Knowing the meaning of a word: shared psychological states and the determination of extensions

Timothy Pritchard. Published in *Mind and Language* 2017, vol. 1.

Abstract: What is it to know the meaning of a word? The traditional view is that it involves the possession of a concept that determines the extension of a word, with the concept corresponding to a single psychological state. Millikan criticizes this view, denying not only that concepts determine extensions but also that sharing a concept means sharing a psychological state. The purpose of this paper is to defend a modified version of the traditional view. I argue that Millikan's claims do not translate directly into a thesis about word meaning. Her arguments relate to an extra-linguistic approach to extension, which we can distinguish from a linguistically oriented notion of extension.

Understanding a word has often been viewed in terms of the possession of a shareable concept, such that (at least ideally) those who use a word will have the same concept as set down by some adult standard. This shared concept is also traditionally taken to correspond to a single psychological state and to play a role in 'determining the extension' of a word. That is to say, what is known in the knowing of a word is thought to provide a criterion for all and only the items to which the word truly applies.

Criticizing this approach, Putnam (1975) famously provided reasons for doubting that anything 'in the head' could determine an extension. More recently, Millikan has emphasized and reinforced the criticism of the traditional view. She has argued not only that extensions are not determined by anything in the head (indeed, she criticizes Putnam for not being externalist enough in this respect; see 2005, pp. 126–129),[1] but also that a shared knowledge of a word—that is, shared grasp of a 'concept'—does not correlate with a single shared psychological state.

I will argue for a position that, at least in outline, corresponds to the traditional position, and yet which does not totally reject the externalist arguments that Millikan develops. A middle ground is possible when we make allowance for the fact that there is a specifically linguistic level of analysis that is not directly beholden to the type of extension that traditional accounts think can be explained, and externalist critiques think cannot be explained, in terms of shared concepts. I will suggest that we need to allow for two notions of determining an extension, one of which concerns broad metaphysical and conceptual questions that relate to the externalist arguments that Millikan puts forward, and the other of which relates to a narrower linguistic perspective that can be used to support a more traditional approach to word meaning.

[1] All references are to Millikan unless otherwise noted.

Millikan's discussion of externalism and concepts is fresh and often deeply critical of previous discussion. She has explicitly addressed the question of what is involved in knowing the meaning of a word, and, in part through a discussion of 'substance', she has made suggestions to help explain how children learn words so readily. I present and discuss the relevant aspects of her discussion, and I use this as a framework for establishing an account of word meaning that is contrary to the one Millikan supports.

In section 1 I discuss Millikan's approach to 'substance' and also the contrasting notion of 'class', and I argue that her externalist observations do not show that some notion of intension has no role to play. In section 2 I consider Millikan's specific approach to concepts, and I argue that, with respect to knowing a word, she is incorrect to deny a place for shared psychological states. In section 3 I support this position by suggesting that associating concepts with the ability to identify, which is central to Millikan's account, is not the appropriate perspective for delineating linguistic understanding. In section 4 I suggest that we should distinguish two notions of what it means to determine an extension, one of which directly relates to linguistic meaning, and the other of which is more relevant to externalist considerations.

1. Substances, classes, and a role for governing intensions[2]

For Millikan, substance is an ontological category, but she typically gives her discussion of substance an epistemological emphasis (1984, p. 275; 2000, p. 26):

> Substances are those things about which you can learn from one encounter something of what to expect on other encounters, where this is no accident but the result of a real connection. (2000, p. 15)

It is the 'real connection' that defines substance (Ryder et al. 2013, p. 131), though Millikan's emphasis on the epistemological consequence can sometimes obscure this. Real connections hold between the members of a substance, and because of these connections the members have likenesses to one another. Hence the consequence that you can learn from one encounter what to expect from other encounters.

There are different kinds of connection, and Millikan distinguishes three kinds of substance depending on the type of connection that is found between the members (2000, ch. 2;

---

[2] For the substance/class distinction, see in particular 2000, ch. 3; 2005, ch. 6. The distinction is Millikan's own and does not correspond to normal usage. In particular, her use of 'class term' and 'classifier' should not be confused with the use of 'classifier' in linguistics to indicate a type of adjective (or use thereof) in which it contrasts with 'qualitative'. Classifying uses of adjectives indicate membership of a group and are not graded. We say 'an annual event' but not 'a very annual event'. Qualitative uses of adjectives are gradable, as in 'a very exciting event'. An adjective can be used in both ways, as with 'a green melon' used to classify the melon (it is not a red melon), and 'a very green melon' where it is used to indicate something about the colour of the melon (cf. Kennedy and McNally, 2010). Millikan's notion of classifying terms corresponds more nearly to the qualitative use, while her notion of substance terms corresponds more to the classifier use.

2005, pp. 109–111). There are 'historical kinds' (such as dogs and chairs), 'eternal kinds' (stuffs such as gold and water), and 'individual kinds' (such as a particular person).

Historical kinds are substances whose members are connected to one another by causal/historical relations. These relations will typically ground important likenesses between the members. For example, members of a species will be alike either, on the basis of asexual reproduction, due to being clones of one another or, with sexual reproduction, due to stability (homeostasis) in the gene pool, this being caused by selection pressure on the characteristics that individuals of that species need to have in order to survive. Hence, blood-relatedness 'is bound to be accompanied by considerable overlap in properties' (2005, p. 107). As well as biological species, Millikan notes that artifacts also can constitute historical kinds. For example, chairs constitute a historical kind, chairs being related to one another through shared function, through the shared requirements of fitting human dimensions and preferences, and through usually being reproduced by copying rather than designed from scratch (2000, pp. 20f.). With eternal kinds, instances are alike not due to factors about reproduction but due to something like a shared inner structure, such as molecular structure, that 'results in a certain selection of properties' (2005, p. 110). This explains why one bit of gold is much like any other bit of gold. An individual substance, such as a person, retains many of its properties over time due to 'natural conservation laws' that tend to 'preserve its properties from day to day' (2005, p. 111).

By describing the connections as real, Millikan intends to convey that they exist independently of anything in our heads. Connections between weasels, or pieces of gold, or an individual over different times, occur because of causal factors that inhere in the world.

It is because there are reasons for uniformities between instances of a substance that we can learn, from one encounter with a substance, what to expect from other encounters with that substance. The reasons typically support a considerable degree of uniformity: from encountering one weasel, or one bit of gold, or Alex on one occasion, you can learn a huge amount about what to expect from other weasels, other bits of gold, or Alex on other occasions. By contrast, encounter with a red thing does not teach you much, if anything at all, about what to expect from encounters with other red things, beyond redness. This is because there are no real connections between red things. Red things do not constitute a substance.

Millikan describes the epistemological consequence of the real connections as one of 'inductive potential'. Studying one instance of a member of a substance 'is likely to yield a considerable amount of probable knowledge' about another instance (2005, p. 107f.). From studying one weasel, we will learn a great deal about what other weasels are like. Such substances display 'rich inductive potential' (2000, p. 17). By contrast, there is no rich inductive potential from studying a particular red thing: from studying a tomato we will not learn much about London buses and rubies.

Classes are defined in such a way that they do not display the rich inductive potential that Millikan typically associates with substances. 'Classes are defined by the members having certain common properties' (2005, p. 107).[3] That is to say, there is a property or set of properties that is correctly diagnostic of all and only the members of a group. There need be no reason why that collection of items constitutes a group besides a decision that those properties will be diagnostic

---

[3] Fuzzy classes are defined by members having overlapping properties or by having many properties in common with a paradigm (2005, p. 107).

in this way. The class of red things, for example, consists of all items that share the property of being red. While there will be reasons why instances of a substance may display some, and perhaps very many, likenesses, there is probably no reason why one red object 'should tend to be like the next [red object] in any respect other than redness' (2005, p. 108). If an item is described as red, we perhaps cannot, by virtue of that classification in itself, retrieve any more information about the item than is already contained in the classification (cf. 2000, p. 37). While there is a vast body of knowledge about weasels, there is no accumulated body of knowledge about red things.

On this basis, Millikan states, incautiously, that 'mere classes are not things one can learn anything at all about by induction' (2005, p. 108). This is misleading, for the way classes are defined does not rule out the possibility that properties used to classify can, at least sometimes, ground uniformities. This will occur where some generalization is evident in the conditions that accompany an instance's manifesting a given classifying property. Taking up an example from Cangelosi and Parisi (1998), Millikan herself notes that white things 'get dirty easily, … show up easily in dim light, stay cool in sunlight, but also tend to blind us, and so forth' (2000, p. 27). Even if the instances of white things have no connections to one another other than all being marked by the shared property of whiteness, this is consistent with there being non-accidental reasons for why some uniformities may, with at least some reliability, hold between members of the group.[4]

Conversely, while paradigm substance terms may display rich inductive potential, Millikan notes that there is 'a long, graded continuum', with substances varying widely in the reliability and multiplicity of the inferences they support (2000, p. 26). Encounter with one weasel provides considerable information about what the next weasel will be like. Encounter with one chair provides quite a lot of guidance about what the next chair may be like, but the range and reliability of inferences drawn from one chair to other chairs will be less than in the weasel case. Holidays[5] also may be viewed as members of a historical kind: holidays display uniformities due to common purpose (the purpose of having a clear break and rest); there is a common constraint, from human needs and preferences, that influences how best that purpose can be satisfied; and there is reproduction by copying (we are influenced by past experience and the experience of

---

[4] Confusingly, Millikan at one point responds to the way in which white things can support generalizations by indicating that 'white' can in fact act as a substance term, where it indicates white things (which *have* properties) rather than the property white (2000, p. 27). The suggestion is surprising. Her motivation for claiming a substance use of 'white' is the epistemological consequence rather than any real connection between white things. Millikan does not normally treat substances as including groupings that happen to display generalizations on account of a classifying property. Rather, substances are groupings based on 'genuine ontological grounds' (2000, p. 39), where these grounds are distinct from the classifying property itself. The ontological structures that makes something a substance are the historical/eternal/individual connections Millikan describes (see Ryder et al., 2013, p. 131).

[5] Millikan states that the category of substances may include 'certain event types', and she gives breakfast, dinner time, siesta time, and war, as examples. Other examples of substances include musical compositions, McDonalds, Californians, Western economies. See 2000, pp. 3, 26–7.

others in how we take holidays). Holidays support some degree of inductive potential, but less so than weasels given the diversity of the ways of taking a holiday.

While Millikan treats the substance/class distinction as one of kind, the categories are not mutually exclusive.[6] The definition of a class makes reference merely to the presence of a property or set of properties that is correctly diagnostic of all and only members of a group. In itself, this definition is independent of whether or not a group has some real connection over and above the presence of a correct diagnostic. Hence, Millikan states:

> Eternal kinds do form classes, all of whose members are alike in a variety of respects. But they are also much more than mere classes, because they are alike in these respects not by accident but in accordance with a causal explanation. (2005, p. 110)

Millikan is here explicitly allowing that a substance may also be a class. Gold is an eternal kind that inheres in the causal order, having a real connection based on a shared atomic structure; but it also turns out that the property of having atomic weight 79 is diagnostic of members of the group.

I suggest that we present Millikan's account in the following way. A class is a whole that is constituted by instances for which a given property or set of properties is diagnostic of all and only members of a group. A substance is a whole whose members are characterized by a connection that is independent of the presence, or absence, of a classifying diagnostic. The connection will display the historical, eternal, or individual factors that Millikan describes. This way of describing substance avoids some potentially misleading emphases that sometimes characterize Millikan's discussion, but we still have enough to capture the distinctions that I believe Millikan wants to preserve.

First, it indirectly accounts for the epistemic contrast between substances and classes, but without appearing to define substances and classes in terms of that potential.

Second, we capture an appropriate contrast between substance and class. They are not mutually exclusive, but it is not on that account inappropriate to say that a substance is different in kind from a class. A substance is never a mere class, even if a classifying condition can be formulated on the basis of what we happen to learn about a substance.

Third, the suggested definition does not place particular stress on substances whose identity can plausibly be viewed as largely determined by the causal order of the world. Millikan in practice tends to concentrate on types of example, such as weasels, for which inherence in the causal order of the world can be emphasized. This emphasis is not so relevant for groupings that do not naturally occur, such as holidays and chairs, but a generalization to such examples is required by Millikan's account (though this is sometimes obscured in her presentation). Millikan herself notes that many artifacts, for example, are 'governed in part by intensions determining their extensions to contain things that have been made for certain uses' (2010, p. 70). Chairs are historical kinds, but they do not constitute a naturally occurring group. Our purposes play a large role in determining them as a group, and this is a governance that, at least in part, is projected

---

[6] Some passages might appear to indicate otherwise. For example, she writes that words such as 'dog', 'chair', 'fruit', 'do not designate classes but units of another kind entirely' (2005, p. 106); the kinds 'that are substances in my sense are not *classes*' (Ryder et al., 2013, p. 133). These are to be read, though, as indicating a distinction in kind without commitment to mutual exclusivity.

from our heads. I have phrased Millikan's account of substance, therefore, in a way that avoids particular emphasis on naturally occurring groupings. The definition merely requires that there be connections (historical, eternal, individual) between the members of the substance. These connections might themselves be related to requirements or roles that cannot be specified purely in terms of the world's causal structures (cf. Prinz, 2013).

We can still, though, maintain an important place for the role that is played by the external world, even where a projected condition has to be taken into account. Suppose that it is correct to look to some partially governing 'intension' with respect to chairs. I here construe 'intension' in a relatively neutral and non-committal sense as indicating some kind of governing 'thought about' chairs that may play a partial role with respect to the determination of extensions, but where the precise nature of that role is left open. For illustrative purposes, I will construe this intension as the requirement that a chair be something that plays the role of providing a sitting possibility, with back support, for one person (this is intended as an approximation that is sufficiently accurate for illustrative purposes). While this is projected on the basis of our own requirements, it does not follow that how chairs are manifested in the sensible world is likewise projected from our heads. It will be a matter of exploration and discovery on our part to find out what materials and what designs are appropriate for manifesting the requirement that chairs should meet. This will in part depend on the nature of our bodies and the various kinds of sitting that are required for different purposes. These constraining factors are not themselves determined by the proposed intension. There is also a historical factor: chairs may be reproduced by copying, so how chairs are manifested will be locally influenced by tastes and traditions. There are, therefore, many external factors that influence the construction and design of chairs. These factors are not projected from our heads.

We can, therefore, distinguish external from internal factors when discussing words and the items those words are used of. This distinction should play a role in an account of linguistic meaning. As I will describe in the next section, Millikan tends to underplay the requirements for 'knowing the meaning' of a word because she not only focuses on the externally determined features of a category, but she also illustrates her account by choosing substance words (such as 'weasel') that lean heavily towards worldly based causal determinations. If we look, instead, to partially governing intensions, there is a place for a shared psychological state, and I will argue that linguistic understanding relates to this perspective rather than to the externally based factors that Millikan stresses.

2. Concepts and conceptions

There is, Millikan remarks, a traditional view that learning what a word means 'is coming to exercise the "same concept" in connection with the word that adults do' (2000, p. 92). Further, this shared concept is traditionally seen as corresponding to a single psychological state (2005, p. 134). The natural order of exposition here is from some prior intuitions or commitments about word meanings to a notion of 'concept' that will play the role indicated by those commitments. Millikan is, implicitly, deeply critical of this, and I interpret this as arising both from not wanting to be committed at the outset to a particular notion of linguistic meaning and from not wanting to be saddled with a traditional notion of concepts. Millikan reworks the notion of a concept into a

form that she is happy with,[7] and on that basis she provides a reassessment as to what an account of linguistic meaning should look like.

Millikan defines a concept as an ability to identify: 'One's concept of a thing might be described as one's ability to identify that thing …' (1984, p. 253). A given substance, such as a particular person or historical kind, may appear to us in many different guises and under many different perspectives and conditions. If we are to accumulate knowledge about the substance, we need to have the ability to reidentify the same substance 'through diverse media and under diverse conditions' (2000, p. 2). It is precisely the task of concept formation to enable us 'to recognize what is objectively the same through enormous diversity of appearances' (2005, p. 131).

This emphasis is very clearly shown in Millikan's application of this task analysis of concepts to more generalized abilities to identify constancies across variation, such as the ability to reidentify three-dimensional forms from different angles or the ability to identify colours under a variety of lighting conditions (1984, p. 255; 2005, p. 67). These abilities are concepts, just as our abilities to reidentify weasels or gold or Alex are concepts.

Abilities can be individuated in different ways. A child and a chemist both have the ability to identify sugar, and in this sense they share an ability. But the chemist has more sophisticated and reliable means for identifying sugar than the child, and in this sense they differ in abilities. Millikan introduces terminology to mark this distinction: the chemist and child have, in one sense, the same 'concept'; but in another sense they differ, and Millikan typically marks this by speaking of different 'conceptions' (2000, p. 11).[8] Conceptions are the particular means used, whereas concepts are the abilities individuated by whatever it is that is identified. Concepts are not individuated by reliability. Even if the child is unreliable compared with the chemist in identifying sugar, there is a sense in which they both have the same 'sugar' concept as long as they can both identify sugar.

Millikan speaks of what is required in order to understand a word in terms of grasping a concept, but with her own take on concepts she is also able to reject the claim that this corresponds to a shared psychological state. For two children to share the concept of their mother is for them both to have the ability to identify her. But their methods might vary—one might use sight, the other sound—and in this respect they have different conceptions of their mother. Millikan suggests that what have traditionally been called concepts are in many ways like her conceptions (2000, p. 12), and her emphasis on the method of identifying shows how one concept can go along with different psychological states.

---

[7] There is the suggestion that she will replace the term 'concept' with 'unicept' to help avoid unwelcome connotations (see Ryder et al., 2013, p. ix.).

[8] Millikan in fact expresses hesitation about the extent to which it is true that different people may have the *same* concept. She treats concepts as 'personal possessions, components of individual people's psychologies. We each have our own concepts' (Ryder et al., 2013, p. 256). But she also allows for a level of individuation at which it is appropriate to talk of sameness (ibid., pp. 131–132), and it is in this way that I will speak of concepts.

Is there such a thing as 'THE adult conception' of a substance? Given the numerous and diverse methods by which it is possible to learn to identify almost any substance, it seems that there could not possibly be. (2000, p. 93)

To know the meaning of a word is to have the appropriate concept. It follows that the meanings of at least many words 'are not determined by methods of application that are necessarily common to all competent users' (2010, p. 45). A word meaning need not correspond to a single psychological state across users.

In other words, by treating the topic of what it is to know a word in terms of ability to identify, and by allowing for different methods of identification, Millikan can argue that the methods people have of using a word need not be shared. While people may share the same 'concept'—this being individuated in terms of the substance identified—this is consistent with there being no shared psychological state that underpins this ability. A second conclusion is that the methods people use need not provide a correct diagnostic for all and only members of the substance. The extension of weasels is determined by the causal structure of the world. Our method of identifying provides a more or less reliable technique for deciding whether or not something is a weasel, but the property or set of properties we use for this purpose does not determine the extension.

The reasoning supports, therefore, the two central claims Millikan makes about an account of what is involved in 'knowing the meaning' of a word: knowing the meaning does not correspond to a particular psychological state nor to something that determines the extension of the word.[9] On the other hand, the examples that most clearly support these claims come from terms for naturally occurring groupings, such as 'weasel'. I will argue that the emphasis that arises from consideration of terms such as 'weasel' is misleading with respect to those substance words used for groupings that are accounted for, at least in part, in terms of what we project from our heads. And as the claims are based on a consideration of substances, we might expect that the two claims do not apply to words that are used for the diagnostic properties typical of classes.

Millikan, though, takes her account to be general across *all* substance terms. Further, by offering an account of how such terms are learnt (described below), she suggests that her account offers insight into the general problem of word acquisition, the problem of how 'it is possible for small children to learn, as Chomsky puts it … "a word an hour" between 18 months and 6 years of age' (2005, p. 119; 2010, p. 46).

Millikan's explanation for how substance terms are acquired is based on the claim that substance terms act very much like (proper) names. Names are easy to acquire, and acquiring them does not seem to correspond to any particular degree of knowledge about the item named.

---

[9] Millikan stresses discontinuity, rather than continuity, with Putnam's views (Putnam, 1975). Putnam appears to allow that the extension of a word can be determined by what is in the heads of experts. As Millikan notes (2005, p. 128), all he need be taken to be denying is that what is in the head of a typical individual determines the extension. Further, with respect to what is in the head with regard to a word, Putnam supports a view (a stereotype approach to concepts) that is consistent with the claim that there is a shared psychological state corresponding to knowledge of the word.

Individual persons are paradigm examples of individual substances (2000, pp. 23–4). If we think of the name of a particular person, there is a sense in which we learn to associate that name—that is, the name understood as naming that person—with whatever knowledge we may have happened to accumulate about that person. In talking of Plato, we are talking of an individual substance about which a body of knowledge has gathered. We carry this knowledge forward to subsequent discussions about Plato, and typically these discussions will be indicated by use of the term 'Plato'. The name itself does not denote any particular property or set of properties. It does, though, act as a locus of identification in the sense that, when that particular name is iterated, we tune in to the fact that it is the named individual substance that we are talking about. We will then draw on some facet or other of what we know of that substance in order to contribute to the discussion.

Words used for naming individual substances are very easily learnt. You can initiate a discussion about someone you know, say Markus, by the purely verbal means of simply using the name 'Markus'. Your hearers may know little or nothing about Markus, but they will be able to pick up on the name and use it to continue talking about Markus. One of the roles that Millikan gives to a concept of Z is that of being the condition for thinking of Z (e.g. 2005, p. 112). So, if being able to talk about Markus involves being able to think about Markus, you must have a concept of Markus. This concept provides both what is required in order to count as knowing the meaning of this name 'Markus' and what is required in order to think of Markus. Because concepts are construed in terms of an ability to identify, it follows that this easy ability we have to pick up on a name such as 'Markus' also gives us an ability to identify and reidentify Markus. Millikan allows for this by suggesting that a response to a substance term is itself a method of identifying. The idea is that when we recognize that a series of uses of 'Markus' are referring to one individual, in such a way that we realize that everything we learn about some Markus or other can be grouped together as pertaining to a single individual substance, we are thereby reidentifying this individual substance in each encounter with 'Markus'. We can, in other words, learn how to reidentify a substance purely through learning that an iterated term corresponds, in each use, to a particular unit around which we accumulate knowledge.

> In many cases there is not much more to having a substance concept than having a word. To have a word is to have a handle on tracking a substance via manifestations of it produced in a particular language community. (2000, p. 90)

In the next section I criticize the appeal that is made to reidentification in our response to a word, but the conclusion I take to be correct: it takes very little to think about something, maybe nothing more, in the limit, than a response to a naming word.

It is easy to introduce such words, and we can extend the point beyond names for individual substances. The tectum, pretectum, and pons are regions of the brain. Pseudomyxoma peritonei, trichopoliodystrophy, and neuromyelitis optica are names of diseases. By this simple introduction, I have handed you six concepts that you may not have had before (cf. 2000, p. 94, where Millikan hands us the concept of African dormice). In Millikan's terminology, you now have the ability to track these substances in discussion. This ability constitutes what it is to possess the respective concepts, irrespective of the fact that you may know very little about the named substances. In particular, there is no need to look to some adult standard, some relatively

expert knowledge of how to identify any of these substances, in order to justify the claim that you can use these words to talk about these substances. Because we are able to recognize a substance 'merely by learning a word for it', Millikan takes this as explanatory for 'why it is possible for small children to learn … "a word an hour" between 18 months and 6 years of age' (2005, p. 119).

In spite of these claims, two facts stand out. First, this account is most easily motivated by particular examples of substance words—exemplified here by words for diseases, species, and regions of the brain, but generalizing to words for natural groupings, discovered and delineated in scientific inquiry, about which typical language users may have very limited knowledge—that have characteristics that are not shared by the majority of substance words that we use in everyday speech. Second, while there is a sense in which this type of substance word is easy to learn, it is also a type of word that can lead to barriers to comprehension.

These two points can be shown by consideration of any typical text. If you run through practically any text and count the clear examples of words that name a particular species, or disease, or relatively abstruse part of the body, or some other distinction that reflects quite specialist knowledge and that points to a natural grouping discovered by science, such words will almost always be in a very small minority and often may not occur at all. Texts that contain a fair proportion of such words are noticeable precisely because they are unusual. They are also noticeable because they are, typically, forbidding and opaque to most readers, irrespective of how easily the words, taken individually, can be acquired. For example, consider the following passage:

> The neuroanatomical substrates for this organization are association cortices of different orders, both sensory and motor, some limbic structures (entorhinal cortex, hippocampus, amygdala, cingulate cortices), and the nestriatum/cerebellum, which constitute the substrate for convergence zones. (Taken, slightly altered, from Damasio, 1989)

This passage is forbidding largely because of the relatively high proportion of words—though even here they are in a minority—that induce a problem because readers have only an outline grasp of the substances spoken of. The ease we have in picking up such words, which Millikan emphasizes, does not offset this. Compare how, in ordinary conversation, if your interlocutor starts naming various people whom you do not know, the ease we may display in picking up such names does not offset the communicative block that use of the names may induce. In spite of the way in which names can indeed be easy to introduce, they also place a burden on memory and knowledge that can make them difficult to process.

We find relief in words that are not proper names and that do not act to label items about which we are largely ignorant. Yet by tending to use such words when illustrating her thesis, Millikan arguably introduces an emphasis that is not appropriate for the majority of substance words that we actually use in discourse. In other words, it may be misleading to describe what is involved in knowing the meaning of a word when the motivating examples tend to be words like 'weasel' and 'tectum'.

In particular, as I have already argued, for many words that can be taken to denote substances, such as 'holiday' and 'chair', there is a clear motivation for appealing to some sort of governing intension that, arguably, needs to be grasped before the words are properly understood.

By concentrating on natural groupings, about which many speakers of English are highly ignorant, Millikan underplays this. If we emphasize our general ignorance, we can see what motivates Millikan to put an emphasis on merely 'having a word', and being able to track its use, as sufficient for grasping a substance concept. This emphasis leads away from any clear motivation to posit a shared governing intension. A motivation is much more clearly seen when we shift the type of substance word that we consider.

In practice, Millikan herself does not treat substance words, even for natural groupings, as mere placeholders for a 'something we know almost nothing of' about which a body of knowledge accumulates. She also states that, as well as ability to identify, there is a second pole to a substance concept, namely, grasp of a 'substance template' (2000, p. 73). The intuition is that if you ask 'What shape is gold?' or 'What atomic number does a weasel have?' you show that you do not grasp the concepts of gold and weasels at all. Knowing a concept at least involves knowing what sort of questions it makes sense to ask.

It is not clear how substance templates fit into Millikan's overall account. They seem to answer to an important aspect of what it is to understand a word, in that for most, and perhaps all, words, an understanding of a word plausibly requires some understanding of what the word is used of. And yet this already takes us beyond the emphasis in which substance words are mere labels (and initially learnt as such) for some item about which our knowledge can grow. The appeal to substance templates suggests that Millikan recognizes a lacuna in her account. I take Millikan's mention of substance templates, along with the role she allows, with at least some words, for a partially governing 'intension' (2010, p. 70), as symptomatic of a requirement for a different type of account for expounding what it is to know the meaning of a word.

With respect to psychological states and extensions, Millikan makes the same claims for class words as she does for substance words: knowing the meaning of a class word does not correlate with a shared psychological state or with something that determines an extension. The basis for these claims, though, is somewhat different than the one used for substance words.

Millikan's claims about class words are, initially, surprising. Given how Millikan introduced the substance/class distinction, we might expect there to be straightforward disanalogy between what is involved in knowing the meaning of a class term and what is involved in knowing the meaning of a substance term. Sometimes Millikan herself appears to indicate this. She writes: 'Because substances are not classes … to have a concept of a certain substance is not to have a certain set of properties in mind' (2005, p. 112). This suggests, by the implied contrast, that to have a concept of a class is to have a certain set of properties in mind, namely, the properties in terms of which the class is defined. Further, Millikan expounds the claim that 'substance concepts are not just classifiers' as the claim that 'conceptions of substances are not what determine the extensions of their corresponding concepts …' (2000, p. 48). Again, the implied contrast suggests that when we are dealing with class terms, a conception that identifies some property (and we might expect the conception to correspond to a shared psychological state) does determine the extension of the class. Thus, we might expect Millikan to describe class concepts as corresponding to a shared psychological state that determines extensions.

But Millikan does not speak of class concepts in this way. Millikan supports a different conclusion about class terms by making use of the fine-grained ways in which methods of identifying can be individuated. Millikan allows conceptions to be individuated in terms of the precise method used to identify a particular property. I might identify weasels by their shape, but

I can do this by touch as well as by sight. This can count as two conceptions of weasels. To understand 'square' it is enough if a child can recognize squares, but there are different ways of doing this: the child can use sight, hands, feet, and so on (cf. 2010, p. 62). We can come to appreciate a distance (in this sense identify a distance) by ocular disparity, by tension in the focusing muscles of the eye, by partial occlusion of one object by another, by the known size of objects, by atmospheric haze, and so on (2005, p. 67).

On this basis, Millikan makes two claims that largely parallel what is claimed for substance terms: there is an absence of 'an agreed-upon decisive method for determining extension' even for paradigm class terms (see 2010, p. 50); and, given this, these terms 'do not have defining handed-down methods of application' (2010, p. 62). Millikan seems to have two points in mind: there are different methods (for identifying squares, for example), without any one method in particular being marked out by the community; and, for each method, our abilities in implementing it are fallible (touch, for example, would not in practice enable us to identify a very large square object). Hence, there is no *one* method that determines the extension, where 'determining an extension' is construed in terms of identifying all and only members of the class.

But from a different perspective both points can be denied without strain. It is plausible to maintain a distinction between, on the one hand, the differing skills and techniques that we may in practice use for identifying squares, and, on the other, a shared notion of what squareness consists in.[10] By focussing on the precise methods of identification used, Millikan ignores the role for something that is shared in our understanding of 'square'. And given that the property of squareness determines the class, there is a sense in which we can say that our concept of squareness determines the extension. This is to treat 'determining an extension' not in terms of the actual identifications we are in a position to make but in terms of a correlation between square items and our understanding of the word 'square'.[11]

I do not think, therefore, that Millikan has given compelling reasons to treat class concepts in the way she does. There are grounds for the claim that class concepts display both characteristics that Millikan is keen to deny: they correspond to shared psychological states, and they determine the extension of the class.

3. Understanding and identification

Millikan argues for a close relation between an account of what is required to understand a term and the ability to identify that which a term denotes. There are, though, reasons to deny that this close relation holds.

When we recognize that a series of uses of a substance term is referring to one kind of substance, such that everything we learn in that series of uses can be grouped together as pertaining to a single substance, Millikan says that we are thereby *reidentifying* this individual substance in each encounter with the substance term. For example, if we recognize that uses of

---

[10] I share an understanding of breaststroke with other swimmers, having studied the stroke using videos and information from the internet. My skill at implementing that knowledge, and the precise way I implement it, differs though very much from that of other swimmers.

[11] This introduces a distinction between 'conception' and the hands-on implementation of a method of identifying. This distinction is developed in section 3 of this paper.

'molybdenum' pertain to a single substance, we thereby reidentify molybdenum each time we encounter 'molybdenum'. Describing this in terms of an ability to reidentify is, though, somewhat strained. I can identify that an item (whether described or perceived) is a piece of molybdenum, and I can identify that a situation (whether described or perceived) is caused by molybdenum. It is not implausible to describe this as an ability to identify and reidentify molybdenum. But to identify and reidentify that you are talking about molybdenum on the basis of hearing the word 'molybdenum' is not, prima facie, to reidentify molybdenum. A more obvious explanation is that it is simply to understand the word 'molybdenum', and that this is distinct from the more demanding ability involved in recognizing that such-and-such is a piece of molybdenum or that such-and-such was influenced by the presence of molybdenum.

Millikan does argue for her point, on independent grounds, by suggesting that hearing a word for a substance is just as direct a way of *perceiving* the substance as when we react to the light reflected off an instance of the substance itself. Hearing a word, Millikan suggests, is just another form in which information about a thing arrives at our sensory surfaces (see 2000, ch. 6; 2004, ch. 9). This in turn depends on the claim that a word is a bearer of a type of natural information (see 2004, chs. 3, 6, 8). In Millikan's sense, an intentional sign, such as a monkey's alarm call, will bear natural information about eagles (say) when the use of the call correlates well enough with actual instances of eagles for it to be profitable for the community of monkeys to use and respond to the sign in the appropriate manner (evasive behaviour). The continued use of the call by the monkeys is explained by a semantic mapping (a mapping to the presence of an eagle) that accounts for the monkeys' response to the sign. Even if the call does not always correctly correlate with instances of eagles, it does so often enough to explain why the sign, and the corresponding response, is replicated in the monkey community. The background idea is that signs are effective when they lead to behaviour that is adapted to the 'existence of the signified affairs' (2004, p. 76)—here, behaviour adapted to the presence of an eagle.

Millikan extends this general picture to the words of natural language: linguistic items are 'the parade cases of intentional signs' (2004, p. 45), and words for substances survive when tokens of a word type are replicated on the basis of a semantic mapping that is often enough satisfied to explain why the sign is used as it is. Use of 'dog', for example, on this account, is maintained because of a semantic mapping, to instances of dogs, that is often enough satisfied for the continued use of 'dog' to be worthwhile. The word 'dog' is 'informative' in the sense that it correlates sufficiently with instances of dogs so as to lead to behaviour that is adapted to the presence of a dog.

This is a surprising claim. While Millikan's discussion of monkey calls, and the like, is convincing, she does not provide clear reasons for extending the approach to an account of individual words of a language. In the following passage, we do not so much get a reason for the claim as a sleight of hand in which two differing senses of 'information' appear to be at work:

> Consider, for example, a speaker whose purpose in using the word 'dog' is to communicate about or to call attention to facts that concern dogs. … Such a speaker will eventually stop trying to use the word 'dog' for this purpose if there is no evidence that it ever has this effect on hearers. Similarly, a hearer whose language-understanding faculties turn his mind to dogs with the purpose of collecting information about dogs

whenever speakers use the word 'dog' will soon unlearn this response if speakers never use the word 'dog' such that it carries information about dogs. (2004, p. 25)

On the face of it, this passage seems uncontroversial, appearing to say nothing more than that use of 'dog' as a word for dogs will die out if the community does not use it to speak about dogs. To 'convey information', in this sense, means nothing more than conveying 'facts that concern dogs'. But when, at the end of the passage, Millikan states that use of the word 'dog' should (with some appropriate frequency) be 'such that it carries information about dogs', she is in fact claiming that use of 'dog' will die out if it does not carry *natural* information about dogs. That is to say, 'dog' will be preserved in use only if its use correlates appropriately with states of affairs in which a dog is instantiated. This is to 'convey information' in the sense of providing an appropriately robust indication of the presence of a dog in the vicinity. The passage seems innocent, but it is in fact making a substantial and controversial claim.[12]

Prima facie, our continued use of 'dog' does not depend on the existence of an appropriate frequency of correlations between use of 'dog' and occurrences of instances of dogs, given that our response to the word need not involve an appropriate behavioural adaptation to the presence of a dog. For many terms, the continued use of which need not be in doubt, it may be the case that there are no such correlations at all or that what a correlation may amount to is unclear. For 'Plato', any correlations are restricted to the past; for 'Higgs boson' and 'dark energy', it is far from obvious what it might mean for a use of these terms to correlate with an instance; and for 'ether', a term still in use, there are no such correlations.

Yet it is a link to natural information that underpins Millikan's suggestion that understanding a substance word can be construed in terms of *identifying* instances of the substance. Consider how we may speak of an expert as identifying the presence of an eagle on the basis of hearing a particular bird sound or particular call given by a monkey. If, though, we remove the link to natural information, we thereby remove any obvious motivation for focusing on identification.

Positive reasons for distinguishing, in our understanding of a word, a factor that is not elucidated directly in terms of methods of identification can be provided from considerations already given. As noted, Millikan allows that we may share an intension that plays some sort of governing role with respect to the chair substance. For 'chair' I have suggested, as an illustrative approximation, something that provides a sitting possibility with back support for one person. But even if we do hypothesize this as a governing intension, it is not itself a method of identification nor does it determine those methods. While the intension is projected on the basis of our own requirements, it does not follow that how chairs are manifested in the sensible world is likewise projected from our heads; hence, it does not follow that the method we use to identify a chair follows directly from our grasp of the intension itself. It will be a matter of exploration and discovery on our part to find out what materials and what designs are appropriate, and locally used, for manifesting the requirement that chairs should meet. We become sensitive, at least for the chairs in our own culture, to the perceptible characteristics of chairs as determined by these

---

[12] This interpretation is not self-evident from the quotation itself, but it follows from Millikan's extensive discussion elsewhere. For the link between 'conveying information' and natural information, see 2004, ch. 3.

external factors. Where we identify the presence of such characteristics, we use them as a normally reliable diagnostic for indicating the presence of a chair. But these methods of identification, which may vary between people, are not themselves determined by the governing intension. Neither do these methods determine the extension.

We can, in other words, combine a claim about a shared understanding, hypothesized to correlate with a shared psychological state, with an observation about the multiplicity of methods of identifying.

The examples illustrate the consistency of positing a shared understanding of a word—where this understanding is adequate for providing an answer to what is required to 'know the meaning'—while allowing for more complex factors that play a role in 'determining the extension'. In effect, this is to posit a distinction between what is required to understand a word and what is required to identify how a word is being applied in a given instance. A shared intension (correlating with a shared psychological state) is consistent with different ways of identifying.[13]

4. Determining extensions

Finally, I would like to suggest that there is an important ambiguity in what might be meant by 'determining an extension' with respect to substance words.

On the one hand, there is the notion of a comprehensive decision as to which items in the world are such and such (weasels or chairs, for example). Such a (putative) decision may be informed by complex considerations, whether biological, in the case of weasels, or perhaps relating (inter alia) to whether or not 'real' chairs should always be artifacts.

We can distinguish this from a second notion of extension, which I phrase in terms of the range of applications that is 'linguistically licensed' by a word. Extension in this second sense is determined solely by the word's linguistic meaning, without taking into account the broader criteria that relate to the first sense of extension.

A similar distinction is implicit in a great deal of work on the context sensitivity of language. Bringing the distinction to attention helps create a space for a notion of linguistic meaning in which the meanings of an important spectrum of words, including substance words, do determine extensions.

We may reasonably doubt that there will always be a correct way, even in principle, of determining extensions in the first sense. Perhaps there is some relatively clear fact of the matter with respect to which mammals belong to the weasel substance, though even this will be susceptible to decisions as to what delineates a particular biological classification. With chairs, we may even more reasonably doubt that there is a fact of the matter, waiting there to be discovered, as to which items in the world precisely constitute the set of all chairs. For example, is there a fact of the matter as to whether a suitably arranged pile of suitcases, or a suitable tree stump, can truly count as being a chair? Most chairs, of course, are artifacts, constructed

---

[13] Represented as functions from worlds to extensions (Montague, 1970), intensions do not directly indicate methods of identifying. A connection may have been suggested by, first, the association of intensions with Fregean senses (Carnap, 1956) and, second, interpreting grasp of a Fregean sense as an ability to identify (cf. Dummett, 1978, p. 130).

deliberately for the chair purpose. It does not follow that there will be a principled basis on which to make a final decision as to whether an item, in order for it to really *be* a chair, has to be an artifact, and deliberately constructed for the chair purpose.[14]

The second notion of extension depends on the suggestion that what is involved in 'knowing the meaning' of a word provides a criterion for use that allows for a wider degree of application than might be warranted when we take into account the specific requirements and concerns that may be operative in a given context of use. To take just a single example from the many that theorists have discussed, Sperber and Wilson (1998) note that if Mary asks Peter to open the washing machine

> she will probably be asking him to open the lid of the machine. However, if Peter is a plumber, she might be asking him to unscrew the back … (p. 186)

The word 'open' fits, linguistically speaking, to either use, but one use rather than another may be appropriate in a given situation. Correspondingly, if the washing machine is described as 'open', the relevant state of the washing machine that justifies this description may vary depending on whether we are addressing an engineer or a laundry person. When addressing a laundry person, it may be deemed false, or at least inappropriate, to describe the machine as open if the back has been unscrewed but the drum is still shut. The word 'open' taken by itself does not answer to these further questions as to whether a particular use is appropriate.

Theorists differ over whether inappropriateness of use corresponds to falsity. That is to say, is it false, as well as inappropriate, to say 'The washing machine is open' when addressing someone wanting to do the laundry, if only the back of the machine is open? Borg (2004), for example, distinguishes what it is to grasp an element of a truth condition (for which read, in my terms, 'know a word meaning') from what is required to verify whether and how the condition is manifested in a particular instance. Borg would not, though, see inappropriate use as giving rise to falsity, given that she takes truth only to require consistency with the element of the 'truth condition' that a particular word conveys.

I abstract from the debate as to whether truth and falsity, or merely inappropriateness, is at issue. The substantive point remains: the range of applications that a word in principle allows for may be wider than the range of applications that are appropriate/true for a given use. It is this distinction that I wish to capture by the two notions of extension.

Examples can easily be found from our use of substance words. For example, consider 'poison' and 'milk'. There can be contexts in which these words would not be appropriately used (and might be said to be used falsely) of items that in other contexts could appropriately receive these predications. What is linguistically allowable, so to speak, may be broader than what is allowable given further considerations. Vitamin D can be used as rat poison, but it will often not

---

[14] Many paths are artifacts, but items do not have to be artifacts in order to truly be paths. In general, I see no reason to believe that artifact status is encoded at the linguistic level for individual words. Even given an assumption about artifactual status, there are complexities involved with making decisions on the sortal identity of a given artifact. See the editorial introduction (Carrara and Mingardo, 2013) to the special edition of *Review of Philosophy and Psychology* 2013, vol. 3(4); also Thomasson, 2007.

be appropriate, and may be deemed false, to describe vitamin D as a poison. There are many stuffs for which the term 'milk' is linguistically licensed but which, in some contexts, may be deemed as not in fact examples of milk at all (e.g. almond milk, rice milk, and so on).

To further illustrate this distinction, consider the suggestion, already used, that 'chair' is linguistically licensed to apply to items that offer a sitting possibility, with back support, for one. By hypothesis, there would be no linguistic error in using 'chair' of a pile of suitcases or tree stump, as long as they provide a sitting possibility, with back support, for one. This would not, though, rule out the possibility that further debate might arise as to whether these really *are* chairs. This debate might arise for practical, cultural, or legal reasons that are not themselves addressed by the linguistic specification. By hypothesis, and I think highly plausibly, the participants in the debate would share the same understanding of the word 'chair' but would make different judgements about further aspects of application.

If this suggestion is on the right lines, we can distinguish broad metaphysical and conceptual discussions about what things 'really are' weasels and chairs (the first sense of determining an extension, with the 'really are' corresponding to perhaps complex judgements relating to particular purposes) from the different question as to the possible range of usage of a word as determined solely on the basis of the word's meaning and what is involved in knowing that meaning (the second sense of extension, in which we look only to the basic linguistic specification and what that licenses). The first notion of extension is not linguistically determined.

In the second sense of extension, it is appropriate to describe what speakers know of substance words as determining the extension of the word.[15] This is an attenuated notion of extension, in that it may not correspond to judgements about truth and falsity (or appropriateness) as determined on other grounds.

Taken in this attenuated sense, it may even be the case that for terms such as 'weasel' we can isolate a shared understanding (corresponding to a shared psychological state) that, in a sense, determines an extension. As a suggested description of a shared understanding of 'weasel', this could be something as simple as: small mammal of a distinct species that is named 'weasel'. There are hints of a similar perspective in Millikan. She seems to indicate a required minimum when she says that a person may know as little as that malaria is a disease or that a weasel is a small mammal (1984, pp. 156–7; 2010, p. 70). With respect to what is linguistically licensed, the suggestion is that the extension of 'weasel' may be determined simply on the grounds of whatever usage the community happens to have alighted on. This will not in practice be arbitrary. Constraints are imposed by the shared understanding of 'mammal' and by the usage of related

---

[15] A referee raises the pertinent question: if we are allowing that an intension can determine an extension, does not this provide a diagnostic for all and only members, hence entail that chairs (etc.) are classes? The answer is yes, but two clarifying points need to be made. First, this observation arises only when we isolate a specifically linguistic level of analysis, in which we concentrate merely on the word 'chair' and its potential range of application. Actual use of a word is typically constrained by further factors. As a result, stating that 'chairs are a class' will be misleading, as this statement will typically be subject to further constraints on 'chair' that will undermine the reliability of the linguistically provided diagnostic. Second, as Millikan notes, being a class is consistent with also being a substance.

words such as 'stoat' and 'mink'. But this, clearly enough, does not provide a scientific criterion for isolating a particular species. By hypothesis, though, even if the word happens to be used by the community such that it is applied to examples that biologists come to decide do not count as weasels (or, conversely, if the community usage misses out on some examples that should count as weasels), this need not, by that fact, constitute a linguistic error; the 'error', as regarded by biologists, need not be regarded as a misunderstanding of the *word* by the community.

In summary, the suggestion is that the resources that language brings to bear at the individual word level for substance terms need not answer to projects aimed at delineating some privileged extension consisting only of items that 'really are' (or 'truly are') examples of such and such. What is known at the individual word level need not provide criteria able to give a final decision on these questions. We do not answer the broader questions by investigating what we know of the individual *word* used for those items. We can, though, legitimately talk of substance words as determining an extension when we restrict this to what is licensed purely on linguistic grounds, and we can relate this to a shared psychological state possessed by those who 'know the meaning' of a word.[16]

Department of Linguistics
University College London

References

Borg, E. 2004: *Minimal Semantics*. Oxford: Clarendon Press.

Cangelosi, A. and Parisi, D. 1998: Concepts in artificial organisms. *Behavioral and Brain Sciences*, 21(1), 68–9.

Carnap, R. 1956: *Meaning and Necessity* 2nd ed. Chicago: The University of Chicago Press.

Carrara, M. and Mingardo, D. 2013: Artifact categorization. Trends and problems. *Review of Philosophy and Psychology*, 4(3), 351–373.

Damasio, A. 1989: Time-locked multiregional retroactivation: a systems-level proposal for the neural substrates of recall and recognition. *Cognition*, 33(1), 25–62.

Dummett, M. 1978: *Truth and Other Enigmas*. London: Duckworth.

Kennedy, C. and McNally, L. 2010: Color, context, and compositionality. *Synthese*, 174(1), 79–98.

Millikan, R. 1984: *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.

Millikan, R. 2000: *On Clear and Confused Ideas*. Cambridge: Cambridge University Press.

Millikan, R. 2004: *Varieties of Meaning*. Cambridge, MA: MIT Press.

Millikan, R. 2005: *Language: A Biological Model*. Oxford: Oxford University Press.

Millikan, R. 2010: On knowing the meaning; with a coda on swampman. *Mind*, 119(473), 43–81.

---

[16] A referee notes that a more common account given by contextualists is that linguistic meaning constrains, but does not determine, extension. This, in fact, is consistent with my claim. The 'constraint' corresponds to that which underpins the linguistically licensed extension. In use, some constriction on what is deemed appropriate (or true) is often in play, and this corresponds to the first notion of extension in which linguistic meaning is not the final arbiter.

Montague, R. 1970: Universal grammar. In R. Montague *Formal Philosophy*. New Haven: Yale University Press, 1974, 222–246.

Prinz, J. 2013: Are Millikan's concepts inside out? In Ryder et al. 2013, 198–215.

Putnam, H. 1975: *Mind, Language and Reality*. Cambridge: Cambridge University Press.

Ryder, D., Kingsbury, J., and Williford, K. (eds) 2013: *Millikan and Her Critics*. Wiley-Blackwell: John Wiley & Sons.

Sperber, D. and Wilson, D. 1998: The mapping between the mental and the public lexicon. In P. Carruthers, and J. Boucher (eds), *Language and Thought: Interdisciplinary Themes*. Cambridge: Cambridge University Press, 184–200.

Thomasson, A. 2007: Artifacts and human concepts. In E. Margolis and S. Laurence (eds), *Creations of the Mind. Theories of Artifacts and their Representation.* Oxford: Oxford University Press, 52–73.