Getting Clear about Equivocal Concepts

Review of: Ruth Millikan, *On Clear and Confused Ideas*

For *Disputatio*

NICHOLAS SHEA

I. Introduction

Just how far can externalism go? In this exciting new book Ruth Millikan explores a radically externalist treatment of empirical concepts (Millikan 2000). For the last thirty years philosophy of mind's ties to meaning internalism have been loosened. The theory of content has swung uncomfortably on its moorings in a fickle current, straining against opposing ties to mind and world. In this book Millikan casts conceptual content adrift from the thinker: what determines the content of a concept is not cognitively accessible. She has only the stanchion of the world to hold her theory fast. She hopes that the tide will turn, and the theory of meaning will come stably to rest downstream of this anchor. This book is a bold exploration of how that might be achieved.

The book covers too much ground to be summarised in a review article. So I intend only to give a brief uncritical overview of the topics Millikan discusses. That is section II. Then in section III I will take up three of her topics, in order to make short observations on each. They are:-

(1) What it is for a thinker to know what he is thinking of.
(2) How it can be that misidentifying produces equivocal concepts.
(3) Soft natural information.

For discussion of some of the other issues raised in the book see Papineau and Shea (forthcoming).[1]

II. Overview

The first half of the book is a theory of one type of empirical concept, those of substances, in roughly the Aristotelian sense. A substance for Millikan is a category over which a variety of predicates are co-projectable in virtue of some single ontological ground. For example, all portions of gold have the same melting point, density at 0^0C and chemical properties (but not size or shape). This is because they all have the same atomic number. And all zebras have roughly the same size, shape and colour; and share those characteristics in virtue of lines of descent which trace back through similar environmental conditions to a single common ancestor.

[1] That review raises questions about two of Millikan's central themes which are not discussed in the present essay: (1) her rejection of Fregean senses; and (2) her argument that in thought there is no distinction to be drawn between equating concepts (mental identities like MARK TWAIN = SAMUEL CLEMENTS) and symbol-typing rules which treat symbols of the two concepts as symbols of the same symbol-type.

Humans are able to keep track of substances, as can many other animals. This is useful: things which share the common ontological ground of similarity will, as a result, stably exhibit some group of properties. By identifying a substance an animal can project those properties over instances of it: once they find that substance S has property P on one occasion, if they can reidentify S, then they can expect to encounter property P again on that occasion. For example, an animal which has a concept of one of its conspecifics can learn behavioural dispositions of that individual (that he is aggressive, say) and then expect the same disposition when he reidentifies that individual.

What then are concepts of these substances? For Millikan, they are not some abstract object, nor are they a mental word applied to the substance. In the first instance, Millikan does not equate concepts with particulars at all. Rather, concepts are abilities — the ability to identify the substance of which it is a concept.[2] Correlatively, the referent of a concept is whatever substance it is the purpose of that concept to identify. Of course, Millikan thinks that these purposes are teleological: what the ability was selected for. But she wants her theory to be compatible with other kinds of biological purposes. In short, a concept is an ability to identify, its referent is what it is its purpose to identify, and its purpose is that which accounts for past successes with the ability.

Deliberately missing from this picture is how we actually pick out substances – our descriptions of zebras or ways that we actually identify gold. Millikan calls these the 'conceptions' associated with a concept. Conceptions play no role in determining the extension of a concept. They are just how we happen, more or less fallibly, to try to identify the substance. The denial of the assumption that descriptions fix the reference of a concept is a familiar move. Millikan goes further and denies any meaning rationalism: the thinker does not in general (except in very exceptional cases) have cognitive grasp of that which determines the reference of her concepts. Rather, the reference is arrived at in virtue of facts about the thinker and her history. A consequence of Millikan's theory is that we can have concepts with very minimal associated conceptions. Provided we can identify the substance by some means for some practical purpose, we have the concept. Such practical purposes will require the thinker to project properties over instances of the concept.

For humans one such means is via language. Simply learning a word for a substance may provide a means of identifying it, via other language users. Even this minimal competence is useful provided the thinker knows what sort of properties ought to project over instances of the substance. For that, he needs a substance 'template': a specification of the type of properties which are projectable for that substance (eg, density but not shape for chemical elements). Since the function of concepts is to allow some range of properties to be projected over substances, there will be a substance template associated with every substance concept (including those arrived at independently of language). Learning a word together with its substance template is enough to identify the substance again and to project properties over its instances, and so gives you the concept. This is often very easy. Let me offer you the hypothetical concept ALANIUM. You

[2] Not all abilities to identify qualify. The thinker must also have practical uses for the ability and a 'substance template' (see 2 paragraphs below).

can probably guess that it is a chemical element, and a metal. So now you have a good idea what properties you can project from one encounter to another. And you can identify it again by recognising the word commonly used for it.

Millikan's theory of the acquisition of substance concepts through language provides an interesting explanation of why thinkers defer to experts. Learning a new word which falls under an existing substance template provides only a very minimal conception of that substance — only one very restricted means of identifying it. Good substance concepts are abilities to identify the substance in diverse ways in lots of circumstances. By learning a word a thinker has acquired a 'seed crystal' around which his ability with a new substance concept may grow. In order to improve his concept, the thinker has to learn about other ways of identifying the substance. He must learn these from experts, or from anyone else who knows other means of identifying the substance. Notice that, unlike Burge's account, this deference is not constitutive of the concept. Rather, Millikan's theory can explain why deference occurs.

The second half of the book argues for various theses about empirical concepts in general, which together constitute a radically externalist theory of concepts. Some of the material is taken from published papers, although here Millikan shows how the material fits together into a coherent theory. A prominent theme is her attack on neo-Fregean senses or modes of presentation. She argues that there are no common, individuable means of identifying a substance which are shared by all those who have a given concept, nor is there a stable intersubjective 'sense' to the concept in addition to its reference. In the process Millikan explains in what some common cognitive abilities consist: to reidentify a substance, to equate two concepts and to know what you are thinking. One of her central arguments concerns identity judgements, such as the statement that Hesperus is Phosphorus. Millikan hypothesises that such statements function so as to alter the thinker's cognitive apparatus — two concepts are merged. Their content and function cannot be captured by means of some proposition, as is commonly supposed. Millikan argues further that, for the brain, there is no difference in principle between treating two concepts as identical and re-typing the symbols for those concepts as being symbols of the same. Vehicles which are intrinsically quite different may be treated as symbols of the same substance. These issues are discussed in Papineau and Shea (forthcoming), which raises some questions about the viability of Millikan's approach.

The vigilant will have noticed that concepts have now been associated with symbol types. But concepts were supposed to be abilities of the whole organism, however realised. Where did symbols, presumably conceived of as vehicles of mental representation, slip in? One aim of this review is to show how these two uses of 'concept' are connected (see the end of subsection III.(1) below). The short answer is that Millikan employs a very liberal notion of 'representation' which allows that every ability to identify a substance will be mediated by an internal representation. That representation can be considered as a symbol for the referent of the concept. So 'concept' can be used to refer to the representation type, as well as the ability which it supports. These two uses of 'concept' pick out different ontological types. On the one hand it refers to an ability. On the other to a mental representation, which is part of the mechanism which realises that ability.

However, for most purposes no confusion should arise in moving between uses, since they individuate concepts with the same grain of analysis.

The rest of this paper is in three sections. In the first I outline Millikan's theory of what it is for a thinker to know what she is thinking of. Section (2) considers confused ideas: how it is that reasoning using concepts can lead to equivocation in what they refer to. Section (3) outlines Millikan's useful notion of 'soft' information, which she explains in an appendix to the book. That appendix is part of a larger project to elucidate the connections between Millikan's theory of concepts and informational theories of content, like that of Fred Dretske; a project which can only be touched on here.

III. Three Issues

(1) What it is for a thinker to know what he is thinking of

A traditional way of thinking about concepts holds that the conceptions we associate with a concept determine the reference of that concept. For example, for Frege sense determines reference. This gives a straightforward sense in which a thinker knows what he is thinking of: he grasps the sense and that determines the reference. But Millikan denies that conceptions play any role in fixing the reference of a concept. So she feels that she owes some account of how in her externalist theory a thinker nevertheless knows what he is thinking of. She provides a detailed theory of this phenomenon, which I survey in this section. The theory is pitched at three levels. The first is to argue that the requirement of 'knowing what you're thinking of' is very easily satisfied. On the only plausible understanding of that capacity, almost any user of concepts will satisfy it. The second part of the theory is to explain how, empirically, thinkers can come to develop that capacity.[3] The third stage is to see how as theorists of concepts we can reflect on our abilities so as to be reassured that our concepts are largely univocal, and not permeated with redundancy, equivocation and the confusion of the title of the book.

(i) Knowing what you're thinking of

Millikan denies that concept users do have or must have justification when applying a concept. Rather, they just exercise abilities to identify. Non-human animals can thus have concepts in exactly the same way as us. Millikan's account of 'knowing what you're thinking of' applies to all users of concepts, animal and human. In what, then, does this capacity consist?

Of course, we cannot hold a concept up against its referent (the substance itself) and compare them in thought, since substances only enter into thought through concepts. Nor is Millikan happy with the idea that we can satisfy this requirement by first using then mentioning the concept — the parallel in language would be ' "horse" means horse'. In fact, it is not clear what she takes the

[3] In the book these parts of the theory are presented in the opposite order. Chapter 13 explains in what the capacity to know what you are thinking of consists, and chapter 7 sets out how thinkers can develop such capacities.

requirement to be. It stems from something like the worries about externalism and self-knowledge which Gareth Evans addressed in *The Varieties of Reference* (1976). In any case, she feels she owes some kind of explanation in this area, and it is as follows. A thinker knows what he is thinking of when he uses a concept as a middle term in an amplificatory inference. For example when reasoning: yesterday's strawberries were tasty, here are some strawberries, so they are tasty. So to know what you are thinking of is not to possess propositional knowledge (of course), but rather to have the ability to put together two thoughts of a substance as being about the same, and then conclude something new.

How is that different from simply possessing the concept? After all, possessing a concept is having an ability to identify, whose function is to project properties over encounters. So the thinker must be able to reidentify the substance for the concept to fulfil its purpose. Does possessing a concept thus entail knowing what you're thinking of? The answer seems to be 'nearly, but not quite', since there may be ways of reidentifying a substance which fall short of using it as a middle term in mediate inference. Simultaneously identifying a substance through two sensory modalities requires reidentification of the substance (ie, co-identifying the outputs of the two perceptual systems), but arguably without any form of mediate inference. However, the difference is minimal. For example, Millikan argues that even such subpersonal activity as using binocular images to perceive depth requires a co-identification of the content of the two images (and hence knowing what you're thinking of).

In short, the theory appears to be as follows. Concepts are abilities to identify, and thus reidentify. A special case of reidentification is pairing two uses of a concept as a middle term in mediate inference. In that case the thinker 'knows what he is thinking of'. It is not clear just how minimal this requirement is. Is it a mediate inference to identify a substance and then apply a known property to that substance, eg, here is a cat, cats like fish, therefore he likes fish? If so, then any system which identifies substances by means of mental representations will, on occasion, satisfy the requirement of 'knowing what it is thinking of'. So perhaps the dialectic is as follows. Millikan feels that she must answer doubts that externalists can account for thinkers 'knowing what they're thinking of'. She then argues that the only sense which can be made of this requirement is that it requires the capacity to pair two concepts of a substance as a middle term in mediate inference. It seems to be a consequence of this theory that any cognitive system that identifies substances by means of mental representations will satisfy the requirement of knowing what it is thinking of. So be it – so much for the anti-externalist requirement.

Notice that Millikan's explanation of knowing what you are thinking of slips from thinking of concepts as abilities, however instantiated, to thinking of them as mental representations: symbols which are used in inferences.[4] This is not, however, an objection to the account. When an ability to identify is mediated by an internal representation it is legitimate to consider that representation as a symbol for the concept. That is, the conceptual vehicle can be considered as a mental symbol type, where mental symbols are typed by their content. And content derives from the purpose of the ability which the symbol helps to realise.

[4] This move was discussed at the end of section II above.

Say that on a particular occasion I identify a dog, and suppose that involves forming a mental representation. Then that token representation belongs to the symbol type DOG, deriving from the purpose of the ability. In a domain in which abilities to identify are supported by mental representations we can think of concepts either as abilities typed by purposes or as mental representations typed the same way. Put another way, where an ability to identify has a vehicle, a concept can be thought of as the vehicle typed by content.

(ii) Developing the capacity to know what you're thinking of

Now that the capacity to know what you're thinking of has been characterised, it can be asked how thinkers ever come to have that capacity. How do they come to be able to reidentify a substance and pair two tokens of a concept of it as a middle term in mediate inference? In chapter 7 Millikan gives an account of how concepts are so 'tuned'. That is by having diverse means of identifying a given substance in a variety of circumstances. Mistakes in identifying then show up in contradictions – properties which should apply to the substance will fail to apply to the misidentified instance. Adjusting to such contradictions allows animals to tune their concepts to a given substance, that is, to increase the range and variety of circumstances in which they can correctly identify it.

The idea is not that the thinker looks for contradictions or explicitly represents the law of non-contradiction. All that is required, rather, is a learning mechanism which is sensitive to the existence of contradictions and failed projections. The fact that misidentifications will throw up such contradictions in practice provides the basis on which conceptual development can be sensitive to such misidentifications. Similarly, a thinker starting out with a confused or redundant concept is not stuck with it. The conceptual confusion may show up in practice, and so can be corrected. This gives the thinker the material to become better at reidentifying substances. According to the characterisation in the last section, that is equivalent to becoming better at knowing what you're thinking of.

(iii) How we can know that our ideas are not confused

Finally, it is worth noting that Millikan thinks there is another level of sophistication above knowing what you're thinking of, which is the level of judgement, of true thought or cognition, which probably only humans have attained. Millikan characterises this with something like Evans' (1976) generality constraint. Humans are distinctive in that many of their mental representations are designed to take part in a wide range of mediate inferences. We gather information for its own sake, and then later use it for unenvisaged purposes. Thus, she argues, many human concepts must be embodied in such a form that they are available to take part in inferences with any other concept. There is not scope in this review to consider whether Millikan is right about this characterisation of what is distinctive about human cognition.

So far I have outlined Millikan's account of in what the capacity to know what you are thinking of consists, and her ideas about how thinkers come to improve those capacities. That is all at the level of unreflective mechanism. But her theory gives us the materials to say something about how the theorist can have

warrant for thinking her concepts are not confused; which I will call an 'epistemology of conceptual abilities'.

The potential worry with the story so far is that having an ability does not ensure you will use it correctly. You ride a bike, but sometimes you fall off. You can identify cows, but sometimes you get it wrong. Millikan has purpose, hence reference, determined historically, in virtue of the actual selectional history of the concept.[5] But that account only entails that you are able correctly to identify cows (say) in those actual historical circumstances. If those circumstances did not include picking out cows on dark nights, you may well misidentify cows on dark nights. There is no reason why your ability should extend much beyond the circumstances of your past successes with the concept. And those circumstances might be far from statistically normal. For example, although some northern hemisphere aquatic animals called paramecia can identify the direction of deadly oxygen-rich water (using internal magnetosomes that detect the earth's field), and swim away from it, they will systematically get it wrong if put in the Southern Ocean where the magnetic field is reversed. How can we tell that our concepts are not like that?

Millikan gives us some answers. As humans we can derive assurance from the fact that we have multiple ways of identifying each substance. Indeed, the utility of a substance concept depends proportionately on the number and variety of different ways we have of identifying its referent. Each time two different component abilities agree that the concept applies, we have an assurance that we are getting it right, that we are within the range of circumstances where the component abilities correctly identify the substance. Different means of identifying a substance could disagree on the same exemplar, in which case either or both must be outside its normal range of operation. That would warn us that something is wrong with our means of identifying. For example, if it looks like a cat but sounds like a dog, one of our ways of identifying must be going wrong. On the other hand, if it looks, sounds, feels and smells like a dog, we have additional assurance that we have correctly identified the substance.

It is not just through different sensory modalities that we can triangulate to the same substance. We may look at the same thing from two different perspectives, or through two different sets of mediating circumstances. For example, the temperature of a liquid can be measured by both a thermocouple and a mercury thermometer. The fact that the two agree provides support that a single property, *temperature,* is being measured. This is not undermined by the fact that theories are used to produce the readings given by each measuring device. The theory is part of the measuring system. Confirmation consists in the fact that the outputs of the two systems consistently agree.

It should be clear by now that Millikan is not claiming that in using our concepts we usually have grounds for thinking that we are correctly identifying their referents. Rather, we just go ahead and do it. There is no guarantee against unsuspected errors. Millikan's approach is to show that there are ways in which it

[5] Since most human concepts are learned, reference is determined by a combination of selectional and learning history: see the theory of derived and adapted proper function in Millikan (1984).

becomes apparent to us when we incorrectly identify. In particular, in such cases different means of identifying will often disagree. It is not that when we are getting it right, we have grounds for justifying our application of the concepts, rather that when we are getting it wrong we have ways of telling, and often do notice. Which makes it less likely that we will make the kind of stable errors in applying our concepts that a paramecium would make in the Southern Ocean. So Millikan is not giving us an epistemology of judgement: saying how we have justification for our beliefs. In general she is suspicious of that idea, and in her picture we often just apply our concepts without justification. However, she does give us as theorists of concepts some justification for our practices in using our conceptual abilities. Given that we have diverse means for identifying substances, we would often notice when we misidentify: different component abilities would disagree. Similarly, when a concept shows no variety in the observational circumstances in which it applies, we should beware that the concept may be empty. Reflecting on the fact that usually we have diverse means of identifying which agree on cases does then give us justification for the use of that conceptual ability. You could mark the distinction by calling it an epistemology of conceptual abilities, rather than an epistemology of judgement.

However, notice that this means of auditing our conceptual abilities does need some quite sophisticated machinery. We must be able to make judgements of the sort, if p then p, where p states something that is the case. For example, if the cat is fluffy [seen], then the cat is fluffy [felt]. So we have to be able to represent states of affairs in subject-predicate form. Furthermore, if we are to recognise contradictions we have to be able to tell when the contrary of a predicate applies: the object on my lap [felt] is a cat, but the object on my lap [seen] is not a cat (… error!). That requires having the operation of negation and the ability to identify the contraries of predicates. All of which are highly sophisticated conceptual capacities. So the abilities by which we as theorists are able reflective to derive some warrant for our conceptual abilities are probably not shared by many other animals.

(2) Producing Equivocal Concepts

The last section explained Millikan's account of how it can be that and how we check that a concept is not empty or systematically misapplied. Now I will consider another kind of failure – equivocation – when a thinker cannot distinguish between two substances and so has a concept which applies to both. There is a type of conceptual audit in which this shows up – we should notice that contradictions are systematically correlated with particular perspectives or means of identifying. The question for this section is how such equivocal concepts could ever arise.

Recall that our means of identifying a substance (our conceptions) do not fix the referent of a concept. That is the centrepiece of Millikan's denial of meaning rationalism. Rather, content is fixed by actual history. So the fact that a thinker would systematically misidentify a substance in some counterfactual circumstance does not alter the referent of the concept. Even if I would always judge horses on dark nights to be cows, that does not broaden the content of my concept to COW OR HORSE ON A DARK NIGHT.

However, Millikan insists that equivocal concepts do arise. They are the 'confused ideas' of the title to the book. Thus:

> 'If not soon corrected, mistaking the identity of an object of thought produces equivocation in thought, hence the beginning, at least, of change in the object of thought.'
> (Millikan 2000, ch. 13)

So it seems that actual misidentifications, rather than counterfactual ones, can alter the referent of a concept. How can that be? Surely if content is fixed by history, then even actual misidentifications should not be considered to impact upon content determination. The answer is that actual misidentifications become part of the history of that concept. Thus sufficient misidentifications can alter the history of the concept and thus what it is a concept of. Indeed, over time a concept of one substance could become equivocal and then shift to become a univocal concept of a second substance, even within the history of an individual thinker.

Millikan is on difficult ground here. It seems as if she is slipping towards verificationism. Not what could be called counterfactual verificationism, where dispositions to categorise determine content, but a kind of actualist verificationism in which the content of a concept is held to be all and only those things to which it has been actually applied. Perhaps there is theoretical room for Millikan to avoid this trap. If her account in Millikan (1984) of how learning fixes content can be sustained, then she would have the materials to formulate a distinction between misidentifications which are part of the development of a new concept, and those which are mere errors. However, none of that fine detail is found in the present book, which leaves a considerable worry about whether the theory ends up shading into verificationism about conceptual content.

Millikan also explains a second mechanism by which concepts can become equivocal, which is consistent with a strict historicism about content determination. This is the idea that equivocation arises when two concepts, which are not concepts of the same substance, are equated (treated as identical). For example, someone might wrongly conclude that their next door neighbour in Chelsea, Madge was in fact the pop star Madonna. According to Millikan's theory, in deciding that Madonna = Madge the thinker makes a functional change, merging their MADGE and MADONNA concepts and treating all the beliefs that have about either as true of a single individual. Millikan argues that the thinker ends up with an equivocal concept which is ambiguous between Madge and Madonna.

How that can be is at first puzzling. Externalists sometimes argue that the contents of concepts can be kept distinct, even if thinkers muddle them up. Each thought retains the content derived from its peculiar history. For example, consider Harriet who was brought up in Canada and used 'public school' to refer to State schools. She then moves to Britain where, after a while, she also learns and uses the phrase 'public school', which in Britain refers to private schools. Unreflectively she takes the words to mean the same in both countries. The externalist about content has the resources to explain how the contents of her two concepts can be kept separate. When Harriet uses 'public school' in a way which refers back to her Canadian memories it means State school. When her use

connects with British memories or activities it means private school. This particular example depends upon the role of public language in determining content, but Millikan wants to deny a parallel phenomenon can occur in the case of thoughts. If Harriet treats the two concepts as the same in thought then she confuses them and ends up with a single concept whose content is equivocal between private and State schools. How is that position consistent with Millikan's externalism?

The answer to this puzzle is, I think, straightforward if you can keep hold of the distinctions — so hold tight. Equivocation arises when two different concepts are mistakenly paired together as a middle term for mediate inference. In getting clear about this, it is important to recall the observation in the last section that, where mental representations are the vehicles of conceptual abilities, Millikan sometimes identifies concepts with types of those vehicles. At the start of an inference, two different identification abilities produce two different concepts (vehicles). Each concept has a different and unequivocal extension, deriving from the different purposes of the abilities which produced them. But then the two concepts are paired as a middle term in a mediate inference, producing a conclusion which also contains a concept (vehicle). What is the extension of that concept? As usual, the answer is given by looking at actual history. But we have a bifurcating history and hence two different purposes. So the vehicle token in the conclusion of the piece of reasoning has an ambiguous content: it applies to both substance A and substance B. It is equivocal. If a thinker systematically treats the outputs of two different conceptual abilities as identical, then those abilities effectively merge. The thinker no longer has two abilities, one to identify A and the other to identify B. Rather, she has a single ability to identify {A or B}.[6]

To return to our example, Harriet ends up with an equivocal concept of PUBLIC SCHOOL. A radical externalist wants to say that one use may have an American history and thus mean state school, another a British history and thus mean private school. By their separate histories the symbols have different contents, although the thinker does not distinguish between them. Millikan points out that there is no possibility of tracing each token use of this concept back to a separate history once the concepts have been paired as a middle term in mediate inference. In effect, each token has a history running back to both American and British uses. So Millikan disagrees with those externalists who argue that the concepts remain distinct. Notice however that equivocal concepts only arise in this way when treated the same by a single thinker. Identical treatment by two different thinkers does not, on Millikan's theory, lead to equivocation; so Putnam's original idea is preserved, namely that Oscar and Toscar on different planets can have water thoughts with different contents even though they are intrinsic duplicates and so treat their water concepts in the same way.

So Millikan reaches the right result. She can deny that two of a thinker's concepts could have different contents if he treated them as identical for all purposes. Indeed, that follows from fixing the contents of concepts by their purposes. Her externalism agrees with a common intuition which is otherwise would be used as a rebuttal by anti-externalists.

---

[6] Which is different from having the disjunctive content A OR B, which would require the separate representation of the components.

Millikan's treatment of equivocal concepts also fits nicely with her account of identity judgments. Recall that what it is to accept an identity claim like Mark Twain = Samuel Clements, in her view, is to treat mental tokens of the two concepts as indistinguishable in mental processing. That is just what is going on in pairing two concepts together as a middle term in mediate inference. Mental reasoning treats the two separate concept tokens as the same in order to reach its conclusion. So mistaken identity judgments produce equivocal concepts. For example, the practice of 17th century scientists of identifying the intensive quantity temperature with the extensive quantity heat gave them a confused concept HEAT/TEMPERATURE. Notice that on Millikan's theory identity judgments are not propositional, rather they are the means of achieving a functional change in the reasoning architecture – they are effectively the merging of two concepts (if concepts are read as abilities, the two distinguishable abilities become one, if concepts are read as vehicles, the two distinct mental symbol types become typed as identical). There is not scope here to criticise Millikan's controversial theory of identity judgments, but just to observe that it is consistent with her theory of equivocal concepts.

A similar story can be told about redundant concepts. Two distinct conceptual vehicle types will actually co-refer if they are produced by mechanisms whose purpose is to identify the very same substance. If those vehicles are not treated as of the same type in later processing, then there is a problem with the thinker's conceptual set up. He has redundant concepts. He fails to make inferences which he could otherwise make. Millikan thinks there is a kind of conceptual auditing in which redundancy will show up, in the sense of the previous section: typically the two concepts will accumulate the same properties in their associated conceptions, without also having contradictory properties. However, to remedy redundancy takes a functional change: the thinker must mark the two vehicle types as identical, so that they may be taken as such when needed and used as a middle term in inference. Merging redundant concepts cannot on any view be described propositionally. Millikan describes it functionally. This has a nice parallel with her story of what happens in the cases of identity judgements which are described propositionally (Mark Twain = Samuel Clements). These too are functional changes, mergings of files.

Millikan's ideas about redundant concepts are deeply anti-Fregean. Her view is that a thinker can rationally have two concepts with the same content which he fails to identify. Fregeans individuate content at the level of cognitive significance, making that move impossible. For a Fregean, if it is rationally possible to think different thoughts involving two concepts, then they must have different contents (different senses). Millikan embraces the anti-Fregean consequence of her theory and explains at length why she thinks there are no such things as Fregean senses. That is a highly ambitious project, but I will not question it here. Instead I will look briefly at Millikan's discussion of the relations between the contents of concepts and their vehicles. This is the topic which Millikan refers to as externalising / internalising sames / sameness etc. (chapter 9). The basic idea is that, as a thinker can have redundant concepts he cannot assume that concepts with different vehicles have different contents. Equivalently, sameness of content does not imply sameness of vehicle.

Care is needed here about how the vehicles are typed. Recall that where an ability to identify a substance is mediated by a mental representation, we type together all such representations (ie, by their content). That vehicle type is a symbol for the substance. Equivalently, mental representations typed that way are concepts. This is a fully externalist typing. Representations are not typed in virtue of any of their intrinsic properties, but rather in virtue of what substances they are representations of. So Millikan rightly warns against assuming that any intrinsic properties of mental symbols carry over to their contents. Symbols which are intrinsically indiscernible may have different contents if they are used to identify different substances, and symbols which are intrinsically discernible may have the same content. The previous paragraph made the point that even symbols which are treated differently by the thinker may not have different content. That is a different level of typing again: it is typing symbols on the basis of how they are treated in processing. On this typing, symbols which have intrinsic differences to which the thinker is insensitive are still typed together: they are treated as the same. It is a kind of internal extrinsic typing. The representations are typed on the basis of a relational property – how they would be treated in processing – but one that is defined only internally to the thinker. In short, there are three ways of typing mental representations in play here:

(1) In terms of their contents.
= fully externalistically
(2) In terms of the intrinsic properties of the representation.
= fully intrinsically
(3) In terms of how they are treated in mental processing
(all representations which are treated the same in processing form an equivalence class, which is a typing)
= internal-relationally

In general Millikan is right to point out that sameness / difference at any of the levels (1) – (3) does not carry over to any of the other levels. However there is an important exception to this principle, which Millikan does not explicitly draw out in the book. It arises from her theory of identity and equivocation which I set out above. If two concepts (symbols) of different substances are treated as identical by the thinker and paired together in a mediate inference, then the resulting concept will be equivocal between the two substances. So the two concepts thereby acquire the same (ambiguous) content. That is to say, concepts which are treated by the thinker as concepts of the same become concepts of the same. Of course, given enough confusion the thinker will cease to have a concept at all. But in the usual case this type of confusion will produce a concept with an equivocal content. Thus, as an exception to Millikan's general arguments against internalising / externalising moves, there is a property of conceptual vehicles which carries over to their contents. When vehicles are typed by how they are treated in internal processing, if they are typed as the same then they will have the same content. In terms of the categories above, sameness of (3) implies sameness of (1). (But not the converse: concepts with the same content may nevertheless be treated differently, ie, redundancy.) Equivalently, difference of content (typing (1) above) implies that the thinker must treat the representations differently in processing (typing (3)).

That is not yet to say that the thinker can *tell* when two of his concepts have the same content. It is unlikely to be fully cognitively transparent whether or not two symbols are treated as the same by all processing. But it is at least something which the thinker might be able to tell and might be able to make warranted judgements about. So it is unlike the Fregean test of difference of content in two respects. First, it does not make thinkers' rational judgements constitutive of difference in content. Second, it goes in the opposite direction – same internal treatment implies same content. The Fregean criterion is that different internal treatment implies different content.

Two final observations will help to show how this theory fits into wider debates about externalism. First, as observed above, the Putnam-style thought experiments are preserved. Identical thinkers in different contexts can have thoughts with different contents although they threat their concepts the same. It is only when concepts are treated as identical within a thinker, and paired as a middle term in a mediate inference, that equivocation and hence sameness of content arises. A second externalist thought is also preserved: two concepts in the same thinker may have different contents in virtue of different histories, even if they are intrinsically identical. It is only if they are thereby paired as middle term in mediate inference that they become equivocal. If they are kept separate, for example by arising in different cognitive systems, their separate contents are maintained.

(3) Soft Natural Information

Having built a theory of concepts on the existence of abilities to identify substances, in Appendix B Millikan says something about how we manage to pick up on which substances are which. A first answer is, 'by making use of information found in the environment about those substances'. Millikan improves this glib response by clarifying what kind of 'information' is needed to play this role.

Millikan's starting point is Dretske's notion of information (Dretske 1980). The basic idea is that a sign carries information in virtue of being correlated with what it signifies. Dretske deploys a strong definition requiring, given some background circumstances, the signal to vary with what it signs as a matter of law; so that the occurrence of the sign implies the existence of the signified.[7] In this sense for example, the coloured light emitted by street lamps is a sign of hot sodium gas. That is, given that emission spectrum the probability that sodium is present is 1, without it the probability is less than 1. Millikan observes that this is too strong a concept of information ("law-information") and that perceptual mechanisms actually make use of a weaker and commoner form of information which she calls recurrent soft natural information ("soft information"). The idea is that a sign carries soft natural information whenever it correlates with some environmental feature, but that correlation need only extend through some local domain, and may not be globally applicable. All correlations count, not just those which raise the probability of what is signified to 1. The only constraint is that the correlation should extend through the local domain for a single reason: there should be some univocal account of why a correlation found in one part of the local domain extends to other parts. For example, in reacting to a moving shadow and diving

[7] Dretske requires only a one-way dependency, from signs to what is signified.

underground a mouse is making use of soft natural information. The shadow carries the soft information that there is a predator overhead, although there is no law-like connection or conditional probability of 1 between the two. Millikan's central insight is that perceptual and cognitive mechanisms can make use of this more common soft information to identify substances, and not just Dretske's rarer law-information.

There are actually three reasons why something weaker than law-information is still useful in identifying substances. The first is because a particular sign-signified correlation may only hold within a limited geographical or temporal domain – mist on the North Welsh coast indicates imminent rain; identical mist on the desert coast of Namibia carries no such hope. So there is no correlation as a matter of law, but still a correlation which is useful in identifying within a particular domain. This may be partly catered for within Dretske's law-information by means of his 'channel conditions', the conditions which must be satisfied for there to be a strict correlation between sign and signified. Being in a particular geographical area might be one of the channel conditions. But there is a logical problem here: the sign can never carry both information about the signified and the information that its channel conditions are satisfied. So channel conditions have to be presupposed by the thinker at some stage, without any basis in law-information, on pain of regress. So the satisfaction of channel conditions is just like actually falling within some appropriate local domain: the thinker need not be able to tell when they fall within that domain.

Second, even within a particular domain the sign may only serve to increase the probability of what it signifies: the shadow on the ground may be a predator, and that is good enough for the mouse. And notice that information is only of any use if it does actually recur within a domain (which law-information need not, since it might be a one-off occurrence). Third, there can be no laws about individuals as such, since laws are necessarily general. So there can be no law-information about individuals. There can however exist grounded correlations concerning individuals, eg, a face with such-and-such features correlates strongly with the presence of Johnny. Millikan's soft information can therefore include information about individuals, where law-information cannot.

In short, natural selection is likely to have designed perceptual and cognitive mechanism to make use of any circumstances where a sign can be used. All this requires is that there is some grounded correlation between sign and signified. That is Millikan's notion of soft information: it is carried by signs whose types are correlated with something in the environment (the signified), there being a reason, grounded in natural necessity, why this correlation extends through a period of time of from one part of a locale to another (Millikan 2000, pp. 236-237). The correlation need only be strong enough to be useful to natural selection. That test will vary depending upon what the information is used for. When a mouse mistakes a shadow for a predator there is relatively little cost, compared to the cost of being eaten. So in some domains false positives may be very common, and the sign-signified correlation correspondingly weak. A further merit of Millikan's proposal is that it caters for this variability.

So how are thinkers able to identify substances? Answer: by making use of soft information in their environment. This suggests an interesting proposal about intentionality in general. There is lots of soft information, but only intentional systems are designed to produce signs carrying soft information. Conversely, it is plausible that any intentional representation must at least carry soft information about what it signifies. Might then intentional representation come to just this: carrying soft information as a matter of natural purpose? That would be an interesting way of uniting Millikan's recent theory of soft information with the traditional teleosemantics of (Millikan 1984).[8] Whether such unity can be achieved remains to be seen.[9]

Nicholas Shea
King's College, London
25th October 2001

Dretske, F. (1980), *Knowledge and the Flow of Information* (Cambridge, M.A., MIT Press).

Evans, G. (1976), *The Varieties of Reference*.

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. (forthcoming), book based on her Jean Nicod lectures to be delivered in May/June 2002, Cambridge University Press.

Papineau, D. and N. Shea (forthcoming), 'Review of *On Clear and Confused Ideas*', *Philosophy and Phenomenological Research*.

[8] Millikan (forthcoming) explores in detail the connections between the teleosemantics of Millikan (1984) and soft natural information, in the theory of what she calls 'recurrent local natural signs'.

[9] I am very grateful to Ruth Millikan, David Papineau and Richard Samuels for comments on an earlier draft of this paper.