual properties must be determined by our capacities to respond to these properties when they are made manifest to our senses through natural information.<sup>3</sup> To have a concept of square, for example, involves the capacity to respond to natural information concerning the presence of square things as it impinges on one's senses. Similarly for concepts of other perceptible shapes, of sizes, of colors, of textures, of softness and hardness, of heaviness and lightness, of lengths and distances. So let us consider how some of these properties are in fact recognized.

What is involved in being able to recognize, for example, shapes? Think of the variety of proximal visual stimulations to which a given shape may give rise when viewed from various angles, from different distances, under different lighting conditions, through various media such as mist or water, when colored different ways, when partially occluded and so forth. How shape constancy is achieved by the visual system, the capacity to recognize the same shape as the same under a wide range of conditions, is a problem of nearly unimaginable complexity on which psychologists of perception are still hard at work. And shape is also perceived by the haptic systems. You can feel the shape of a small object in your hand in a variety of ways, for example, with these fingers or those, when the object is turned this way or that way in your hand, perhaps by using two hands, either merely by holding the object or by actively feeling or stroking it. You can perceive larger shapes (say, in the dark) by exploring with larger motions that involve your arms, body and perhaps legs, and by employing the touching surfaces of a wide variety of your body parts. This kind of perception of shape, involving the coordination of information about the exact positions of one's body parts with information about what touches these parts, is of such a complex nature that, to my knowledge, psychologists have not even attempted to study it.

Similarly, how color constancy, texture constancy, size constancy, distance constancy and sound constancy are achieved are enormously complicated matters. (We are adept at identifying sounds, especially speech sounds, as the same sound at origin whether near or far, through air or through water, muffled or distorted and so forth.) In each of these cases in which perceptual constancy is achieved it is abundantly clear that no single rule is applied. Different clues are used by the perceptual systems in different circumstances, separately or together. For example, depth is perceived with the help at

least of ocular disparity, tension in the focusing muscles, occlusion of one object by another, knowledge of the size of objects viewed, and atmospheric haze. We also recognize distances by touch and stretch using many different parts of the body, and by ear we recognize fairly well the distances from ourselves of things that make noises. The blind can often tell where nearby walls are located by reflected sound. Measuring distances with a ruler or a tape measure or just a string, or measuring as a surveyor does by triangulation, or measuring with an odometer or a micrometer or by timing the return of light are also ways of determining distances.

That all of these ways of determining a particular distance are ways of recording one and the same property obviously is not something determined by reflection alone. Coordination of these diverse ways of identifying one and the same property has been achieved through long experience, experience of the race during evolution, experience of the growing child resulting in perceptual tuning, experience in measuring and calculating with the use of a wide variety of instruments. Evidence that our concepts of distances are univocal concepts is deeply empirical.

Several more points deserve to be made here. First, none of these ways of telling distances is infallible, certainly none is known to be infallible a priori. Second, none of these ways of telling distances is any more definitional of our concepts of distances than any other. No one of them defines distance, the others being merely correlated. Yet each adds something to our concepts of distances, nor could we have distance concepts at all were we not in command at least of some of these methods of recognition. Third, it should be clear that these various ways of telling distances do not correspond to a collection of prior concepts of properties which are then judged to concern one and the same property.

If there were such a collection of prior concepts, with which we would make judgments about distance-as-perceived-thusly<sub>1</sub> versus distance-as-perceived-thusly<sub>2</sub> and so forth, presumably these concepts would be countable. But ways of perceiving a property are not countable, not just because they are too numerous, but because they are not the right kind of thing to count. Think, for example, of the myriads of different perceptual data structures, in meandering continuous patterns merging into one another, any one of which might lead you to judge that you had perceived, by sight and/or by feel, something of a particular shape, say, the shape of a hammer. These could no more be counted than, say, the number of areas there are on a sheet of paper. That basic perceptual properties are thought of through myriads of different modes of presentation that are then judged to be presentations of the same is a hopeless idea.

The situation is similar with thoughts of individuals. Traditionally it is supposed that to think of an individual is either to capture that individual in one's mind with some description that uniquely identifies it or to be able to recognize it perceptually. But uncountably many different descriptions will fit any individual uniquely and there are uncountably many ways that any individual might be recognized in perception, for example, (if it is a person) by family members or close friends. A family member might be recognized, say, from front, back or side, by the stance of their body, by voice, by characteristic expressions or doings, under each of myriad different lighting or sound mediating conditions, and so forth. There are innumerable alternative methods that might result in thinking of the same individual. Different people can have quite different kinds of concepts of the same individual by using quite different descriptions or methods of recognition, and a single person may be in command of innumerable different ways of

identifying the same individual. Surely the ways I have of recognizing each of my daughters are not countable. Nor do any of these methods constitute a definition of any of my daughters for me. Natasha, say, is not defined for me by the way I recognize her, by the look of her face (from this angle or that), by the sound of her voice (when she is happy or sad) and so forth. She doesn't have a definition, either an appearance or a set of properties, that make her be who she is. None of the ways I can recognize her either in perception or by description is more important than any other in determining who my Natasha thought is a thought of.

Nor is it determined a priori that my Tasha thought is of any one definite person, that these various methods of recognition all converge on the same thing. Indeed, for any single way of identifying Natasha it is not determined a priori even that this single method always captures the same person. Whatever appearances I go by, it is not determined a priori that I will never encounter someone else who has that appearance as well. Similarly, for descriptions: descriptions are never known a priori to fit one and only one individual; they can be empty and they can fail to be unique. Thoughts apparently of individuals can be equivocal (Tweedledum mixed with Tweedledee) and they can be empty (Santa Clause) without one's knowing this a priori.

Thoughts of individuals can also be redundant without one's knowing it. You can have two different thoughts of Samuel Clemens without knowing it, or fail to know that two people named Dr. Jones are in fact the same. When this latter happens there will indeed be separate ways that you go about identifying Dr. Jones that feed into your separate mental terms for him, separate congeries of overlapping methods. But these will not be different ways to think of Dr. Jones, different modes of presentation of him, but only

separate ways that thoughts of him are stimulated. If you discover that this is really one and the same man, that will not result in a new belief taking up residence in your head, a special kind of belief called an "identity belief". The result will be that you merge your two concepts of Dr. Jones into one, now bringing all your methods to a single focus. The result will be that you now know somewhat better than you did before just who it is you are thinking of when you think of Dr. Jones.

Thoughts of biological kinds can be considered in somewhat the same light. This is what J. S. Mill said about them:

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).

We now know what Mill did not know, namely why this is the case. Biological species are not mere classes. The members of a species are not bound into a unit by possessing certain defining properties in common. Members of the same species originate in the same gene pool, but gene pools typically contain alleles for all or most genes. Typically, there are no genes that every member of the species has in common with every other member. On the other hand, as Mill observed, the members of a species do tend to be like one another in an enormous number of respects. This is partly because they originate from genes that have been replicated from one another and like genes (in like genetic context) often produce like phenotypes. It is also owing to a variety of factors tending to produce homeostasis in the gene pool so that novel genes entering the pool are

unlikely to survive unless they produce extremely minor changes. 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 have common and overlapping properties for a good reason. There is a good reason why one member of the species will probably be like the next in very numerous respects. This is why "a hundred generations have not exhausted the common properties of animals or plants."

Because biological species are not classed collected together merely by some set of common or overlapping properties, the extension of the concept of a species cannot be determined merely by a conjunctive or disjunctive set of properties represented in the mind. Moreover, just as no common way of identifying an individual or a property is required of all who think of that individual or property, there is no central set of properties that everyone must use in order to identify a given species. Typically there are very numerous properties which, either taken alone or in small sets, are each diagnostic of the kind. And just as with individuals and with properties, it is true that each person may have very many alternative ways of recognizing a species. Consider how many different ways you have of recognizing the presence of a dog or a cat or a horse. For many familiar species, one's ability to recognize the species may be constantly improving, as one learns to recognize it by a wider and wider diversity of diagnostic signs and under a wider diversity of conditions.

Nor do any of the particular methods that a person uses for recognizing a species constitute some sort of final criterion of encounter with that species. None is more "definitional" than any other, even for that individual. So the same situation obtains here as with concepts of individuals and of empirical properties. There is no a priori guarantee that

it is really the same species, the same glued-together unity, that one's various ways of identifying are reaching. Nor is there an a priori guarantee that any one way that one tries to identify a species always reaches the same kind. It is not known a priori that my ways of identifying dogs, say, don't lump two or more species together under one concept, making that concept equivocal. It is not even given a priori that they don't fail ever to re-identify anything objective at all, that they are not empty. That my ways of identifying reach an actual species is not a priori.

Because there is a reason why the members of a species are like one another, various kinds of inductions drawn over the members of a species will mostly yield true conclusions for a reason. That these conclusions turn out true is not accidental. Thus, again following Mill, species are what can be called "real kinds;" they are not merely nominal kinds. Elsewhere<sup>4</sup> I have argued that there is a variety of different principles that can cement the members of real kinds together such that there is a reason why one member of the kind is likely to be like another. Some real kinds, such as the various biological species, are "historical kinds," their members being alike because something like copying has been going on against the background of some relevant ongoing historical environment. Copying from one another or from the same original plan is why the restaurants within a given restaurant chain tend to be alike, why various renditions of Beethoven's 5th or of The Irish Washerwoman are much alike, why "Greek salads" tend to be similar, why Gothic Cathedrals have similar plans, why American doctors have so many bits of knowledge and also so many attitudes in common, and so forth. A different but more familiar example of a principle that binds real kinds together binds "natural kinds" (in Putnam's sense) together. The members of natural kinds 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.

I have argued<sup>5</sup> that the majority of kinds that are recognized by natural language are real kinds. This is because only real kinds can be genuine subjects of knowledge. It is only when individuals are banded together such that there is a reason why each individual should be like the others in various respects that we can obtain knowledge about this unit as such, unless, of course, by examining each member separately. Thoughts of units of this kind are the seeds on which all empirical knowledge is built, for all empirical knowledge is inductive.

But more important for our purposes here, it is always possible to have a concept of a real kind in any of a variety of different ways, using a variety of different techniques for recognizing its members, either separately or together. This is obvious, for example, in the case of Putnam-style natural kinds. There are, in general, many different techniques for detecting any particular chemical element or compound, many reliable tests for it. Of course the basic structure of the element or compound may be known as a result of scientific investigation (i.e., from experience) but one doesn't just look and see, say, that a substance is composed of atoms with 16 protons! Historical kinds are less often talked about, so let me finish by giving some examples of ways of identifying members of historical kinds that will illustrate why they are never defined by just one method of recognition.

Consider again the various biological species. Aristotle thought that what bound

individual organisms into the same species was that they had a single (Aristotelian) "form" in common. Roughly, he thought of species as being Putnam-style natural kinds. After the Darwinian/Mendelian revolution we think we know better. We know that species are historical kinds: dogs are dogs because they partake not of the same form but of the same gene pool. Aristotle's mistake illustrates a very important point, namely, that you can have perfectly good concepts of a real kinds without necessarily having an understanding of what holds these kinds together. Surely Aristotle was as capable of having thoughts of dogkind or humankind as we are. In order to think of a real kind you don't have to grasp it's basic principle of unity. But suppose that you do grasp this principle. So you decide to use as your basic criterion for whether a creature is a dog that it was born of a dog (better, of two dogs). Now how will you tell whether it was born of a dog? Well, first you have to be able to recognize whether its mother was a dog (and its father)! Knowing the principle of unity that binds the members of a historical kind into a unit takes you nowhere toward recognizing these members. We are back where we started then; there are many equally good ways of recognizing dogs, none of them definitional, and they converge on the same species as a matter only of empirical fact.

Now consider renditions of The Irish Washer Woman. What makes a playing of The Irish Washer Woman into a playing of The Irish Washer Woman is that it is copied from an earlier rendition of The Irish Washer Woman or played from a score copied from earlier renditions so forth. But that doesn't tell you how to recognize The Irish Washer Woman. On the other hand, if you can recognize this tune it is likely that no more than ten consecutive notes anywhere in it will be enough for you to identify it. A tune has no definition. It can be played badly or well, by this instrument or that, with missed notes and sour notes, with

variations. But a tune, a ballad, a symphony, an opera is a real kind, and despite having no definition, and it can, of course, be thought of perfectly well.

Similarly, a Gothic cathedral is one because it has been copied from other Gothic cathedral. Otherwise it's not actually Gothic. But knowing that doesn't tell you anything about the character of Gothic cathedrals. You have to see one or hear it described. Gothic cathedrals are pretty easily recognizable by any number of features, but Gothic Cathedrals don't have a definition. That the various different patterns of features by which they might be reliably identified are diagnostic of a single architectural style is an a posteriori matter, a matter of causal/historical connections.

As a final and more sophisticated example, consider western medical doctors. Children recognize doctors by their stethoscopes and tongue depressors, by the fact that they are taken to see doctors when they are hurt or ill, by the fact that people call them "doctor" and talk about having going to them when they were hurt or ill, and so forth. Adults know which are the doctors by where their names are listed in the yellow pages and by the signs on their office doors, because they say they are doctors and because other people say they are and so forth. But all of that is superficial, you will say; what really makes a doctor into a doctor, in the western world anyway, is that he or she has been trained in an accredited medical school, passed certain examinations and fulfilled various other requirements such as residency requirements, and been licensed by the appropriate legal authority in some country or state to practice medicine. In North America, for example, medical schools are accredited by the The Liaison Committee on Medical Education (LCME) which is "sponsored by the Association of American Medical Colleges and the American Medical Association." But, first, notice that children and probably many adults

don't know that fact, nor would they, for the most part, even know there were such formal facts about doctors, or even know who to ask to find out the details of these facts (I found a medical student to ask, who didn't know exactly but sent me to a website), yet this ignorance does not inhibit them from thinking thoughts of doctors. Second, consider what "accredited" means or what "licensed by a legal authority" means, and so forth. Well, that's complicated, and different things are entailed in different countries and states. But in general, it will involve certain actions on the part of certain institutions that have been granted certain authorities by certain political bodies. In the abstract, that a certain person is a medical doctor will rest on the actions of certain social bodies. Take one of them. Take, for example, the American Medical Association. Which social body is that? How are we to recognize it? Social bodies don't have definitions any more than people do. An individual social body is composed of an historically situated group of people who bear certain complex relations to one another. Identifying the American Medical Association is in many ways like identifying an individual person, or perhaps like identifying activities of the species dog. There are lots of ways to identify this organization or its activities, but none are definitional, and the fact that these various ways all connect with the same organization is an empirical matter, not something known a priori.

Do medical doctors, western style, then form a real kind? There is much knowledge and many skills that they mostly have in common, and also many attitudes and practices, and these similarities obtain for a good reason. Doctors have learned from one another, from teachers who have learned from one another, from the same traditions, indeed, from many of the same textbooks and journals. Their techniques and attitudes have been passed from teacher to student and from colleague to colleague, across national lines as

well as within them. There are good reasons why certain generalizations apply to most or many western doctors and good empirical reasons why doctors can be identified --though fallibly, of course-- in any number of ways. Modern western medical doctors do not merely form a class. They form a real kind about which a good deal can be learned.

The case of western doctors is a rather complicated case. But think back now to the earlier discussions of empirical properties, of individuals and of biological species. These paradigms guide us easily to the following general conclusions.

Abilities to identify and reidentify appearances of the same objective thing as appearances of the same constitute a substantial part of the possession of any empirical concept. Whether these concepts are of empirical properties, of individuals or of real kinds, the abilities to reidentify that underlie them rest on the natural laws that structure natural information. An ability to recognize something is, obviously, not contained in your head alone, anymore than the ability to ride a bicycle is contained in your head. It depends on causal interactions between you and what you perceive, on the way channels of natural information are structured, and so forth. Perfecting the ability to collect this information accurately and efficiently --an ability originally derived through evolutionary history, then through perceptual learning, and finally through experience in making judgments based on perception, linguistic input and inference-- is at every stage an empirical matter. The tests by which we tune our abilities to recognize what is objectively the same as the same are empirical tests all the way down.<sup>6</sup> Learning what is the same as what is at the base of all conceptual development, and conceptual development is a rich and structured interaction between the organism and its environment. To be thinking at all is already to be employing abilities that are deeply embedded in the world.

To discover that Mark Twain is Samuel Clemens or that woodchucks are groundhogs is just one further small step in conceptual development, one more small step in the fallible process of learning what information is carried through what media. Having made this step helps to perfect one's conceptual repertoire; it does not add a necessary step or a new premise in one's valid reasoning from knowledge of properties of Twain to properties of Clemens or from properties of woodchucks to properties of groundhogs. Learning to identify things in a new ways is not storing away special beliefs called "identity beliefs," but improving one's concepts, improving one's basic abilities to think at all.

Like all other abilities that rest partly on the structure of the world outside the organism, of course these abilities are not infallible. It is always possible that an empirical concept binds together information about things that are not the same, hence becomes equivocal, the test for this being, in general, further experience. Ultimately, then, that an empirical concept is not prey to ambiguities is known to one only a posteriori that of one's mediate inferences are valid is known in the same way. One's rationality depends at every point on the complex causal/informational structure of the empirical world. Rationality is firmly embedded in the world outside the mind.

## REFERENCES

- Millikan R. G. 1984 Language, Thought and Other Biological Categories (Cambridge MA: The MIT Press)
- Millikan R. G. 2000 On Clear and Confused Ideas (Cambridge UK: Cambridge University Press)
- Millikan R. G. 2004 Varieties of Meaning (Cambridge MA: The MIT Press)

Russell B.1912 The Problems of Philosophy (New York: H. Holt, Inc.)

Dummett M. 1978 Truth and Other Enigmas (London: Duckworth)

Strawson, P.F. 1974. Subject and Predicate in Logic and Grammar. London: Methuen and Co.

---

<sup>1</sup> Millikan 2000

<sup>2</sup> See Strawson 1972; Millikan 2000 chapter 8 and following.

<sup>3</sup> On natural information, see Millikan 2004, chapters 3 and 4.

<sup>4</sup> Millikan 2000

<sup>5</sup> Millikan 2000

<sup>6</sup> These tests are described in some detail in Millikan 1984 chapters 18 and 19, 2000 chapter 7, 2004 chapter 19.