

**On Knowing the Meaning;
With a Coda on Swampman**
Ruth Garrett Millikan

Abstract

I give an analysis of how empirical terms do their work in communication and the gathering of knowledge that is fully externalist and that covers the full range of empirical terms. It rests on claims about ontology. A result is that armchair analysis fails as a tool for examining meanings of 'basic' empirical terms because their meanings are not determined by common methods or criteria of application passed from old to new users, by conventionally determined 'intensions.' Nor do methods of application used by individual speakers constitute definitive reference-determining intensions for their idiolect terms or associated concepts. Conventional intensions of non-basic empirical terms ultimately rest with basic empirical concepts, so no empirical meaning is ever left merely 'in the head.' I discuss the nature of lexical definition, why empirical meanings cannot ultimately be modeled as functions from possible worlds to extensions, and traps into which armchair analysis of meaning can lead us. The coda explains how 'Swampman' examples used against teleosemantic theories of content illustrate such traps.

**On Knowing the Meaning;
With a Coda on Swampman**
Ruth Garrett Millikan

Dictionaries aim to provide definitions that contain sufficient information to impart an understanding of the term. It is a fact about us language users that we somehow come to understand and use a potential infinity of sentences containing a term once we are given a certain small amount of information about the term. Exactly how this happens is a large mystery. But it does happen, and dictionaries exploit the fact. Note that dictionary entries are not unique. Different dictionaries can give different bits of information and yet be equally effective in explaining the meanings of terms. (Stanford Encyclopedia of Philosophy, 'Definitions,' Anil Gupta, Apr 10, 2008)

1. What I will argue

I will suggest a general analysis of what determines reference or extension for an empirical term, one that is fully externalist and quite different from Putnam's analysis in 'The Meaning of "Meaning"' (1975). It applies to the full range of empirical terms and implies that the meanings of very many empirical terms involve no reference-restricting mental organizations at all that are common to 'the heads' of competent users. For example, users of the term 'water,' need not think of water as 'the same liquid as ... ' something else. To uncover the meaning, in the sense of what restricts the reference, of an empirical term, is to uncover whatever it is that synchronizes usage, supporting agreement in judgments with others and also oneself, hence enabling communication and the successful acquisition of inductive knowledge, thus explaining proliferation of the term's tokens. I suggest that this synchronization is never accomplished entirely through or because of the use of handed-down conventional criteria for application or methods of recognition. How often it is accomplished without the conventional involvement of any prior limiting concepts or methods at all — how many empirical terms lack conventional intensions¹ hence are 'basic' — is a question to be examined.

These claims about empirical meaning rest on an analysis of how empirical terms manage to do the work they do in communication and the gathering of knowledge, which analysis rests in turn and on a suitable ontology. Basic empirical terms are learned in a way that does not pass on to the learner any particular manner of thinking of or identifying their extensions, but only some kind of grip or other, perhaps completely idiosyncratic, on a naturally bonded extension having multiple handles. Individual users may each use a variety of alternative handles, none of which are criterial or definitional either within the public language or within the user's idiolect. Ties bonding the term's extension together inhere in the causal order rather than being projected from people's heads. They pull together a sort of clot of extension, anchoring the term, but often without delineating for it any sharp boundaries. (Where boundary problems might arise in communication, we gracefully turn to different, or more, words.) Natural bonds tie the

¹ I intend to be using the term 'intension' more or less as it was used in Putnam 1973 and 1975, but widely enough to cover paradigm and exemplar theories and also conceptual role theories or 'theory theories' of the concepts traditionally supposed to govern correct application of empirical terms.

extension of such basic extensional terms together in a way that manages to synchronize usage, support agreement in judgments, hence communication, hence proliferation of the term's tokens, without uniformity in methods of recognition or criterial intensions in the heads of users. More fundamental, they support the agreement of term users with themselves, thus allowing the accumulation of stable knowledge and the application of it. The nature of these bonds may be unknown or even misunderstood by the users of the term, though subject to empirical investigation yielding 'real' definitions.

A corollary is that purely a priori or 'armchair' analysis, is not the right tool for examining meanings of basic empirical terms. This is because their meanings are not determined by methods of application that are necessarily common to all competent users. And because non-basic empirical terms always rest ultimately on empirical concepts that are basic, these being held together and in place by empirical-world bonds, the meanings of empirical terms can never ultimately be modeled merely as functions from possible worlds to extensions. All empirical meaning is, in the end, immutably embedded in the actual world.

I will also discuss various implications of these positions, including the nature of lexical definition, and traps into which the traditional example/counter example method of conceptual analysis can sometimes lead us. A coda explains why 'Swampman' examples are such traps, hence do not tell against teleosemantic theories of content.

2. Introducing the problem (with a weasel)

In Plato's *Meno* the question is raised how Socrates and Meno will recognize a correct definition of *virtue* when they find it if they don't already know what virtue is, in which case, of course, they needn't be looking. Sometimes called *the paradox of analysis*, the difficulty can be put this way. Either a word and its analysis have the same meaning, in which case the analysis expresses a trivial identity that could not tell us anything new, or else the word and its analysis do not have the same meaning, in which case the analysis is wrong.

Since Ayer 1936, a traditional response has been that knowing what a word means is not a matter of having something before one's mind when one uses it, but merely knowing how to follow the conventional linguistic rules of application that determine its meaning. Fumerton, for example, compares these rules to the 'extremely complex syntactical rules of sentence structure,' noting that 'there is a world of difference between following rules and being able to formulate the rules one is following' (1989, p. 113). A different response is to reject the very idea that words, for the most part, admit of conceptual analyses or have conventional definitions:

A major problem with the definition story was the lack of convincing examples; nobody has a bullet-proof definition of, as it might be, 'cow' or 'table' or 'irrigation' or 'pronoun' on offer; not linguists, not philosophers, least of all English-speakers as such. (Fodor 1998, p. 92—93)

Moreover, a word's definition is supposed to be analytically equivalent to it, but no reconstruction of the analytic/synthetic distinction has been well received since Quine's devastating attacks.

There is a less discussed problem as well. Between the age of one and high school graduation, the average American child learns 45,000 words (not counting proper names or words easily derived from other words, such as 'reapply' derived from 'apply')

(Bloom 2000, p. 6). That comes to at least ten words a day, many of these learned from only one encounter (Chomsky 1995). How on earth could they be learning conventional definitions for all of these words, or conventional criteria for their application, or conventional inferential connections among them? Mothers do not attempt to impart to their toddlers how they recognize that it is a kitty or a bit of applesauce on which they are commenting, nor can the child observe how they have processed and used the sensory information that lead to their linguistic responses. Indeed, blind toddlers with sighted parents learn words at a normal rate (Bloom 2000, Ch.1). Moreover, very soon the young child begins to learn the bulk of its new words from linguistic context alone. Your own vocabulary was probably obtained largely from reading in context. These simple facts about word learning seem to undermine not only the classical necessary-and-sufficient-definition theories of word meaning, but also exemplar theories and paradigm theories and various conceptual role or 'theory theories' of word meaning.

Still, it is well to be cautious. For example, the measure that is generally used to count a person's vocabulary size is pretty loose. Whether a person 'knows a word' is decided by whether they can pick the definition that fits it from among, say, five taken from the dictionary. Suppose, for example, that you are given the word 'weasel' and asked to choose from among some dictionary's definitions for 'weasel,' 'quark,' 'whistle,' 'conviction' and 'besides.' If you knowingly pick the right definition, does that show that you know the meaning of 'weasel'? Certainly you might pick the right one yet not be able to recognize a weasel, not be able, for instance, to tell a weasel from a mink. Indeed, I was in that position just a summer ago when our cat brought one in. (Foolish but brave cat!) Do you have to be able to tell a weasel to know what 'weasel' means? What exactly IS required for a person to 'know the meaning' of a word like 'weasel'?

That is the central question for this essay. I will begin by elaborating a bit on difficulties encountered when trying to answer. Then I will propose a fresh way to understand what empirically-grounded extensional meaning is, and what it is to 'know' one.

The vocabulary tests could, of course, be made more fine grained. For example, we might ask subjects to choose for the word 'weasel' from among the definitions of 'weasel,' 'stoat,' 'ermine,' 'polecat' and 'mink.' For 'weasel' the *Oxford English Dictionary* says

A carnivorous animal (*putorius nivalis*), the smallest European species of the genus (of the order *mustelidae*) which includes the polecat, stoat, etc. It is remarkable for its slender brownish body and for its ferocity and bloodthirstiness.²

Even without the explicit exclusion of polecat and stoat, I'm sure I would have had no trouble picking this definition. Indeed, excluding the Latin, I could probably have come up with more or less the same myself. Yet I didn't recognize that weasel, right there in front of me! Does it follow that I didn't now what 'weasel' means, or not? Why?

This brings to mind Putnam's claim (1973) that there is a 'division of linguistic labor' whereby certain experts are in charge of really knowing the meanings of certain words while the rest of us take out loans from the experts. Perhaps the problem of how precise

² There is a problem here about American versus British usage. I will be cheating a little in this essay, since only in British usage does the word 'weasel' capture a single species, whereas my weasel was American.

rules or good paradigms for so many words could be learned is moot, because ordinary people use these words without knowing these ways. They learn just enough about the real meanings, say, to tell weasels from such things as quarks, whistles and convictions. Only the experts stand ready with ways by which to tell a weasel from scratch. But there is a puzzle about these experts as well. Surely we must intend that it is their usage we are following and not that of some other persons as ignorant as ourselves. We must intend to follow the usage of the weasel experts, not, say, the mink or the stoat experts. But this seems to return us to Meno's paradox. We know enough about weasels to know, perhaps, that the experts on weasels are biologists, perhaps, more exactly, field biologists, indeed, quite exactly, the ones who know a lot about weasels. But how are we to tell which ones know a lot about weasels rather than, say, mink or stoats, if we don't ourselves know what weasels are? In practice, I suppose, we ask, and they tell us it's 'weasels' they know about. But this seems to make them merely self-proclaimed keepers of the correct rules for the English term 'weasel.' Why should we trust them? More puzzling: *what are we trusting them about?*

Setting that aside for the moment, suppose we examine the methods these experts come up with. We ask them, for example, how they tell weasels from mink. Perhaps they do this by looking at the toes. (Mink have partially webbed toes). Or perhaps they distinguish by looking at the teeth. (Mink don't cut throats). Or perhaps by looking at the fur. (There's a reason why mink are made into fur coats). Or perhaps by the set of the eyes. (Weasels are sneaky and mean). Indeed, shouldn't we expect different experts to have different ways of recognizing weasels, say, the paleontologist telling by the bones, while the field ethologist tells by the stance and the behavior, the trapper by the fur and the toes? And here comes a young expert claiming that a good way to be sure of a weasel is by certain characteristics of its musk glands, another who describes peculiarities of the claws? Surely, the more of an expert on weasels a person is the more different ways they are likely to know how to tell a weasel. On the other hand, surely not every expert knows all these ways. How many ways and which ones does he have to know really to know the meaning of 'weasel'?

These modern experts may agree, of course, that what's really essential to weasels is not this or that trait but ancestry.³ Weasels are, as such, all on the same twig of the evolutionary tree and on a different twig from, for example, mink. The various traits weasels usually have are not definitional but merely symptoms of ancestry. For this reason it is always possible, even, that an expert should be wrong about whether some particular individual is a weasel. It might be an irregular specimen. Or the experts may not realize that there exists another unrelated species easily confusable with the weasel. The main thing about a weasel is only that it has to be born of another weasel (well, two), which in turn was born of another weasel And how do we complete this recursive rule for telling weasels? What else is required for the expert to know the meaning of 'weasel'?

It is not only names of organisms that cause this kind of worry. What is it that 'the experts' on the meaning of 'vitamin B₁₂' know? Perhaps they know the chemical

³ There are contemporary debates about how biologists should understand the notion of a species. Mostly these debates have little bearing in the case of common sexually-reproducing animal species such as weasels and Ernst Mayr's 'evolutionary species concept,' to which I refer here, is typically used for these species in practice.

formula. But any child could uncomprehendingly memorize the formula. Do the experts also have to know how to recognize vitamin B₁₂ in the laboratory? As with chemicals generally, however, there are any number of different ways that vitamin B₁₂ might be identified in the laboratory. Which are the ones used by the special keepers of the semantic rules for 'vitamin B₁₂'? And do the experts have to be competent actually to perform these particular tests in order adequately to understand the meaning of 'vitamin B₁₂'? Do they have to be good chemical cooks? Do they have to be able perform these tests from scratch, for example, without first reading the labels on any bottles in the laboratory?⁴ More mundanely, do we defer to experts on the meaning of 'sugar'? If so, all the same problems arise.

Putnam (1973, 1975) offered what is sometimes known as a 'causal' theory of the meaning of terms like 'tiger' (cf., 'weasel') and 'water' (cf., 'vitamin B₁₂, 'sugar'). This theory concerned the meanings of what he called 'natural kind terms.' Examples anchoring a natural kind term were said to be determined 'indexically,' rather than by conventional defining properties, as, roughly, what had caused our thoughts of these kinds. Other members of the kind were determined as whatever was, in fact, of the *same* natural kind as these examples, a matter for empirical inquiry. The theory assumed, however, that competent speakers of a language know which of their terms are natural kind terms, and that they understand what it would be for one specimen or sample to be of 'the same natural kind' (for example, 'the same liquid') as another. But not only have there been disputes among both experts and philosophers over which are the natural kind terms, there is no agreement on what constitutes sameness of kind. Aristotle thought that what made tigers all be of one kind was a common form, and Putnam thought it was common genes. According to contemporary biologists, neither was right. (For most species, no distinctive genes are found in the genome of every member of the species.) So Putnam's theory does not seem to settle the question what it is to 'know the meaning' of a term like 'weasel.'

More important, this kind of problem clearly is not special to natural kind terms. Absence of an agreed-upon decisive method for determining extension seems to characterize masses of other empirical concepts. Consider, for example, 'square.' How do you determine whether something is in the extension of 'square'? Most of us know that squares have four equal straight sides and equal diagonals. But won't it be enough for the child to understand 'square' that she is able to recognize squares by sight or touch, then learn the geometry later? If that's enough, is it necessary even to be able to recognize squares by sight? Do blind children have to rely on experts for their use of 'square'? Suppose someone knows what square things look like but not how they feel. Or they know how squares feel when small in the hand but not when large in their arms. Or they can tell with their feet but they have no hands? To know what 'square' means, must you be able to tell whether a large piece of land, say a county, is square? That will definitely require you to think of lengths or distances, of course, as would thinking that squares have four straight sides and diagonals of equal length.

⁴ Does 'vitamin' have a conventional intension which helps to delimit the extension of the term 'vitamin B₁₂'? If we found it was of no nutritional value, for example, would we say there was in fact no such thing as vitamin B₁₂? I'm not sure that we know. But what currently guides the usage of a term does not always accord with its origins or its compositional structure.

Consider then the prior question what it is to understand words like ‘distance’ and ‘length’ — to have ‘our concept’ of length or distance? Length is a property with a definite metric.⁵ Does one have to know this or is it enough to recognize just that lengths can be longer and shorter? Think of all the ways there are of recognizing length or distance. Psychologists tell us that depth perception, the recognition of distance from oneself, is achieved by using, alternatively or redundantly, at least ocular disparity, tension in the focusing muscles of the eye, partial occlusion of one object by another, the known size of objects, and atmospheric haze. Does the fellow with only one eye who can’t tell by ocular disparity fail truly to understand what the word ‘distance’ means? Rough distances of things that make noise are recognized by ear. Do the deaf then fail to know what ‘distance’ means? Distances of proximate things can be recognized by reaching, perhaps with the arms or the legs or with a stick. A length may be recognized by first recognizing distance from one’s self, then (unconsciously) reckoning by retinal size, or by comparing with other things of known length at the same distance. Lengths are also recognized by measuring with rulers, tape measures, dividers, micrometers, by triangulation (surveying), or by timing the return of light. Which if any of the hardly numerable ways of telling distances figure in the semantic rules for the word ‘distance’ or the word ‘length’? Which are embedded in the definition of ‘square’? And how do parents pass on to children (at ten words a day) which ways are genuine parts of ‘our concepts’ of length, of distance and of square?

Reflection shows, I believe, that this sort of problem generalizes without limit. Uncompounded, empirically-grounded, extensional terms typically are like ‘weasel,’ ‘sugar,’ ‘square,’ ‘length’ and ‘distance.’ That there is a certain way of recognizing members of the extensions of these terms that is definitional, that provides their meanings in the language, that would correspond, say, to ‘the possession conditions’ for ‘our concepts’ associated with them, is surely a philosopher’s fiction. Indeed, these terms seem to lack any handed-down, conventional, extension-determining criteria or application methods projected from the heads of all competent users. Put briefly, they do not seem to have ‘intensions.’⁶

3. Constructing a theory

Traditional philosophical theory has it that the meaning of an extensional term is given by an intension, a means its competent users conventionally employ for delimiting its extension, that is passed on when the term is learned by new users. Knowing its meaning is knowing (how) to conform with these conventional means of delimitation. The program of ‘conceptual analysis,’ for example, the method of example and counter-example, is a way of attempting to make intensions of this kind explicit. Many contemporary non-classical theories of extensional meaning, such as paradigm theories and exemplar theories and conceptual role theories, even contemporary causal theories, assume the correctness of this very general theory of extensional meaning. They assume that what delimits the extensions of empirical terms is ultimately projected from the heads, from the psychology, from the intentions, of users.⁷ Without pausing to examine the full stock of these ultimately internalist theories, I propose simply to

⁵ Understanding of specific measures of length, such as two feet, does, of course, depend on this.

⁶ See footnote 1 above.

⁷ This assumption called ‘the seed assumption’ and is discussed at length in Millikan 2001.

construct what I take to be a fully externalist theory about the meaning, about what determines the extensions, of our basic empirically-grounded extensional terms, and about what it is to 'know their meanings.'

Unlike Putnam in 'The Meaning of "Meaning",' I can not use the armchair method of example and counter-example, calling on our a priori intuitions about what is or would be in the extension of this term or that, for that method assumes exactly what I aim to disprove. Instead I will attempt to construct a theory using another traditional philosophical method — by sifting and organizing certain matters of common knowledge, mixed with other observations easily made and with some information from contemporary science, then forming an explanatory hypothesis in an attempt to explain the data. In this case the hypothesis will be about what work certain linguistic forms, certain kinds of words and phrases, manage to do and *how* and *why* they are able to do it. As philosophers usually do, I will construct this theory from the armchair, but in the way that the theoretical sciences also do, by reflecting on certain empirical data in an effort to understand them. How language works, I take it, is a broadly empirical question, to be investigated like any other.

To develop this theory, I will need first to invoke a more general theory of linguistic meaning from earlier writings of mine, but placing the emphasis this time on the meanings of extensional terms.⁸ And somewhat later, I will need to review some ontology, also defended more adequately in earlier writings. I will try hard to make these supporting positions sufficiently clear, but without repeating a lengthy exposition or defense.

4. The meaning of 'meaning'

Fifty years ago, Wilfrid Sellars offered a very sensible theory about what 'means' means in the context of, as he called it, 'the translation rubric' — the form "'X" means Y' (1956, 1963). A datum that especially interested him was that this rubric can be used with any kind of word at all. Thus 'Hund' in German means *dog*, 'et' in Latin means *and*, 'hellas!' in French means *alas*, 'niet' in Russian means *no* and 'Chicago est grand' in French means *Chicago is large*. Sellars's hypothesis was that this form gives a functional classification of 'X,' instructing the hearer to use the form 'X' in the way that he already knows to use the exhibited (not mentioned, not used) form 'Y.' The translation rubric comes out true whenever 'X' and 'Y' have 'the same use' in the language or in their respective languages, what counts as sameness of use depending on context: 'differences which are irrelevant in one context of inquiry may be relevant in another' (Sellars 1963, p. 110n.). Sellars's theory is interesting in that it supports itself. Meaning is fundamentally use; the fundamental use of the word 'means' is to induce the hearer to use the word on its left side as he already knows to use the word on its right; we thus have a consistent description of what 'means' means.

Unless we can articulate a general theory of what a linguistic 'use' is, however, Sellars's theory remains superficial. Sellars took 'uses' to involve a community conforming to rules for playing Wittgenstein-style language games including, especially, rules of response to perceptual data with outer and inner language, rules of inference,

⁸ In my 1984 and, especially, 2000, I addressed the nature of empirical property concepts and empirical substance concepts, but a claim of the current essay is that explaining the meaning of a term is not the same thing as describing a concept to which it corresponds.

and behavioral responses to inner and outer language. And he gave what he called a 'normative' account of these rules. They were sustained by social sanctions, adults training the children by administering frowns and smiles until they were consistently following the rules. I am questioning that the meanings of basic extensional terms are determined either by conventional rules of response to perceptual data or conventional inferences from prior concepts. I will not be defending a Sellars-style functionalist theory of language and thought. Also, modern child language studies show that children generally learn language not only at a surpassing rate but with no hint of instruction. Language learning definitely is not driven by social sanction. Happily, however, accepting the heart of Sellars's position on the means rubric need not commit us to Sellars's own way of interpreting 'use.' I have proposed a different interpretation, and a different theory of the norms involved in language use.

I have invoked an analogy between the use or function of a language form and the biological function possessed by an adaptation, for example, the function of an organ or instinctive behavior (1984, 2004, Part III, 2005). Looked at this way, the function — I call it a 'stabilizing function' — of a conventional language form is roughly its survival value. It is an effect it has had that encouraged speakers to keep reproducing it and hearers to keep responding to it in a roughly uniform way, each relying on the settled dispositions of the others. For example, the sentence forms that are labeled 'indicative' in a language have as one stabilizing function to produce in hearers true beliefs, their content varying with the words arranged into this form. Indicative forms survive, in large part,⁹ because the imparting of true beliefs is often of interest to speakers and hearers alike. Thus speakers reproduce indicative forms when their purpose is to impart beliefs, and trusting hearers respond to the forms by believing what has been said with the purpose of acquiring knowledge.¹⁰ Conversely, if the use of the indicative by speakers for lying or misinforming became sufficiently rampant and random, hearers would soon stop believing indicative mood sentences — it would become impossible for a person even to learn this part of the language — and this would cause speakers to stop using the indicative. Like a biological function, *a linguistic function need not be performed regularly*. But it does need to be performed often enough that its linguistic vehicle continues to be reproduced. Without that, the vehicle, the linguistic form, will either be lost or change its function. (Full details and defense are in the references mentioned above.)

If Sellars's analysis of the means rubric is transposed in this way, meaning is, simply, stabilizing function. For example, the (or at least a) stabilizing function of "'x" means y' is to produce in the hearer a disposition to use 'x' in the way he already uses 'y.' To say this is (not to give, but) to *describe* the (or perhaps a) meaning of "'x" means y.'

Besides this most basic dimension of meaning, I have proposed that a secondary kind of meaning for many terms and phrases concerns items in the world to which they necessarily correspond when managing to serve their stabilizing functions in accordance with historically normal explanations for this achievement.¹¹ That is, many

⁹ Indicatives usually have other functions as well, for example, in English the indicative is used for issuing strict orders.

¹⁰ These purposes are not usually intentions, if intentions are represented. A variety of different kinds of unrepresented purposes is discussed in Millikan 2004, Ch.1.

¹¹ See especially Millikan 1984 and 2005, Ch.3.

terms and phrases have referents or extensions in conjunction with their stabilizing functions. To describe these referents or extensions not accidentally, but according to what actually accounts for how they figure in explanations for the survival of their corresponding empirical terms, is part of describing their meanings. Using the example above, to succeed in serving its stabilizing function in the normal way, in accordance with the historically typical explanation for success hence survival, the word ‘means’ in the clause “‘x’ means y’ needs to correspond to the relation that holds between two terms when their stabilizing functions are the same.¹²

I will now try further to explain and support this position on meaning by showing more carefully how it applies to basic, empirically-grounded extensional terms and casts light on the problem what it is to know their meanings.

5. Applying the theory to empirical terms

How then should we describe the meanings of empirically-grounded extensional terms such as ‘weasel’ and ‘square’? A start is to say that a stabilizing function of such a word, ‘X,’ is helping to implant satisfied intentional attitudes of one kind or another toward X or toward Xes — toward that which is in the extension of ‘X’ — other details of these attitudes depending on linguistic and extra-linguistic context, on surrounding words, grammatical mood, and so forth. For example, a stabilizing function of ‘weasel’ is to help implant true or satisfied intentional attitudes towards weasels. Moreover, part of the explanation of how it has normally performed that function is by corresponding to weasels in phrases and sentences that carry information about weasels or that cause attitudes or actions toward weasels. What is it, however, that has bound together the members of this extension in such a way that it can be cited, as such, as a group, collectively, in a *natural explanation* of the survival of the term ‘weasel’? We need to spell out, explicitly, how corresponding to various members of that group, *as delimited in that way*, has helped to explain survival. The way to start, I propose, is by looking for some kind of selection pressure that might impinge on individual words, helping to determine their stabilizing functions as involving empirical meanings, and in so doing, also determining in what various ways explanatory extensions for them can be determined.

There is a hint in Wittgenstein:

If language is to be a means of communication there must be agreement not only in definitions but also (queer as it may sound) in judgments. This seems to abolish logic, but does not do so. —It is one thing to describe methods of measurement, and another to obtain and state results of measurement. But what we call ‘measuring’ is partly determined by a certain constancy in results of measurement. (*Philosophical Investigations*, paragraph 242)

If we measure it, and then measure it again, we should get the same result, if I measure it again, and also if you do. That, Wittgenstein says, is part of what makes it count as *measuring* and thus, he implies, the fact that this or that particular operation can be count as a measuring operation is an empirical matter. The same thing must have the

¹² Of interest is that it does not follow that this rubric has the same stabilizing function as does the form “‘x’” has the same meaning as “‘y,’” in which the term ‘y’ is mentioned rather than used. Truth conditions do not determine functions. See, for example, Millikan 2004, Ch.7 and 2005, Ch.3.

same length if we or anyone measures it again, and the same weight and the same specific gravity. If a mineral sample is determined to be harder than another because it scratches the other, then it should still be harder if we try the scratch test again. Similarly, if our methods determine it to be a weasel they should not also determine it to be a mink, and if they determine that a sample is vitamin B₁₂ they should not also determine it to be Vitamin C, or sugar. Were it not for these consistencies, Wittgenstein claims, our empirical terms would be useless for communication.

Far more important, however, as it seems to me, the supposed concepts with which they were associated would be useless for gathering knowledge. It would not be possible to form stable beliefs, accumulating knowledge over time that could be applied on new occasions. Agreement in judgments with oneself is even more fundamental than agreement with others (Millikan 1984, chapters 15 ff., 2000 chapters 1—7). These two kinds of agreement constitute our most basic evidence that our 'methods of measurement' are measuring anything *real*.

Granted that a central function of language involves (1) the communication of (2) knowledge, then, we should expect there to be strong selection pressures weeding out terms with a persistently disruptive tendency to produce inconsistent judgments, contradictions in thought, strong disagreements among otherwise compatible persons about the truth of sentences in which they occur. For a descriptive extensional term to survive without change in how people are using it, there has to be enough internal and external agreement on its applications, enough agreement on its extension. And this, Wittgenstein has cautioned us, is not to be taken for granted, but is determined empirically.

What is deeply penetrating in Wittgenstein's view is that agreement in methods need not be sufficient to determine agreement in judgments. In the passage above he implies, however, that agreement in methods, in definitions, *as well as* agreement in judgments, is required for communication. On the contrary, I think that agreement in methods often is not necessary to determining agreement in judgments, hence not for communication, nor for acquiring stable knowledge. Moreover, as a practical matter, their own agreement in judgment with others is often the only thing actually discerned by language learners and users as a check on their usage, hence the only factor (of this kind) controlling proliferation of an extensional term's tokens. So it cannot be part of such a term's stabilizing function to implant intentional attitudes towards its extension or members of its extension (toward X or toward Xes) *as recognized in any particular way, or as thought of under any particular description*. Our basic extensional terms do not have handed-down conventional intensions. My job then is to explain *how it can be* that agreement in judgments, agreement on extension, often does not require agreement in methods.

6. Ontology

6.1 *Some ontology: real kinds*

Suppose that we were able fully to describe every individual in the natural world by locating it at a point within a single huge multidimensional logical space each dimension of which represented contiguous determinates from a different determinable range. And suppose that the world were such that every point in this huge logical space was equally likely to be occupied. That is, between any two actual individuals there would run any

number of paths in the space along which other individuals might be randomly realized. So there might be animals that were half way between rabbits and snails and others three fourths or nine tenths of the way, along any of innumerable different paths of incremental similarity, and likewise between computers and mountains, between newspapers and rivers, between eggplants and egg beaters, and between these and whatever you might dream up in your imagination. You would like to invent a language to use in this world, a rabbit is on hand, and you would like to name a kind to put it in. How will you define a leak-proof logical enclosure around this rabbit to include just exactly what you wish to put in the extension of 'rabbit'?

You might decide that a 'rabbit' must have a shape within certain limits (you should define these exactly), it must have fur and big front teeth and it must hop and eat vegetables (defined exactly). You now explain to others how you will use the term 'rabbit.' But no matter how 'sensibly' and strictly you have defined 'rabbit,' in this world knowing something is a 'rabbit' will yield no information not already directly contained in your definition. If not explicitly excluded, the next rabbit might talk Sanskrit, run by clockwork, sense its world using sonar, and broadcast radio waves carrying fake news on Tuesdays. Defining 'kind words' would have no point in such a world, for no knowledge could be accumulated around them. Three points then about kinds in the actual world.

1. As a matter of empirical fact, logical space is almost entirely empty. Huge gaps exist separating tiny clumps or clots of actual individuals or carving out miniature humps, peaks or ridges on the surfaces of these clumps. There are, for example, well over two million separate extant species of animals, but they don't generally fade into each other, let alone into shoe horns and alarm clocks. Our most basic kind words do not name classes. They name these clots or various protrusions from these clots.

2. Each of these clots or protrusions is located almost entirely along some very short path on each of a huge number of dimensions of logical space. That is, the things in each tend to be like one another in a huge number of different ways. Often a large number of these ways distinguish items in a clot from items in all other clots very reliably, yielding many alternative reliable methods of recognizing things in the clot, or reliable, perhaps, within the region of space-time wherein a certain language community lives.

3. Because the clotting is a contingent matter of empirical, not logical, fact, in strictness the extensions of basic kind terms are determined only for the actual world, or only for regions within it. The sources of these terms' meaningfulness, explaining how they can continue to serve communication and knowledge-gathering functions despite alternative methods of recognition, are rooted in the actual world. These sources are, in part, causal necessities, allowing only a few kinds of logically possible individuals actually to be realized. But they also include clotting mechanisms that determine certain individuals within specific space-time regions to be *caused* each closely to resemble the next, the restriction on displayed property mixes depending also on historical contingency. Terms anchored to these clots or to protrusions on them may lose their purchase entirely if one attempts to apply them in other possible worlds, even where these worlds have the same causal structure as the actual one.

Actual weasels, for example, are all, roughly speaking, copies of one another. They are all produced from the same gene pool, the genes having been copied from one

another or from common ancestor genes. Absent a strong change in selection pressures, any gene pool energetically exhibits homeostasis, for besides the rarity of mutation, new genes entering the pool are nearly always immediately selected out for lack of fit with the incumbents. (Mutant weasels born with a taste mainly for carrots will not survive since their digestive systems will not handle carrots. Natural selection occurs in part owing to countless subtle and less subtle variations on this sort of theme.) Further, natural conservation laws coupled with similar dynamics governing other biological species keep the historical environment in which weasels appear roughly uniform over time so that selection pressures remain fairly steady. There are good local reasons why weasels form a clump.

Weasels form a 'clump' in that they have very numerous subtle and less subtle common properties, or largely overlapping properties, often with clean gaps separating these properties, or diagnostic sets of these, from the properties of other species. They are also clumped in space and time, all being found in one roughly continuous space-time region, having appeared on earth in a definite era and being due, undoubtedly, to disappear in another. Indeed, just like an individual, the weasel clump is located in our space-time or 'historical' order essentially. It is formed, it is clotted together, by causal relations between its elements coupled with causal relations linking it to a specific historical environment. This is what has kept the word 'weasel,' despite the variety of ways people recognize weasels, both from causing disagreement in judgments, and from sliding off onto other species.¹³

The clump that is a commonly-named species may exhibit few if any properties common to all and only all of its members. In most species, for example, every distinctive gene has alleles. There can always be malformed and mutilated individuals, and many don't survive for good reason. Yet for every species there are traits without end that are possessed by nearly every member, others that are possessed by most members, others possessed by many, and so forth. And there are, for most species, very many traits and alternative sets of traits, each of which is highly diagnostic of the species. There are many different ways to tell a weasel, including by its bones. But none need be 100% reliable; certainly none are definitional.

Commonly-named species are, in general, examples of what I have elsewhere called 'historical kinds' (1999, 2000a, Ch.2, 2005, Ch.6),¹⁴ clusters whose unity is formed, hence that exist as kinds, essentially in the historical order. Historical kinds are clumps or protrusions formed through processes of duplication or reproduction, or by the role that numerically identical items have played in their causal histories, aided by principles of conservation, or they are classes or subclasses of these. Historical kinds include

¹³ Current debates about how species concepts should be defined in biology have no bearing on the above points which are all undebated empirical truths about weasels and most other commonly named macroscopic sexually-reproducing animal species in our actual world. The species debates current in theoretical biology occur, in part, because different kinds of forces can be responsible for the clusterings found in different forms of life causing unified kinds to form according to somewhat differing principles. Differing principles may form overlapping historical kinds, some may form kinds that merge together along certain quite specific dimensions making principled divisions difficult or impossible, and so forth.

¹⁴ They are not all held together in exactly the same way as weasels, hence the debates over species concepts mentioned in the footnote above.

Historical kinds are also described in Millikan 1984, Ch.16, but not under that name.

models of cars and of computers (copied one from one another or originating from numerically the same plan), folk tales, published books, Gothic architecture, peanut butter, milk, McDonald's restaurants, bus number 14 in the city of Verona (which historical kind this phrase is a description of may change occasionally) and musical pieces. (Consider, for example, what makes *The Star Spangled Banner* sung by a classroom of first graders into a rendition of *The Star Spangled Banner*. Not, presumably, what it sounds like, but what it is being copied from. (Compare Nussbaum 2006.) Looking with a sharper eye, the 'untouchables' of traditional Indian culture, physicians and lawyers in the western world (they have learned the same traditions and attitudes from the same sources and from one another) also form historical kinds. And so forth. (The point here is to illustrate the ubiquity of historical kinds. Some of the terms I've used to describe them above may, of course, have handed-down intensions that help to stabilize their usage.)

Historical kinds need not be isolated clumps; they may correspond just to easily discernable protrusions, bulges or bumps. Biological species, for example, tend to be bulges, having emerged through rather unstable and narrow necks from earlier species, and the Gothic style emerged from earlier styles and then influenced later styles, merging with them broadly to form a sort of bump or rounded peak within architectural forms. Historical kinds or subkinds may merge seamlessly with kinds that are significantly different along specifiable dimensions, one kind merging with the next to occupy a connected worm-like or amoeba-like property space. Thus words anchored in part or wholly by historical kinds often do not have clear boundaries. Extensional terms often don't need clear boundaries, however. When misunderstanding threatens, we simply use different, or a few more, words.

Another sort of kind to which words can cling are what Putnam terms 'natural kinds' (1972). Members of Putnam kinds have a common inner nature from which their more superficial properties flow. Paradigms are the chemical elements and compounds. Samples of a Putnam kind, such as a particular chemical, often have a great many properties in common, such as color, texture, degree of malleability, thermal and electrical conductivity, specific heat, freezing and boiling points, tensile strength, spectral properties, various subatomic properties, distinctive dispositions to combine or decompose chemically, and, we shouldn't neglect, contexts in which they are typically found. Accordingly, there are many different ways to identify each of these kinds. There are many ways, for example, to tell vitamin B12. Effective ways to tell may be local, and need not be infallible. I tell vitamin B12 by the label on the bottle. The label might not be as reliable as some chemical tests, but then doing the tests would surely involve someone's reading labels on bottles in the laboratory.

6.2 *Some ontology: basic properties*

I have mentioned the variety of ways that lengths and distances are recognized. Psychologists of perception study the unimaginably complex tasks that the perceptual systems perform in recognizing various other perceptual constancies, for example, in recognizing the same color again in different contexts under diverse lighting conditions, or recognizing the same phoneme (the same vocal gesture) again despite wide acoustical variation, as when formed by a baby, a woman or a man, as partially masked by other sounds or masked by echos or wind. Empirical science has taught us to

recognize the same temperature again using a variety of differently constructed instruments, and so for many other properties that our perceptual systems do not pick up naturally.

The reason we need a variety of different methods for recognizing each basic property is transparent. The causal interactions, in different circumstances, with the complexly structured and diverse media that transmit information about any distal property are complicated and extremely various. They are determined in part by the actual physical laws. They are also determined by actual surrounding conditions, conditions that are in place for purely historical reasons. Learning to reidentify the same properties, given ones' contingent and changing environment, from different perspectives, perhaps through touch, sight, hearing, taste or smell, the use of instruments, through various other kinds of inductive evidence, is an ever continuing project. Certainly, no particular way of identifying a basic property defines its extension, either in the public language or for an individual speaker, for none are infallible. All depend on supporting conditions, which are known, if at all, only empirically and which can always, in principle, change without notice. Exactly as with natural kinds and Putnam kinds, the meanings of basic property terms are anchored through practical experience in making stable judgments, finding ways to agree with ourselves and with others — experience which serves as a fallible indication (not as a final *criterion*) of their objectivity. Such property terms do not have defining handed-down methods of application. They do not have either public or idiolect defining intensions.

It is my belief that all empirical terms either fail to have any handed-down intensions, or that what is handed down are criteria the basic elements of which themselves have no defining intensions. (Arguments to the contrary are solicited.) It is for this reason that I believe that the meanings of all empirical terms are, at the most basic level, immutably embedded in the actual world.

6.3 *Some ontology: more kinds and properties*

Most of the physical substances one encounters day to day are mixtures containing various proportions of pure elements and compounds. Words for many of these mixtures name historical kinds, for example 'milk,' 'wood,' and 'wine.' These kinds are 'historical' because the actual limits on variation in chemical constitution among samples of each of these substances, hence the overlap in properties that they display, results from the way the sources of these samples are intertwined in the causal order. For example, although milk from different mammalian species varies in a number of respects — what nourishes one mammal might starve another — still, the properties of milk overlap considerably, owing not to a semantic rule but to their causally connected origins in the historical order.

Other names for mixed substances, and also many adjectives and adverbs, are anchored to 'peaks.' The peaks themselves are described using modifiers like 'pure,' 'perfect,' 'absolute,' 'complete,' 'total' and so forth: 'pure water,' 'a pure tone,' 'perfectly straight,' 'a perfect fifth,' 'absolutely empty,' 'completely impermeable,' 'totally bald' and so forth. These are natural end points when approaching, often, from any of various property directions (and the recognition of them, or of approximation to them, can often be accomplished many alternative ways, under alternative conditions and so forth). Pure physical elements and compounds are peaks that anchor many words for common

substances. Thus water is whatever is close enough to pure water for the communication needs of the moment. Basic color words in the various languages, these being added to languages in a definite order or very close, correspond, roughly in order of addition, to the height of natural peaks in discrimination distance for the human visual system from neutral grey (Jameson and D'Andrade 1997). What is red or yellow is whatever is close enough to pure red or pure yellow for the communication needs of the moment. Historical kinds may also serve as somewhat rounded peaks in anchoring adjectives and adverbs, for example, 'peach' (the color), 'waxy' and 'furry.' Words anchored by peaks are essentially vague. Their meanings are described with reference to these peaks, but their proliferation has not depended on use within any set distance from these peaks. Useful closeness depends on context. 'The round red one,' in one context, might not be 'a round red one' in another where it competes with things more perfectly round and red. The set of actual historical applications that have helped account for the proliferation of a peak-anchored term will have certain statistical properties but no definite cut off point. This characteristic of vague terms is perfectly simple to understand and makes it obvious, I believe, why the logical moves that produce sorites paradoxes are inappropriately applied. The meanings of these words are not stabilized either by natural or by conventional boundaries between them but by comparative closeness, in context, to naturally stable reference points.

In sum, agreement in applications hence in the judgments expressed with many property and kind terms, agreement both with oneself and with others, rests not on agreement in methods but is an empirical matter, resting on natural law and on clumping and humping and peaks in nature. Hypothesizing that the most basic stabilizing functions accounting for the survival of language that employs empirical terms are helping to carry information about the empirical world and directing attention or attitudes or action towards aspects of it, as intended by speakers, it appears that the stabilizing functions of these terms are to help carry information about these natural peaks and clots and the things that compose them, to direct attention or action towards these or members of their extensions, and so forth. It is because of the clotting that enough agreement occurs in practice to coordinate speakers and hearers of these terms in projects of common interest, hence to keep these terms in circulation. Depending on time and situation, speakers may use any of a very wide variety of indicators to stimulate their applications of an empirical term, every one of these applications constituting merely in an implicit *hypothesis* that what they are identifying is the same property or another part of the same clot. Some of the methods used for a given property or kind may be more reliable than others, but none are definitional. Neither speaker nor hearer generally has reason to care how the other has managed to diagnose a basic property or kind. Neither cares, for example, which perceptual feature-detectors the other has employed to tell weasels from mink, or even red from blue.

What has been accounting for the clotting that gives rise to a certain kind, or accounting for agreement about a certain property, may remain unknown or gravely misunderstood by the users of its name, though subject to empirical investigation. It is by sophisticated empirical methods that we have learned, for example, about gene pools, about the inner natures of Vitamin B₁₂ and sugar, about the natures of temperature and of distance (the nature of space). If the underlying cause supporting agreement for such a term is univocal — if the members of the extension are like one

another all for the same reason — then the term is univocal. It is not ruled out by logic, however, that separate causes of clotting should produce clots occupying the same or closely overlapping areas of property space. Should this happen, so that one portion of such a term's extension resembles the rest by sheer accident, then the term is equivocal, though the term's users may not know that. The explanation for such a term's survival in the language would not be univocal, so its stabilizing function, its meaning, would not be univocal either. Putnam's example of 'jade,' which equivocally names both jadeite and nephrite, is an instance of this kind (see section 11 below.) Before Newton, the term 'weight' was equivocal between weight and mass.

7. Some concrete results

Putnam's most famous example of a natural kind is water, and there are, of course, a great many different (fallible) ways to tell water. What has been holding the term 'water' from slipping this way and that all around n-dimensional property space, what has been supporting agreement in judgments, is not some criterial way everyone has been taught for recognizing water. It is the contingent but lawful clustering of distinctive traits of water all produced by the same molecular structure. Recall Putnam's fable about Twin Earth, which is identical in all ways to Earth except that what they call 'water' on Twin Earth has the molecular structure XYZ. XYZ resembles water in every surface respect by which water was identified on Earth in 1750, before modern chemistry determined water to be H₂O. On the analysis I have suggested, as on Putnam's own, XYZ was not within the extension of our word 'water' even in 1750. The stabilizing function, the current meaning, of a word rests on what has, as a contingent matter of fact, been holding its usage in place, effecting agreement among users and for users with themselves, despite the use of a variety of alternative recognition techniques. Dispositions that people might have to apply a term in new circumstances that would accidentally continue to support agreement but for new reasons are not relevant to its current meaning. On Earth in 1750 the deep structure XYZ had born no responsibility for stabilizing the earth concept or term 'water.' It was not because XYZ could manifest itself in the many alternative characteristic ways by which various people at various times had been identifying water, or in terms of which people had been thinking of water, that they had managed to agree with themselves and others in judgments, but because H₂O manifested itself in these various ways. If earthlings had gone up to Twin Earth in 1750, they would, of course, have mistakenly applied the Earth term 'water' to XYZ, leaving us a legacy in which 'water' had become referentially equivocal. But no one did go up to Twin Earth in 1750, so we were not left any such legacy. What might have accounted for or might in future account for a word's proliferation is not what *has* been accounting for it, and is in no way relevant to its current meaning.

Notice that this argument for Putnam's result doesn't depend, as his did, on some special linguistic rules we have all learned determining that 'water' must always name 'the same liquid' and legislating some common way that we should all understand 'same liquid.' The reference of 'water' is entirely *direct*. It has no defining intension. Keep in mind also that my methodology has been entirely different from Putnam's, who used example and counter-example.¹⁵ I began with various apparent data, formed an

¹⁵ The Twin Earth examples used here and in what follows are not used, as Putnam used them, as examples with which to *test* the theory of meaning on offer but used, rather, to clarify it. I am not interested in invoking

explanatory hypothesis to explain them, and I am now considering the consequences.

Earlier I mentioned that historical kinds are firmly lodged in the causal-space-time order in the same way that individuals are. Similarly, Michael Ghiselin (1974, 1981) and David Hull (for example, in 1978) claimed that species are not classes but instead are like big scattered individuals. The members of a species form a blood-related family, not a class. What makes various portions of an individual's history into portions of one and the same individual's history are the space-time-causal relations among them, determined by conservation laws and principles of homeostasis that preserve many of the individual's properties over extended periods of time. This is empirically known, not legislated or known *a priori*. There have been folks who believed that an individual person could, in certain magical cases, skip over gaps in space or time and then reappear, perhaps even in two places at once. These folks were mistaken; they were not just speaking a different language. The idea that individual objects are, as Sellars put it, 'space-time worms,' is not a matter of conventional definition but a matter of how nature actually works. Similarly, what makes various weasels into members of one and the same historical kind are the space-time-causal relations among them determined by reproduction, by principles of homeostasis, by conservation laws that keep the environment stable and so forth, thus preserving the species' properties over instances and over time. In neither case, for example, is it mere overlapping likenesses in inner molecular structure (or the mere sharing of any other area within property space) that defines the relevant unity. Sameness of inner structure is not what ultimately underlies either the overlapping likenesses of an individual's time-stages or the overlapping likenesses of members of commonly-named animal species. *Causal connections*, conservation, homeostasis and so forth — processes connecting what's in one historical time and place with what's in the next — *are what determine that there are overlapping likenesses of inner structure, not the other way around.*

Keeping this similarity of individuals to species well in mind, consider the following new Twin Earth fables.

Let Twin Earth be exactly as Putnam described it, except that its 'water,' too, is H₂O. Like Earth in every detail, it contains doubles of you and me, and of Bill Clinton, Hillary Clinton and Chelsea Clinton. Now let Earth Chelsea telephone Earth Hillary to ask after her father's plans for the day. By some cosmic accident, she reaches Twin Hillary instead (after all, Twin Hillary has the very same telephone number). Twin Hillary and Earth Chelsea agree that Daddy got too much sun at the beach with them last weekend and Twin Hillary warns that he is still feeling a bit crisp. She gives Earth Chelsea Twin Bill's schedule for the day, mentioning what he is wearing and where he has planned lunch. Subsequently, Earth Chelsea finds her father at lunch, confirms that he is wearing a light tan suit with no tie, that he is still a bit sunburned and so forth. She agrees with Twin Hillary in these various judgments — or, more cautiously, she approves of various like-sounding sentences. Assuming that the likeness of Twin Earth to Earth is completely accidental, however, nothing *explains* this agreement. It is completely accidental. Had Chelsea been talking to her own mother instead of to Twin Hillary, there would have been a good explanation for agreement. Contact with the

intuitions about what terms like 'water' and 'weasel' mean as evidence for the theory. Indeed, in the case of the term 'weasel' (coming up next) I am likely to tromp heavily on the intuitions of many people.

same man, holding on to his own properties through space and over time, would have been responsible for the women's agreeing judgments. Similarly, Earth Bill Clinton himself has been responsible for the continued largely-agreeing uses — uses not causing widespread contradictions — of our conventional name for him. The name 'Bill Clinton' that we use here on Earth has in its extension only one man. The same-sounding name on Twin Earth is a homonym.¹⁶

Now compare weasels. Suppose farmer Jones, here in Connecticut, telephones state agricultural consultant Barnes for advice on preventing the weasels from killing more of his chickens, but he reaches Twin-Barns instead. They talk for some time, agreeing, apparently, on many things, for example, that weasels seem to enter through holes hardly big enough for a mouse, that they kill by slitting the chicken's throat, that weasels may kill many more chickens than they eat, that they can't be caught in ordinary traps, that they come in the night and leave a musky smell, and so forth. Suppose also that consultants Earth Barnes and Twin Barnes recognize weasels by their teeth and toes and bones and so forth while farmer Earth Jones, like myself, would not have known a weasel had the cat brought it in but instead recognizes the smell, the characteristic way of killing, and so forth.

Now if Earth Jones were talking to Earth Barnes, their agreement in judgments would be explainable. They might never have encountered any of the same individual weasels, nor do they use the same ways of identifying weasels. But there is a good reason why the weasels they have encountered, and those encountered by others in the Earth language community, have had so many properties in common. Agreement in judgments on earth-weasels has resulted from the ways these weasels are causally linked through reproduction, homeostasis in the gene pool and steady selection pressures in the numerically same ongoing environment. But there are no causal connections between earth-weasels and twin-weasels. Of course one earth-weasel is like another for the same reason one twin-weasel is like another; the mechanisms causing the clotting are identical. Compare the connections among the life-stages of each of the two Bill Clintons. But an earthling weasel is like a twin-weasel for no reason at all. Of course, we might tell a story about why Twin Earth is exactly like Earth. That will make earth-weasels like twin-weasels for a reason, but not for the same reason earth-weasels are like one another. It will also make life stages of the two Bill Clintons alike for a reason, but not for the *same* reason that each stage of Earth Bill Clinton is like the next. The stabilizing function of our word 'weasel,' what currently supports the continuing proliferation of its tokens, is conveying information about earth-weasels. Its twin word would be a homonym.

Again, this result doesn't depend on some linguistic rule we have all learned requiring 'weasel' to be a historical-kind term. Aristotle's term for weasels (supposing he had one) was also bound to the weasels on Earth, even though what he *thought* bound them together was a common Aristotelian form, so he would mistakenly have *thought* twin-weasels were just weasels.

But perhaps you also think that twin-weasels would just be weasels? Then you too

¹⁶ There may also be homonyms of 'Bill Clinton' on Earth, of course, if there are other people with a same-sounding name.

are mistaken about what makes weasels be weasels!¹⁷ Will you reply that you jolly well know what you mean by the word ‘weasel,’ and there simply is no such limitation on how you would apply it? My claim is that you are not in charge of the meaning of your word ‘weasel.’ What it means isn’t determined merely by how you or anyone else is disposed to apply it. It has a definite meaning only in so far as ‘the results of measurement [with it] are constant’ *for a reason*. What it means, if anything, is determined by what has actually been responsible for the constancy, the consistency, of those results — presumably the weasels themselves, an individual self-perpetuating, earthly clump, constrained by properties of its actual, contingently ongoing environment.

8. Knowing the meanings of words

Are we to conclude then that Aristotle didn’t know what ‘weasel’ (its Greek equivalent) meant because he thought weasels were whatever had a certain Aristotelian form? What is required to ‘know the meaning’ of a basic extensional term?

It is fairly usual in the philosophical tradition to assume that ‘the meaning’ of an extensional empirical term is what determines its extension. On the above analysis what determines the extension of a term is something that helps to explain how it happens that enough of the time, people have not disagreed on its applications either with themselves or with one another. Sometimes part of the explanation may be that usage has been partially channeled through an intension handed down from user to user, or among those heavyweights who have been influential, for one reason or another, in determining the course of its more general usage. But I have argued that the explanation always makes reference, in the end, to what is not found in the organization of the heads of users, but in the structure of the empirical world outside. At the bottom level, then, what determines its extension is not the knowing of anything by its users. Moreover, for terms whose usage is guided in part by intensions, the fact that they are so guided is a matter of the details of actual social history and influence, not a matter ultimately discerned merely through armchair reflection. People can fail to have any grasp at all of who the ‘experts’ are influencing the course of usage of a term that is still firmly within their (correct) usage. It would be interesting to survey, for example, how many ordinary Americans know that a ‘church’ ‘has to be’ a place where those who believe in the divinity of Christ worship, and among those who can’t say, how many of them can say where to find the experts on the matter. (Most will probably say, ‘the dictionary’!) Still, it is likely that effectively all of these people’s actual day to day uses of the word ‘church’ are to places of worship used by Christians.

Because it is a matter of detailed social-psychological history whether or not a particular empirical term or phrase actually is basic or whether it has an intension, and if it has one, by what means of transfer that intension manages to channel the bulk of its actual usages, I cannot, of course, offer a general criterion by which to distinguish a priori which terms are basic and which have or tend to have intensions. We can only make intelligent guesses. For example, I have just made a guess about the word ‘church.’ I would also guess, for another example, that terms for tools and many other artifacts are governed in part by intensions determining their extensions to contain

¹⁷ Again, the question what makes weasels be weasels should not be thought to hinge on the outcome of current controversies about how ‘species’ should best be used as a technical term in biology. The meaning of the traditional term ‘weasel’ is a completely independent matter.

things that have been made for certain uses, and that by some means this becomes understood by nearly all who employ or effectively understand these terms.

What then is usually involved when we say that a person ‘knows the meaning’ of a certain empirical term, for example, when it is said that sixteen year olds know the meanings of 45,000 words? The close analogy that I described above between the way individuals and historical kinds are each clotted together suggest an answer that seems to apply to many empirical terms. Knowing what it is that bears a certain common name may often be like knowing who it is that bears a certain proper name, for example, like knowing who Nicolo Paganini was. What counts as ‘knowing who’ depends on the purposes at hand (Boër and Lycan 1986). Knowing just about any definite description or alias of a person might count as knowing who that person is in some pragmatic circumstances. Being able to recognize that person or to find them readily may also count. Yet other times, ‘Who is that clown?’ may be an appropriate question even when we are staring right at him. Knowing an indefinite description, such as ‘a famous violinist,’ will also serve sometimes for knowing who. That is who Nicolo Paganini was. In like manner, ‘knowing what ... ,’ as in ‘knowing what a weasel is’ might require as little as knowing that weasels are small mammals, or that they are the animals that get into chicken coops and slit throats, or it might require as much as being able to identify them in the field. The thing that it cannot require is the ability to give an account of the rules by which competent English speakers apply the term ‘weasel,’ because there are no such rules.¹⁸

Perusal of any same-language dictionary shows clearly that ‘knowing meanings,’ in the dictionary’s sense, need not involve either being able to *give* the meaning (as required for filling in the “‘X’ means Y’ rubric), or knowing ways that most ordinary users recognize extensions. Dictionaries are often content merely to *describe* extensions — the clots — mentioning salient facts about them, just as one might *describe* a person to tell ‘who they are.’ Dictionaries may also offer descriptions currently accepted by the scientific community.

barley *n.* 1. A grass in the genus *Hordeum*, native to temperate regions, having flowers in terminal, often long-awned spikes. (American Heritage Dictionary)

Vitamin _{B12} *n.* A complex compound containing cobalt, found especially in liver and widely used to treat pernicious anemia. (Ibid.)

Compare:

Lincoln, Abraham, 16th president of the U.S. (1860-1864) Assassinated (Webster’s International Dictionary, second edition)

¹⁸ Many more details are in (Millikan 1984, Ch.9). I have also argued that when language use proceeds in the manner that keeps language forms reproducing, words carry ‘locally recurrent natural information’ of the same kind that is the staple of ordinary perception (2000a, Ch.6, 2004, chapters 3—11, 2005, Ch.10). Understanding descriptive language, when all goes well, is merely one more form of perception of the world, indeed, a form of direct perception, in the same sense (which needs to be carefully described) in which vision is normally ‘direct.’ It follows that learning the significance of various linguistic forms is, in many ways, just like learning the significance of various visual forms or tactual forms or auditory forms. We don’t marvel that children can learn so easily what thousands of things look like!

Paganini, Nicolo, Italian violinist (ibid)

Notice that the definition of 'barley' does not guarantee that no other plants fit this description as well. And if vitamin B₁₂ ceases to be widely used to treat pernicious anemia, this will not change the meaning of the term 'vitamin B₁₂.'

brown n. Any of a group of colors between red and yellow in hue that are medium to low in lightness and low to moderate in saturation. (American Heritage Dictionary.)

vinegar n. An impure dilute solution of acetic acid obtained by fermentation beyond the alcohol stage and used as a condiment and preservative. (ibid.)

These descriptions do not attempt to show how people typically know to apply these terms or what comes to mind when they are used. Notice too that these definitions could not be used in the "X" means Y' rubric. It would be very odd, for example, to translate the German word 'brun' with the above definition for 'brown.' Nor are common ways of thinking about or recognizing members of a naturally-bonded extension mandatory ways, ways that are built into the language, ways essential for agreement in usage. After all, English was the language that Helen Keller understood and used. There are various catadromous species of eels (ones that breed in the ocean but spend the rest of their adult lives in freshwater) for which no one, until fairly recently, had identified the larval stages as being of the same species as the adults. When it was discovered what these larvae became when grown, the definitions of their names given in dictionaries changed but in the common way of speaking their meanings did not change. Nor did the way the extensions of their names were delimited change. What changed was only empirical information had about the kind.¹⁹

9. Disclaimers

Before moving on, I want to make quite explicit that in advocating the above views I am not denying any of the following. (1) Without doubt there is a high degree of overlap among the various ways that most adults are able to recognize the extensions of the most common extensional terms, such as 'red' and 'cat' and 'applesauce.' (2) There is sometimes universal agreement among adults that certain ways of recognizing the applications of a basic extensional term are better, even uniquely better, than others, for instance, looking at colors in broad (but not too broad) daylight. (3) Scientists sometimes attempt to stabilize scientific terms by reference to 'operational definitions,' such as the galvanomic skin response as an indicator of increased emotion, and such

¹⁹ With the example of these catadromous species in mind, consider the usual definition given by dictionaries for 'vixen,' namely 'female fox.' Does this definition correspond to an intension carried by the term 'vixen' or does it correspond, merely, to something just about everybody knows about vixens? Not only are vixens usually thought of as female foxes, so far as we know, being a female fox corresponds to a 'real definition' — to part of the Lockean real essence — of the vixen. It is the double x chromosome, the female-making genes, that produce the various properties, including size, shape and behavior, by which vixens can be told from dog foxes. But if there were peoples who recognized the difference between vixens and dog foxes only by, say size, general appearance and behavior, but had no general concept for foxes and no general concept of femaleness, their name for vixens would still be correctly defined as 'female fox,' and it would also be natural to say their word meant the same as our word 'vixen.' This suggests that what makes 'vixen' mean the same as 'female fox' is not or not merely that these have the same intension, but the fact that pretty much everyone implicitly knows that being a female fox corresponds to a real definition of the vixen. Contrast the definition of 'brother' as 'male sibling.' Brothers do not form a real kind as such. The word 'brother' does seem, quite definitely, to have being a male sibling as a conventional intension.

definitions can sometimes serve as fairly reliable indicators of the presence of certain real properties or kinds although, of course, there is no guarantee. When better indicators of the same thing are discovered, however, these ‘definitions’ tend to be replaced. (4) There can be terms — classical examples are ‘phlogiston’ and ‘witch’— that only *appear* to be extensional terms, being in fact empty of extension or highly equivocal. A variety of supposed methods of identifying their extensions are handed down from one person to another, but the underlying cause of their proliferation has no firm basis in nature. (5) There are, of course, also fictional terms such as ‘Santa Claus’ and ‘unicorn,’ terms that we pretend are extensional terms. We may hand down pretend useful methods of identifying (Santa’s outfit and beard), pretend descriptions (he often says ‘Ho ho ho!’), perhaps even pretend intensions, for their pretend extensions.

10. Mistaking one’s own meanings

A public language can contain words whose extensions are bonded in ways that the people using the language know nothing about. I have mentioned that both Aristotle and Putnam were wrong about what it is that clots the members of biological species together (section 2 above). In early Greek times, many believed that a new sun arose every day. Thus they believed that what actually is an individual was a kind. Still, the actual bonding underlying usage was the same for their word as for ours. It was the self-perpetuating nature of an individual sun that was responsible for their continuing agreements with themselves and one another on, for example, when suns could be expected to rise and set, when they would be higher or lower on the horizon, and so forth. Some people have wrongly believed that in magical cases, a person can make leaps over space or time and possibly even become double. Suppose that Mack believes this. He might well insist then that your double whom he meets on Twin Earth was really also you. But if Mack has dubbed you ‘Mick,’ and in fact knows nothing about Twin Earth (so the extension of his ‘Mick’ has not become equivocal), still the extension of his ‘Mick’ would include only you. Compare young children, who have very poor if any grasp of time concepts before about age four. Although they can’t yet understand the space-time-worm principle, hence carry no space-time continuity ‘criterion of identity for persons’ around in their heads, surely they mean exactly the same thing that their mothers mean when they say ‘Daddy.’

Consider the name ‘aurora borealis’ that the ancients gave to the northern lights. Most likely they would have mistaken for the *aurora borealis* the spotlights sometimes lighting up the sky at grand openings of shopping malls. It doesn’t follow that the light from these spotlights was in the extension of their term ‘aurora borealis.’ This is not because there was a conventional rule, somehow passed from father to son, requiring that ‘aurora borealis’ refer always to ‘the same underlying cause’ as of ‘those lights up there’ (compare Putnam’s ‘same L as what’s in the lakes and streams around here’). Nor is it merely the fact that their utterings of the term ‘aurora borealis’ were caused by something other than spotlights. (I am not advocating a causal theory of reference.) Their utterances were caused by emanations from the sky of changing patterns of diffused light, and spotlights produce such emanations as well. The reason is that the stability in judgments the ancients were able to make about the *aurora borealis*, the things they were able to confirm about it and agree on, such as its typical shapes, colors and brilliance, ways of moving, where and at what time of year it could be seen, were

stabilized by the bonding tie that is the stability of the Earth's magnetic field lines and their effect on the electrons that travel through. Obviously the ancients had no inkling of the nature of this bonding tie.

The word 'jade' was introduced to name the jadeite carved by Aztec and Mayan artisans in Mexico before the advent of modern chemistry. Later nephrite, which has startlingly similar superficial properties but is completely different chemically, was mistaken for jade.²⁰ This was a mistake — a mistake that soon rendered the term 'jade' equivocal, or, for those who now know about the double meaning, homophonous. 'Jade' is not and has never been just a class term for anything with such and such superficial properties. There are, for example, other substances superficially like that are not called 'jade'. There are two different accidentally converging reasons for the agreement in judgments that still partially sustain use of the term 'jade.'

11. Undermining one's own meanings

Much of traditionally conceptual analysis has proceeded by the a priori method of example and counter example, on the assumption that what determines the extension of a term is projected from the understanding a competent user must have of a handed-down intension. Challenges to proposed definitions are constructed by imagining situations in which one would naturally apply a term to be analyzed but which don't fit the definition proposed, or vice versa. This method is doubly vulnerable. It is vulnerable because a competent user may not correctly understand what is holding together the natural feature, clump, hump or peak to which the term is actually anchored (as with Aristotle and biological species). Second and more insidious, situations may be imagined in which the real-world foundations, the contingently existing clots, that support our words have been dissolved away. We imagine worlds in which the words would, in fact, have lost their meanings, not because they were no longer associated with any methods of application, but because that which underpins 'constancy in the results of measurement' given these methods has been destroyed. We may mistakenly suppose that sentences in our current language are capable of expressing propositions that are 'functions from' *even these* 'possible worlds to truth values.'

Suppose, for example, that although conditions for perception were perfectly normal, some object that looked square when close up and had equal diagonals when measured with a ruler always or often felt triangular, looked octagonal from a distance, and rolled on the ground the way wheels do. Would this object be square? Could *any at all* of the ways that various of us regularly recognize squares give wrong results for an object and yet it be square? Exactly how many of our ways could fail? Exactly which ways? The craziness of these questions results from supposing that contingent features of our natural environment that anchor the very meaning of the word 'square' have been destroyed. It's not just that it is nowhere written which ways of recognizing squares are 'definitional' or 'criterial.' The possible-worlds game simply crashes when you dissolve too much real-world glue. Language and thought are not a layer of ghostly ectoplasm that can be peeled off the real world and laid overtop any possible world at will. The integrity of extensional language and thought is maintained through deep roots in the actual world. Both language and thought — their meanings — will evaporate if one tries to relocate them in worlds that are too far away.

²⁰

The Encyclopedia Britannica 1967

It is particularly easy to undermine meanings when examining terms anchored by clots of an essentially historical nature. These terms are tied not just to the causal structures of the actual world, but to local *historical* contingencies of the world in which the languages containing them have grown up. Having, in imagination, kicked away the planks on which such a word is standing, we argue pointlessly over where it is destined to land. It is true that changes that sever basic referential bonds for words sometimes actually occur. Historical clots in property space may grow or shrink, change shape, appear and disappear. Rapid technological evolution in modern times has required many new decisions to be made on word meaning. But to long for a vocabulary that would fit any world, one that could be used, for example, to carve out meanings in the frictionless world of section 6.1 above, is not coherent, and to suppose that the vocabulary we already have could serve this purpose is quite mad.

12. Coda: On Swampman

Suppose lightning strikes a dead tree in a swamp; I am standing nearby. My body is reduced to its elements, while entirely by coincidence (and out of different molecules) the tree is turned into my physical replica. My replica, The Swampman ... moves into my house and seems to write articles on radical interpretation. No one can tell the difference.

But there is a difference. My replica can't recognize my friends; it can't recognize anything, since it never cognized anything in the first place ... I don't see how my replica can be said to mean anything by the sounds it makes, nor to have any thoughts. (Davidson 1987)²¹

There is a considerable literature on whether Davidson was right about this.²² Several authors, most recently Tanney (2004) and Jackson (2006), have been explicit in arguing that, given our ordinary ways of applying terms for mental events, hence given their ordinary meanings, Davidson had to be wrong. Other authors think that Davidson was right. I will use Davidson's Swampman example to illustrate a possible-worlds case in which the determiner of extension for a term has been dismantled, so that the term no longer has clear empirical meaning. I will argue that in this example, Davidson has removed the natural planks on which our terms for mental events are resting, thus undermining their meanings.

To see this, begin by noticing that for roughly the same reason Swampman is not Davidson, he also is not, in the ordinary sense, a man. It is not the arrangement of molecules that determines something to be either another time-stage of Davidson or another member of his species. The various existent, detailed arrangements of molecules that form actual humans are clumped in logical space throughout an irregular and porous area without logical boundaries. Recall the absence of any distinctive

²¹

Compare also:

Suppose that by some cosmic accident a collection of molecules formerly in random motion were to coalesce to form your exact physical double. ... that being would have no ideas, no beliefs, no intentions, no aspirations, no fears, and no hopes ... This because the evolutionary history of the being would be wrong. (Millikan 1984, p. 93)

²²

See, for example, the articles in *Mind and Language* 11.1, 1996 (including my own, to be improved on in this section, I hope), Johnston 1987, Tanney 2004, Jackson 2006.

genes, the genetic deformities, the mutilations caused by accidents or war. Were it not for the way natural causes actually determine a contingent clumping of properties in a certain space-time region, there would be nothing to determine the limits of humans. Local, natural, this-world history is what defines, in its sloppy way, what is and is not human, defining it for this world only. Notice also that although the usual kinds of inductions running either from Davidson to Swampman or from other humans to Swampman would come out true, they would come out true for no reason. Without having examined Swampman molecule by molecule and without having a molecule by molecule understanding of how Davidson works or how humans work, we could not rationally predict anything at all about Swampman — that he wouldn't suddenly start speaking Mayan, or light up blue at night, or self destruct when the clock struck twelve.²³ That people's various ways of attempting to identify humans and to obtain and use stable knowledge about them would indeed work in the case of Swampman would be in each case another total accident. Of course we might decide we should treat Swampman as a human under moral laws and under the laws of the nation. We might also decide to call him 'Donald Davidson' and hold him responsible for paying Davidson's taxes. But this is not relevant to our current uses of either of the term 'man' or the name 'Donald Davidson.' Speaking quite accurately, in the world of Swampman, the very meaning of 'human' has been undermined, just as the meaning of 'square' has been undermined in a world where things that look square can roll.

Consider next whether Swampman has eyes.

eye, n. 1. An organ of vision or of light sensitivity (American Heritage Dictionary)

organ, n. ... 2. Biology. A differentiated part of an organism, such as an eye, a wing, or a leaf, that performs a specific function (ibid)

Organisms are, as such, members of historical kinds, and their parts have specific functions (in the sense meant here), such as vision or light sensitivity, in so far as their performance of these functions has helped the organism to survive or to reproduce. Thus eyes are kind members within a certain *class* of historical kinds. Swampman's pair of glossy forehead protuberances are not, in the ordinary sense, eyes.

What can we say then about the ontology of mental kinds? Taking beliefs as our example, there has been considerable controversy about whether to say that any animals other than humans have beliefs. However, human beliefs themselves seem to form a historical kind. There has been enough agreement on applications of the notion of belief ('believes that p,' 'believes that Q,' ...) to humans to keep the term 'belief' and its variants in wide circulation despite the great variety of ways we use, of course always fallibly, to identify what beliefs people have. We tell, for example, by noting any of a great variety of behaviors, by hearing what the people themselves say, by knowing where these people have been and what they have done or, in our own cases, by introspection. Something, it seems fair to assume, is accounting for this 'constancy in the results of measurement' — something, presumably, about the human belief-fixing and belief-using systems which, themselves being consolidated as historical kinds, have consolidated the kind *human belief*, putting ragged limits on the variety of its

²³

On the fallacy of supposing rationality to be a projectable property merely as such, see Millikan 2000b.

manifestations. The question about other animals is whether any of them have anything sufficiently like human beliefs to put these all in a single *class* together. This would be a *class* of historical kinds.

What human beliefs and other intentional attitudes are like apart from manifesting themselves the ways by which we usually know them is a matter for empirical inquiry, as is whether any other animals possess anything similar. No matter how this inquiry turns out, however, Swampman is not a candidate for having beliefs. Or, more accurately, in the world of Swampman the ordinary meaning of 'belief' has been undermined; the word is left hanging. Similar arguments would apply to Davidson's words 'recognize,' 'mean' and 'thought.'

Thus the intuitive outcome, in the armchair, of Swampman thought experiments is not sufficient evidence against historical theories of intentionality. This is a case where the a priori method of example and counter example fails.²⁴

RUTH GARRETT MILLIKAN
Philosophy Department U-54
University of Connecticut
Storrs, CT 06269-2054 USA
ruth.Millikan@uconn.edu

References

- Antony, Louise 1996: 'Equal Rights for Swamp-persons'. *Mind and Language*, 11.1, pp. 70—75.
- Ayer, Alfred J. 1936: *Language, Truth and Logic*. London: Victor Gollancz Ltd.
- Bloom, Paul 2000: *How Children Learn the Meanings of Words*. Cambridge MA: The MIT Press.
- Boër, Steven E. and Lycan, William G. 1986: *Knowing Who*. Cambridge MA: The MIT Press.
- Chomsky, Noam 1995: 'Language and Nature'. *Mind*, 104, pp.1—61.
- Davidson Donald 1987: 'Knowing One's Own Mind'. *Proceedings and Addresses of the American Philosophical Association*, 60, pp. 441—58.
- Fodor, Jerry A 1998: *Concepts; Where Cognitive Science Went Wrong*. Oxford: Oxford Clarendon Press.
- Fumerton, Robert 1989: 'Russelling Causal Theories'. In C.W. Savage and C.A. Anderson eds., *Rereading Russell: Essays in Bertrand Russell's Metaphysics and Epistemology, Minnesota Studies in the Philosophy of Science*, 12, pp. 108—18. Minneapolis: University of Minnesota Press.
- Ghiselin, Michael T. 1974: 'A Radical Solution to the Species Problem'. *Systematic Zoology*, 23, pp. 536—44.
- 1981: 'Categories, Life and Thinking'. *Behavioral and Brain Sciences*, 4, pp. 269—83.
- Hull, David L. 1978: 'A Matter of Individuality'. *Philosophy of Science*, 45, pp. 335—60.

²⁴ This material was originally prepared for the Jack Smart Lecture and a follow up at Australian National University. I am grateful for that invitation and for the ensuing discussions. I have also learned from audiences at MIT (especially the graduate students!), at the Universities of North Carolina, Tel Aviv and Queens University Belfast, from my own graduate students, from the faculties of Carleton and St. Olafs, and from Crawford Elder, Michael Lynch and Dan Ryder. Tom Baldwin and two readers for *Mind* added important suggestions.

- Reprinted in Sober 1994, pp.193—216.
- Jackson, Frank 2006: 'The Epistemological Objection to Opaque Teleological Theories of Content'. In G. Macdonald and D. Papineau eds., *Teleosemantics*. Oxford: Oxford University Press.
- Jameson, K. A. and D'Andrade, Roy G. 1997: 'It's Not Really Red, Green, Yellow, Blue: an Inquiry into Cognitive Color Space'. In C. L. Hardin and L. Maffii, eds., *Color Categories in Thought and Language*, pp. 295—319. Cambridge UK: Cambridge University Press.
- Johnston, Mark 1987: 'Human Beings'. *Journal of Philosophy*, 84, pp. 59-83.
- Millikan, Ruth Garrett 1984: *Language, Thought, and Other Biological Categories*. Cambridge MA: The MIT Press.
- 1996: 'On Swampkinds'. *Mind and Language*, 11.1, pp. 103—17.
- 1999: 'Historical kinds and the special sciences'. *Philosophical Studies*, 95, pp. 45—65.
- 2000a: *On Clear and Confused Ideas*. Oxford: Oxford University Press.
- 2000b: 'Reading Mother Nature's Mind'. In D. Ross, A. Brook, and D. Thompson, eds., *Dennett's Philosophy: a Comprehensive Assessment*, pp. 55—75. Cambridge MA: MIT Press.
- 2001: 'Cutting Philosophy of Language Down to Size'. In A. O'Hear, ed., *Philosophy at the New Millennium*, pp. 125—40. Cambridge: Cambridge University Press. Reprinted as Chapter 7 in Millikan 2004.
- 2004: *Varieties of Meaning: The Jean Nicod lectures 2002*. Cambridge MA: The MIT Press.
- 2005: *Language: A Biological Model*. Oxford: Oxford University Press.
- Nussbaum, Charles 2007: *The Musical Representation*. Cambridge MA: The MIT Press.
- Putnam, Hilary 1973: A Meaning and Reference—. *The Journal of Philosophy*. 70, pp. 699-711.
- 1975 'The Meaning of "Meaning"'. In Keith Gunderson, ed., *Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Science*, 7, pp. 131—93. Minneapolis: University of Minnesota Press.
- Sober, Elliott 1994: *Conceptual Issues in Evolutionary Biology*. Cambridge MA: The MIT Press.
- Tanney, Julia 2004: 'On the Conceptual, Psychological, and Moral Status of Zombies, Swamp-beings, and Other "Behaviorally Indistinguishable" Creatures'. *Philosophy and Phenomenological Research*, 69.1, pp.173—86.
- Wittgenstein, Ludwig 1953/2001: *Philosophical Investigations*. Oxford: Blackwell Publishing.