

INTRODUCTION

DAN RYDER, JUSTINE KINGSBURY,
AND KENNETH WILLIFORD

Ruth Millikan's work is notable for its originality, scope, and coherence, and for its unwavering naturalism, with her focus on the proper functions of our cognitive and linguistic mechanisms as the unifying thread. Besides her meticulously worked-out and comprehensive theory of mind and language, she has made important contributions to philosophy of biology (especially to the discussion of biological functions), epistemology (especially on the nature of empirical knowledge), and metaphysics (especially the metaphysics of natural kinds).

Millikan's first book, *Language, Thought and Other Biological Categories* ("LTOBC"), appeared in 1984. It was followed by a steady stream of articles (some of them collected in *White Queen Psychology and Other Essays for Alice* ("WQ," 1993) and by three further books: *On Clear and Confused Ideas: An Essay About Substance Concepts* ("OCCI," 2000), *Varieties of Meaning* ("VOM," 2004), and *Language: A Biological Model* ("LBM," 2005). (A bibliography of Millikan's works can be found on p. xiii.) LTOBC was striking both for the novelty of its ideas and for its density. Many of the views defended in Millikan's copious later philosophical output, along with the arguments for those views, were already there in the 333 pages of LTOBC. It is unsurprising that many found the book hard to get to grips with. It earned her a reputation for being a *difficult* philosopher; a reputation that those who have read her later, more expansive work know to be undeserved, or at least greatly exaggerated.

This volume contains thirteen new essays on the work of Ruth Millikan, along with Millikan's replies to each. The philosophical range of the contributors is a testament to the breadth of Millikan's contribution to philosophy. Karen Neander in "Toward an Informational Teleosemantics," Peter Godfrey-Smith in "Signals, Icons, and Beliefs," Nicholas Shea in "Millikan's Isomorphism Requirement," and Michael Rescorla in

“Millikan on Honeybee Navigation and Communication” take up issues concerning Millikan’s teleological account of mental content. In “Concepts: Useful for Thinking,” Louise Antony takes exception to various features of Millikan’s account of concepts. Richard Fumerton’s “Properties over Substance” is a critical examination of Millikan’s claim (in OCCI) that substance concepts are ontologically and epistemologically prior to other concepts. Mohan Matthen in “Millikan’s Historical Kinds” and Crawford Elder in “Millikan, Realism, and Sameness” discuss Millikan’s account of natural kinds; in Elder’s case this is on the way to considering whether or not Millikan is a realist about objects. Charles Nussbaum writes about a Millikanian approach to naturalizing the law of noncontradiction (“Craning the Ultimate Skyhook: Millikan on the Law of Noncontradiction”).

Millikan has a radically externalist view about meaning, and a number of the essays are about matters arising from this. Jesse Prinz, in “Are Millikan’s Concepts Inside-Out?,” argues for a breaking down of the externalism/internalism divide. Cynthia and Graham Macdonald in “The Epistemology of Meaning” suggest a way in which teleosemantics can be made compatible with the existence of something like Fregean senses, while David Braddon-Mitchell in “Weasels and the *A Priori*” would like to bring Millikan and two-dimensionalism closer together. Finally, Willem deVries, in “All in the Family,” discusses the way in which both Millikan’s and Robert Brandom’s theories are, despite their striking differences, recognizably descendants of the views of their teacher, Wilfrid Sellars, and urges a rapprochement that would bring both closer to their roots.

Proper Functions

The function of a thing, in one sense, is what that thing is *for*, what it is supposed to do. One of the homes for such teleological concepts is in biology. For example, the function of the heart is to pump blood, and the function of the liver is to remove toxins from the bloodstream; arguably what makes something a heart or a liver is its function. Central to Millikan’s philosophy of mind and language is the idea that representations and the mechanisms that produce them have this kind of function as well; hence the provocative title *Language, Thought and Other Biological Categories*.

In the decade or so before the publication of LTOBC, Larry Wright (1976) defended a naturalistic account of teleological functions, and Robert Cummins argued in response that only artifacts have teleological functions. What sense, Cummins asked, can be made of the claim that an object is *supposed to do* something unless that object was made or designed by someone for that purpose? In opposition, Cummins defended a non-teleological, dispositional account of functions in biology, whereby a function is anything a mechanism does that contributes to the capacities of the larger system to which it belongs. A particular capacity is privileged only by our interests as investigators.

An important difference between Cummins-functions and teleological functions is that if something has a Cummins-function, it must actually perform it – the Cummins-function of a thing is always something that it actually succeeds in doing. The teleological function of something, in contrast, is something that it *should* do but not

necessarily something that it actually does. While biologists and physiologists are certainly interested in Cummins-functions, when they talk about *malfunction* (a sub-class of the cases in which something has a function which it fails to perform) and when they investigate evolutionary etiology in order to explain why some characteristic is present (Why do I have a heart? To circulate my blood), they seem to be talking about teleological functions. However, Cummins' challenge needs a response: What is it for something to have a teleological function, absent any creator or designer who made that thing for a purpose?

Millikan's account of teleology is not intended as an analysis of either an everyday or a biologist's notion of function. Rather, she introduces a theoretical definition of what she calls *proper functions*, a definition whose usefulness is to be judged by the explanatory success of the account of mind and language in which it plays a central role. Nevertheless, it meets Cummins' challenge. Millikan's (like Wright's) is an etiological account of function, according to which the proper function of my heart, for example, is to do whatever it is that the hearts of my ancestors did that contributed to the survival and reproduction of my ancestors and consequently the survival and proliferation of hearts of that type:

To put things very roughly, for an item *A* to have a function *F* as a "proper function," it is necessary (and close to sufficient) that one of these two conditions should hold. (1) *A* originated as a "reproduction" (to give one example, as a copy, or a copy of a copy) of some prior item or items that, *due* in part to possession of the properties reproduced, have actually performed *F* in the past, and *A* exists because (causally historically because) of this or these performances. (2) *A* originated as the product of some prior device that, given its circumstances, had performance of *F* as a proper function and that, under those circumstances, normally causes *F* to be performed by *means* of producing an item like *A*. (Millikan 1989b, 288–289)

Functions that fall under (1) above are *direct* proper functions; those that fall under (2) are *derived* proper functions. Proper functions can be *relational*; for example, a chameleon's pigment-changing mechanism has the (direct) proper function of bringing about a certain relationship between the skin of the chameleon and the surface the chameleon is on (in order to perform the further function of concealing it from predators). Continuing with the same example, a particular skin pattern has a proper function *derived* from the proper function of the pigment-changing mechanism that produces it: the proper function of enabling the chameleon to escape detection by predators. The pattern can have a derived proper function even if no chameleon has ever displayed precisely that pattern before.

A consequence of all this is that representations themselves can have proper functions, not just devices that produce them; *learned* representations can have proper functions, not just innate ones; and *unique* never-before-produced representations can have proper functions, not just representation types that are reproduced from earlier types. The caricature of Millikanian teleosemantics as applicable only to frogs reflexively snapping at flies is very far from the truth, something that all our contributors fully appreciate.

Representations: The Basic Teleosemantic Framework

On Millikan's view, different types of representations are characterized by their proper functions, derived from the proper functions of the mechanisms that produce them: the representation producers. Representations are also *used* by other mechanisms: their consumers. (According to Millikan, the proper functions of its consumers are the most important in determining a representation's content.) In his contribution, Peter Godfrey-Smith raises questions about the producer–consumer pairing ubiquitous in Millikan's treatment of representations, from states of bacterial magnetosomes to human linguistic utterances.

Representations fall into two broad types (not mutually exclusive): imperative and indicative. *Imperative* representations have the proper function of enabling a representation consumer to bring about a particular state of the world. For example, a representation in motor cortex may have the proper function of getting the rest of the motor system to bring about a certain arrangement of the limbs. More accurately, it is supposed to bring about a certain arrangement of the limbs according to a particular “mapping rule,” whereby variations in the state of the cortex are supposed to produce corresponding variations in the state of the limbs. It is helping to effect a mapping according to this rule that is selected for, and therefore characterizes a proper function of motor cortex. This means that the representational system is *productive*: it can represent (and have the function of producing) unique states of the limbs never before produced in its ancestors. In his contribution, Nicholas Shea addresses this kind of productivity.

In contrast with imperative representation, “... what makes an item an *indicative* [representation] has nothing to do with what it does; it has to do only with the *conditions under which* it Normally does whatever it does.” Roughly speaking, “Normal conditions” are those that must be mentioned in a minimal, general explanation of how the representation¹ has managed to perform its job in the past. That job is to enable its consumers to perform *their* proper functions. Among the normal conditions for proper performance of the consumer's function will be what the indicative representation represents. The thought is that indicative representations may be used for all sorts of things, but they will in general not be effective unless they stand in the right relationship to the world.

For example, beavers slap their tails in order to signal danger, and send the colony to safety. The representation producer is the slapping beaver (or its slap-producing mechanism), and the representation consumers are the hearing beavers (or their response mechanisms). The proper function of the slap is to enable the hearers to elude danger – this is the effect that resulted in the survival and proliferation of the slapping–fleeing combination, and it is thus the proper function of the consumer mechanism. A Normal condition for performing this function is that the slap coincide with the presence of danger, and it is this particular Normal condition that Millikan identifies as what the slap represents. It forms part of a Normal explanation for how the hearers have eluded danger in the past. (In her contribution, Karen Neander disputes this Normal conditions based analysis, arguing that sensory representations have the function of carrying information.)

¹ Or rather its ancestral incarnations.

A Normal explanation for how the hearers have eluded danger includes other Normal conditions, however. In order to flee, the hearing beavers need adequate oxygen, for example. Why does the slap only represent danger, among all of these Normal conditions? It is danger that the slap “maps” onto, not the presence of oxygen. The time and place of the slap vary, not with anything to do with oxygen, but with the time and place of danger. As in the motor cortex example, all representation involves a “mapping rule,” whereby transformations of the representation map onto variations in the environment represented. Bee dances are perhaps a more obvious example. Different bee dances are systematically related to each other, and the states of affairs onto which they map are similarly systematically related to each other. Transformations of the dance “correspond one-to-one to transformations of the location of nectar relative to hive and sun” (Millikan 1986, 74). Again,

... [w]hich mapping rule (which transformation correlation) is the relevant one to mention – which rule determines what the dance represents – is quite obvious. This rule is determined by the evolutionary history of the bee. It is that in accordance with which the dance must map onto the world in order to function properly in accordance with a Normal explanation, or, what is the same, in order that the mechanisms within watching bees that translate (physicist’s sense) the dance pattern into a direction of flight should perform all of their proper functions (including getting the bees to nectar) in accordance with a Normal explanation. (Millikan 1986, 78–79)

It is a Normal condition for proper performance of a watching bee’s nectar-finder that the dancing bee produce a dance that maps onto the current location of nectar in a certain way. That way is defined by the mapping rule, which is a correspondence between a system of possible representations and related group of possible states of affairs in the world. According to Millikan, this mapping rule plays a part in causally explaining the selection of the producer–consumer system. While the presence of oxygen may also be a Normal condition, it does not feature in any such mapping rule, therefore neither the beaver slap nor the bee dance represent anything about oxygen.

Many primitive representations are simultaneously indicative and imperative: Millikan calls these “pushmi-pullyu representations.” Both beaver tail slaps and bee dances are examples. The slap both says that there is danger, and tells other beavers to take cover; the dance both tells of the location of nectar, and commands the bees to go there. The story for more complex representational systems (such as human mental and linguistic representational systems) is more of the same, except that the indicative and imperative functions are more often separated. Desires and commands are imperative representations, so their function is to cause their consumers to bring about a certain state of affairs. That state of affairs is the content of the desire, or the meaning of the command. Beliefs and indicative utterances are indicative representations, so their proper functioning depends on their mapping onto the world in the right way – the way that historically has enabled their consumers to perform their proper functions. The basic teleosemantic framework reappears in many forms throughout Millikan’s work.

Informational teleosemanticists think that representation-producing systems have the function of producing states that carry natural information, that representations have the function of carrying information, and that the content of a particular representation is the

information it has the function of carrying. Millikan thinks otherwise: although true representations usually do carry natural information, that is not their function. The proper function of a representation is what it is designed to effect, to cause; whether it carries information is a matter of its origins and its current relations.

In “Toward an Informational Teleosemantics,” Karen Neander argues that at least for sensory indicative representations, informational teleosemantics can be made to work. Sensory systems, she suggests, have response functions: functions to respond to something by doing something else. Their function, according to Neander, is to produce representations that carry natural information.

Millikan, in her response, resists this move. The case she focuses on is that of a true representation that has not been produced in a Normal way. For example, someone may infer on the basis of mistaken evidence that Dan Dennett is in the next room. This belief is not a natural sign of that state of affairs, it does not carry natural information about it; nevertheless, it may happen to be true that Dan Dennett is in the next room. In cases like this, where an indicative representation is true by accident, the representation consumer doesn’t care. It performs its function equally well however the representation is produced. Therefore, argues Millikan, the representation-producing mechanism can’t have been selected for producing true representations *in a particular way*, but only for producing true representations, i.e., representations that map onto the world in the particular way required for interpretation by the consumer:

... if producing *true representations* is what *representation-producing systems* are selected for, informational teleosemantics can’t be quite right. The problem to be solved by teleosemantics is not just how false representations are possible, but also how representations can sometimes be representations and be true despite having been arrived at through accident or malfunction.

Peter Godfrey-Smith’s essay “Signals, Icons, and Beliefs” starts with the observation that Millikan’s account of content is a version of what he calls the sender–receiver approach: any entity that has semantic content has it as a consequence of its relations to a producer on one side and an interpreter or consumer on the other. Godfrey-Smith discusses sender–receiver systems in the light of recent work by Brian Skyrms and others on how such systems come into existence and are maintained. He then discusses the content of “messages” or “signals” in such systems, comparing Skyrms’ information-theoretic account with Millikan’s account. Skyrms’ account falls victim to the standard objection to simple informational accounts of content: it fails to allow for the possibility of misrepresentation. Millikan’s account does not suffer from this problem, because for Millikan, the content of a “signal” is (roughly) the state of affairs that must obtain for the receiver of that signal to use it successfully. But Godfrey-Smith points out a different problem for Millikan: that this focus on the receiver/consumer of a representation in some cases yields content-attributions that are “at odds with reasonable-looking intuitions about content.” The cases in question are ones in which the content of a representation on Millikan’s “consumer-determined” view is not the same as the state of affairs which triggers the production of the representation: R is a response to state of affairs X, but R

is useful only because X is reliably correlated with state of affairs Y because X and Y have a common cause. On Millikan's view, the content of R is Y, as Godfrey-Smith puts it, "no matter how remote Y is from the producer's ability to directly track the world." (Paul Pietroski [1992] gives a striking example in which a creature responds to dawn light with a representation meaning, according to Millikan's theory, "no predators this way.") Godfrey-Smith remarks that sometimes this fits with how we usually think of content, and sometimes it does not.

Godfrey-Smith explores how Millikan's account plays out in the case in which producers, representations, and interpreters are all internal to an organism. Raising the possibility that our mental mechanisms may not in fact have a straightforward sender-receiver structure, he considers some psychological research into internal maps (in rats): the evidence (along with a computational model of the rat's mapping system due to Reid and Staddon 1997) suggests that there is no internal map-reader distinct from the map. Godfrey-Smith suggests that this might be the situation with mental representations more generally: the distinction between producer, sign, and consumer might not be clear-cut.

Millikan accepts this last conclusion: there need be no "reader" of our mental representations, and producers and consumers are mechanisms or systems that "can overlap in their components and may operate in several capacities even at the same time." Her response to the query about Pietroski-type cases is that there is no general problem here: there are many intuitive cases where a sign and the state of affairs it is about have a common cause, for example beliefs about the future. The reason the Pietroski case jars with intuition is not this, but because as described, it involves an illicit attribution of sub-personal content to a whole creature (among other reasons).

Thought is productive: we can have an indefinitely large variety of thoughts, including novel ones with no history. Millikan accounts for the productivity of thought by appealing to isomorphisms (one to one mappings) between systems of representations and the systematically related states of affairs that they represent. In "Millikan's Isomorphism Requirement," Nicholas Shea argues that the relation of isomorphism is insufficiently demanding to play the role that Millikan assigns to it. He intends this as a friendly amendment, however, since there's something else that does the job.

Millikan has presented her theoretical picture as though a mere mathematical correspondence – an "isomorphism" is her term – between representations and affairs represented plays a causal/explanatory role. But, Shea points out, mere mathematical correspondence is cheap, and explains nothing. Rather, it is the *natural relations* among representations and the corresponding natural relations existing among states of affairs that do the explaining. An isomorphism does obtain, but comes along for free.

In her reply, Millikan generally accepts this point, saying that what she had in mind by "isomorphism" was indeed a correspondence between natural relations (call this a "natural isomorphism"). She is less concessive toward a related point that Shea makes, concerning the role of isomorphism in securing representational productivity. Shea points out that, depending on how the consumer system is structured, there is no necessity that relations among possible representations or relations among possible represented states of affairs be *natural* relations. Arbitrary mappings can work just as well, as long as the

representation consumer is built to take advantage of them. However, it is only when the relations are natural that the representational system will exhibit productivity, allowing for never-before-represented states of affairs to be represented for the first time in the history of the species.

Millikan agrees with the link between natural relations and productivity, but she argues that Shea's "arbitrary" mappings *do* in fact exhibit productivity. Her key move here is to show that Shea's examples display a natural isomorphism after all. It is an isomorphism that falls under what she calls "substitutional correspondence" (which Shea would consider arbitrary) rather than a more robust "projectable correspondence" (which Shea would call non-arbitrary). Shea, she argues, mistakenly sees productivity only when the correspondence is "projectable." There is a *natural* isomorphism even when the correspondence is merely substitutional.

A final dispute between Shea and Millikan resides in whether a producer–consumer system typically makes use of natural isomorphism as opposed to making use of simple correlational information. In the bee dance case, for instance, Shea argues that the system does not make use of the fact that similar bee dances indicate similar distances away; an arbitrary mapping, with no similarity structure, would do just as well for the job bees give it, which depends only on correlational information. Millikan disagrees (as with Neander); Shea, she says, is focusing on the producer's job, and not the consumer's. Correlational information does not explain how the consumer manages to do its job, because accidentally true representations (e.g., representations that are not causally related to their represented states of affairs, and so do not carry information about them) serve the consumer just as well as information-carrying true representations. "Truth" here picks out the mapping according to the natural isomorphism (which may be merely substitutional) that causally explains selection of the producer–consumer system.

Millikan's theory assigns truth conditions to (indicative) representations by singling out a particular mapping rule or natural isomorphism that helps explain selection of that particular producer–consumer system. It is important to Millikan's view that there is continuity between what she sees as simple representational systems, such as bacteria representing the direction of low oxygen, or bees representing the relative location of nectar, and more complex representational systems like our own.

Michael Rescorla challenges this picture in "Millikan on Honeybee Navigation and Communication." His central challenge comes from a contrast between the biological and psychological sciences. Rescorla points out that Millikan takes representations to be theoretical posits: she is not analyzing the use of the term "representation," whether by ordinary folk or by scientists. If representations are theoretical posits they must earn their explanatory keep. Rescorla argues that they do not do so for simple representational systems. In particular, there is no explanatory need to refer to truth conditions for bees and bacteria. By contrast, there is such a need in scientific psychology, and this fundamental difference is present for all to see within scientific practice. In effect, Rescorla argues that Millikan is likely to be mistaking a superficial similarity between us and the bees and bacteria (functional isomorphisms) for a deep similarity showing our internal states to be of the same kind, a similarity that she mistakenly identifies with the having of truth conditions. Instead, for a state to have truth conditions is probably something completely different.

Millikan replies that a substantial part of her corpus may be thought of as a detailed attempt to show that there really is no gulf here. (Indeed, she originally set out to account for truth conditions in *language*, and the application of the view she developed to simple systems was an unintended byproduct.) In addition, she criticizes Rescorla (and Tyler Burge [2010], who has raised similar questions) for setting invisible goalposts: they want to be shown how truth conditions play an explanatory role in bees and bacteria, *as they do in psychology*, but then they fail to say, non-circularly, what this supposed explanatory role in psychology consists in. As far as she can see, she has satisfied this requirement: on her account, the having of truth conditions (understood as Normal functional isomorphisms) explains how an individual's behavior coordinates with its environment, while accounting for a representation's intentional nature, that is the possibility that it can fail to map.

Concepts

Over time, you have learned lots of things about cats, sometimes in a way that allows you to give voice to this knowledge, and sometimes more implicitly. However, each time you have learned something new, whether by observing a cat or hearing someone talk about one, in order to incorporate this new knowledge correctly into your cognitive system, you must have made use of your ability to identify the proper subject matter: cats. Similarly, when you make use of your knowledge in your practical and theoretical dealings with cats, you must first recognize that it is *cats* that you are dealing with. This achievement is far from trivial, since cats can manifest themselves in myriad ways: as ginger or tabby, as fluffy or sleek, yowling or especially silent, and even through the medium of someone else's speech. In your empirical dealings with the world, this ability to recognize something – an object or a real kind or a property or a relation – as the *same* again is absolutely fundamental, and it is the focus of the attention Millikan pays to concepts.

On Millikan's view, "basic empirical concepts" just *are* abilities to reidentify. Your concept of cats is an ability to reidentify the biological cat kind, implemented by some mechanism that has that particular reidentification as its proper function. Your concept of *square* is an ability to reidentify square things from many distances, under many lighting conditions, by eye or by touch. In the context of judgment, the essence of the act of reidentifying is the performance of inferences where the concept acts like a middle term in a syllogism.

In turn, a person's ability to do A consists in the disposition to succeed in doing A if they try "... under the conditions that accounted for their ... past successes in doing A" (OCCI, 62). These conditions are "Normal conditions" for exercise of that ability, that is, conditions needed in order to explain how exercises of the mechanisms responsible for that disposition had succeeded in those past cases that are responsible for their having been maintained or selected for. Thus an empirical concept is a disposition to successfully reidentify something in Normal conditions. In her contribution, Louise Antony criticizes both Millikan's historical analysis of abilities, and her identification of concepts with abilities to reidentify rather than with mental words in the language of thought. She also addresses Millikan's generalizations concerning the mechanism by which reidentification is achieved, defending the identity judgment against Millikan's attacks.

A reidentification function depends on there being something objective to reidentify. The most fundamental entities that both humans and other animals capable of reidentifying, Millikan calls *substances*. (In his contribution, Richard Fumerton questions the epistemological priority Millikan assigns to substance concepts.) A substance is anything that retains its properties over space and/or time for some underlying reason; thus, things “about which it is possible to learn from one encounter something about what to expect on other encounters” (OCCI, 2). On this definition, a wide range of things can be substances, including individuals, kinds, stuffs, and event types, though the underlying reasons why these retain their properties across encounters with them can vary widely. Different samples of water have similar characteristics because of the stability of water’s chemical structure, commonalities across cats are explained more by common ancestry, whereas an individual object, like a child’s teddy bear, retains its properties because of basic physical conservation laws. In his contribution, Mohan Matthen offers an alternative to Millikan’s account of historical kinds like cats, whereas Crawford Elder takes Millikan to task for failing to maintain a proper realism about individuals.

There are various ways in which the ability to reidentify can be foiled; one of these is when a cognitive system confuses two substances or properties as being the same. The result is a confused or equivocal concept (hence the title *On Clear and Confused Ideas*). For example, someone’s reidentificatory mechanism might collect up information from both jadeite and nephrite, forming the equivocal concept of jade. There is no guarantee that one’s concepts are not confused in this way, but there is a cognitive mechanism whose function is to prevent it. When a concept is confused, like the jade concept, contradictions will tend to arise when the concept is applied in perception and inference. Therefore the cognitive mechanism that irons out contradictions serves to eliminate equivocal concepts. In this way, coherence acts as a test for unequivocal correspondence. In his contribution, Charles Nussbaum examines this key but underappreciated aspect of Millikan’s theory, with an eye on a Millikanian explanation for how the law of noncontradiction gets its grip on the world.

In “Concepts: Useful for Thinking,” Louise Antony takes issue with several features of Millikan’s picture of substance concepts. Antony thinks of concepts as vocabulary items in the language of thought (LOT), and she takes Millikan’s view of substance concepts as abilities to reidentify to be inconsistent with the LOT view. Antony argues that concepts have a wider range of functions than Millikan allows for. In particular, Millikan is mistaken when she says the fundamental role of concepts is to allow for reidentification and thus the collection of information; instead, concepts allow us to *think* about things. In Millikan’s failure to see this, and in her preference for an ability analysis of concept possession, Antony detects an underlying behavioristic bias. In reply, Millikan argues that the *good* parts of the LOT are fully consistent with an understanding of concepts as abilities, and Millikan (she says “Of course!”) wants to emphasize concepts’ important role in thinking; cf. their above-mentioned role in enabling inferences.

On the LOT view, syntactic forms serve as modes of presentation in a roughly Fregean explanation of opacity. In place of equivocal concepts, the LOT account posits the existence of false identity judgments, e.g., where the mental word JADEITE and the mental word NEPHRITE feature in the judgment JADEITE is NEPHRITE. There can also be

true identity judgments, of course. Millikan disparages this view in OCCI, putting in its place an account she also finds in Strawson (1974), whereby acts of identification are not to be assimilated to judgments. Instead, they involve a fundamental process of “sameness marking,” whereby there is some sort of marker whose function is to make the cognitive system treat a spread of information as pertaining to the same thing. A helpful image here is of two folders being merged: instead of an identity judgment, there is a change in the underlying representational language of the cognitive system. Antony denies that the LOT theory ultimately collapses into the Strawsonian view, as Millikan claims. On the contrary, the LOT view is perfectly coherent, and (she notes in passing) it offers the best explanation of the opacity phenomena as well. In response to Antony, Millikan argues that the identity of LOT syntactic vehicles (which is what accomplishes sameness marking) must ultimately be defined functionally. If that is correct, her original remarks concerning LOT still hold, since the supposed identity judgment would functionally merge what was previously two LOT terms. (She addresses Fregean treatments of opacity in her replies to Prinz and the Macdonalds, below.)

Antony and Millikan also participate in an exchange concerning Millikan’s historical analysis of abilities (Antony favors a dispositional account), and equivocal concepts (Antony thinks Millikan’s theory of confused concepts is itself confused). Millikan replies that Antony’s critiques here are based on a misunderstanding of her views, and suggests that Antony’s dispositional analysis is itself equivocal: it incorrectly sees true abilities where there is only accidental success.

Richard Fumerton’s essay “Properties over Substance” is a critical examination of Millikan’s claim (in OCCI) that substance concepts are ontologically and epistemologically prior to other concepts. He distinguishes four different ways in which the conceptual priority claim can be understood:

- 1 We typically acquire substance concepts before we acquire other concepts.
- 2 Other concepts can be analyzed, at least in part, into substance concepts (and not vice versa).
- 3 We can recognize substances much more easily than we can recognize things that belong to other categories.
- 4 We typically develop a vocabulary for labeling substances before we develop a vocabulary for labeling other things.

Fumerton notes that Millikan provides empirical evidence for (4), and points out that our use of common nouns might be a disguised way of talking about things having properties, and that in any case, it isn’t clear that (4) provides reason to believe any of the other three priority claims. Fumerton suggests that what we grasp first and recognize most easily might be complexes of properties, and concludes by wondering whether there might in fact be only a terminological difference between this and Millikan’s view.

Millikan responds by defending the view that “... one might recognize a substance owing to the fact that certain of its properties are manifesting themselves to one’s senses, but do so without employing concepts of those properties.” The defense is that the content of a concept is determined by the use to which it is put. What makes a concept a



MAMA concept is that the systems that use it won't function properly unless its tokening corresponds to the presence of Mama. Even though a baby may recognize Mama by her smell, her voice, etc., responding to those things serves no purpose *except* to make her respond appropriately to Mama. In this sense, our MAMA concept is prior to our concepts of the properties that enable us to recognize her.

On Millikan's view, biological (and other functional) kinds are historical kinds, and similarities between members of a biological kind (a species, for example) are to be explained by the historical relations they bear to one another. In "Millikan's Historical Kinds," Mohan Matthen defends the view that biological kinds are historical kinds against Michael Devitt's (2008) arguments for the revival of biological essentialism, arguing that although, as Devitt suggests, the similarities between members of a species may be able to be ahistorically explained in structural terms, polymorphism (the division of a species into stable dissimilar sub-kinds such as the sexes) cannot. Polymorphism can only be explained by certain historical relations, Matthen argues, and since these same historical relations are relevant also to similarities, it is clear that it is these historical relations (not the structural similarities) that define species membership. Millikan, on Matthen's view, overlooks the importance of polymorphism.

Matthen then argues that there is one respect in which biological kinds may be less historical than Millikan thinks. In cases in which there are two speciated descendant populations that turn out to be able to interbreed, he suggests, it is a mistake to rule out by definition the possibility that they belong to the same species – scientists do not treat the question as though it is settled by definition.

Millikan responds that the issue of how species should be delineated is not one that she wants or needs to take a stand on. However, as she says, she does use species as examples of historical kinds, and she does think that all of the actual groups which have been proposed in recent times as species have been historical kinds. The members of a species are alike because of certain historical relations that hold between them, and it is those relations, rather than the resulting similarities, which make them members of the same species. She suggests here that in cases of polymorphism, "each of the sexes [for example] within a sexual species forms a historical kind nested within the historical kind that is the species itself." In response to the case of the two speciated descendant populations which can interbreed, she suggests that the two populations are part of the same historical kind (as well as each separately constituting a different historical kind), since they do share a history which explains their similarities, both having split off from the same parent population: she declines to take a stand on the question of which of these historical kinds constitutes the species.

Crawford Elder, in "Millikan, Realism, and Sameness," takes up the subtitle of LTOBC, *New Foundations for Realism*, and argues that Millikan provides only half the foundation. He endorses Millikan's realist account of kinds as the only viable one, but contends that the apparatus she presents in chapters 16 and 17 meant to ground realism about individual objects and their persistence through time ultimately fails. Realism about objects requires that they have mind-independent origins and extinctions. According to Millikan's view, however, there can be large numbers of overlapping objects, including (for instance) "this adolescent," "this human being," and "this mass of human tissue,"



each of which has *different* temporal beginnings and endings due to their differing persistence conditions. As a consequence, Millikan says “a great deal of room is left for decision on our part as to how to divide the world into ... temporally extended wholes” (LTOBC, 293). This, Elder maintains, is a pretty stark statement of *anti*-realism about object persistence. He recommends that Millikan adopt a view whereby the membership conditions for natural kinds double as persistence conditions for their members, along with independently motivated tighter conditions on natural kindhood that will eliminate the problematic overlaps entailed by her current view.

The main thread of Millikan’s reply says that the overlaps are not, in the end, problematic for realism. On her view, the persistence of objects is to be treated exactly parallel to her treatment of the identity of a historical kind across its members: “The individual remains the same individual because it hangs on to its properties over time in accordance with natural conservation laws or other causal mechanisms.” Yes, there are many overlapping objects, but this does not endanger their mind-independence. Nor does it mean that we typically make some arbitrary intentional decision concerning what we’re talking or thinking about, as the quotation in the previous paragraph seems to suggest. Linguistic and conceptual functions are naturally selected so as to pick out certain individuals rather than others. Typically, this “picking out” will be rather vague; while there are natural divisions where many properties change at once (e.g., death), our selected linguistic and conceptual purposes are served even though the natural demarcation is not fully determinate. There is a healthy vagueness in our thought and talk about the many overlapping, but nevertheless real, individuals in the world.

Charles Nussbaum’s “Craning the Ultimate Skyhook: Millikan on the Law of Noncontradiction” explores some less charted territory for Millikan, in particular concerning what she ought to say about a naturalistic ground for the law of noncontradiction, as well as how to treat modal notions. He begins by asking where noncontradiction’s ontologically legislative authority comes from. Why does the world invariably obey this law? After showing that a number of theories (including modal realism and ersatzism) fail to answer the question, he explains what he takes to be Millikan’s answer. According to Millikan, properties come in ranges of contraries, i.e., there is a determinable with a range of determinates each of which is contrary to all the others. This contrariety is a matter of natural necessity: determinates under a determinable are incompatible by natural law. Negation is, fundamentally, what asserts this contrariety, whereas the principle of noncontradiction represents it more generally, as a hypothesis “about a certain kind of skeletal structure to be sought” (“Reply to Nussbaum,” this volume). This necessity is what explains the “grip” of logic on the world. It isn’t that logic constrains the world’s structure; instead it reflects it. Millikan identifies exactly what logic reflects: natural property incompatibility within a categorical structure (i.e., some ontology or other involving specific substances and the property ranges [determinables] that characterize them).

Up to this point, Millikan largely accepts Nussbaum’s description and extension of her account of noncontradiction. This is to get rid of the Myth of the Logical Given. But she is less inclined to accept Nussbaum’s proposal for a gradualist account of how human beings acquired their capacity to reason according to noncontradiction, now seen as representing a very general worldly structure. She doubts Nussbaum’s suggestion that this



capacity is somehow derived from similar sensitivities to contrary spaces in perceptual processing, because perception does not have subject–predicate structure. Rather, non-contradiction has the status of a high-level theoretical postulate, albeit one built in by evolution. (Nussbaum also invites Millikan to say a few words about the modalities, and she obliges, driving an interesting wedge between necessity, which she of course takes to be grounded in real-world structure, and possibility, which she does not. “P is possible” has a merely epistemic function: it serves only to move the mind in a certain direction, and has no truth condition.)

Externalism, Language, and Meaning Rationalism

On Millikan’s view, there are usually an indefinitely large number of ways of reidentifying something, depending on the circumstances, and different people (including different experts) may diverge drastically in the particular ways in which they reidentify that thing or property. In Millikan’s terminology, they may differ drastically in their “conceptions” of it. Consider how a child might identify a weasel compared to how a paleontologist might, and then consider how the ways of identifying a weasel that are available to Helen Keller will be different again. Nevertheless, the conceptions found in the child, the paleontologist, and Helen Keller (not to mention the changing conceptions in a single person through time) all share a function, a particular reidentification function. What they all share – a concept, or “unicept” (see Dennett’s foreword) – is externalistically individuated or relational. In his contribution, Jesse Prinz argues that Millikan’s externalism is overblown, and that internal psychological factors make a greater contribution to concept identity than Millikan acknowledges.

The diversity of reidentification methods is of fundamental importance to Millikan’s corpus, from her discussions of perception (perceptual constancy mechanisms) to language, and we now turn to the proper functions of the latter. The function of a language form is what it does that explains its proliferation and survival, of course. This is the language form’s “stabilizing function.” Like internal representations, some language forms perform their proper function involving cooperating producers (speakers) and consumers (hearers) by mapping onto the world according to a particular mapping rule – these are “semantic mapping rules” or simply “semantic mappings.” Semantic mappings determine truth conditions (for indicative forms) or satisfaction conditions (for imperative forms), but two linguistic forms might share a truth condition while differing in semantic mapping (e.g., “It’s raining” vs. “Rain is falling here now.”)

Stabilizing function, semantic mapping, and truth conditions are thus three different (though related) aspects of linguistic meaning.² Much of Millikan’s philosophy of language may be understood as an application of these aspects of meaning to various language forms. For example, the form “A is B” has a stabilizing function to get hearers to merge

² Unfortunately, in LTOBC Millikan used the term “sense” for semantic mapping, and “intension” for conception, a potential source of confusion. (Still more confusing, having introduced “unicepts,” she has recently been calling conceptions “input methods.”) In this volume, we avoid the older LTOBC terminology.



their concept files for A and B, as described earlier, but its truth condition is that “A” and “B” refer to the same thing. It is not part of the stabilizing function of “A is B” that the hearer form a belief about words, however. By contrast, that *is* part of the stabilizing function of “‘A’ has the same referent as ‘B’.” In addition, these two language forms differ in their semantic mapping rules. For example, “‘A’ has the same referent as ‘B’” has a place occupied by the term ‘referent’, and it mentions (rather than uses) the terms ‘A’ and ‘B’. “A is B” has neither of these features.

Millikan wields this three-pronged account of meaning to arrive at a comprehensive and nuanced treatment of a wide variety of language forms, including some hotly debated in the philosophy of language like definite descriptions, indexicals and demonstratives, propositional attitude reports, “exists,” and “means.” Each of these language forms and all of its subtypes, down to univocal terms, forms a lineage propagated by copying through years, centuries, or millennia of language use, forming a complex tree (or web) of descent similar to that found in the biological world. If a proper function of a language form is inherited via one of these lineages, that function is an aspect of the expression’s linguistic meaning. If not – if, for example, the proper function in question is derived from the speaker’s purpose alone – then that aspect is pragmatic. Contextual factors may play a role in either.

The proper function of an ordinary referential term will be to cause a representation of that term’s referent in the hearer’s mind (as part of a particular true belief, if the term is used in an indicative sentence). Details of the hearer’s conception of that referent will normally be irrelevant to successful coordination between speaker and hearer; rare exceptions are when a definition must be carefully passed down from user to user (perhaps “royal flush”). As a result, conceptions may exhibit little or no overlap among competent members of the same language community, nor do they need to in order for the relevant language forms to perform their proper functions. The only rules that a speaker must grasp are mapping rules, and this grasp may be minimal: the fix that a particular language user has on a term’s referent may amount to as little as simply knowing the term. This knowledge is the entering wedge that allows the language user to begin gathering information about the referent – filling the relevant file folder – upon hearing someone use the term in an indicative sentence.

It follows that in Millikan’s philosophy of language, there is nothing like Fregean sense. “The meanings that characterize the public part of a language,” she says, “are fully extensional” (LBM, 54). Both Jesse Prinz and Cynthia and Graham Macdonald take issue with Millikan’s eschewal of Fregean senses. As he does in the conceptual case, Prinz argues that psychological factors play a more important role than Millikan acknowledges, while the Macdonalds are more concerned with the role senses play in rationalizing behavior. They fear that Millikan’s view makes our knowledge of our own meanings so impoverished that it is incompatible with our status as rational agents.

David Braddon-Mitchell comes at the issue from two-dimensional semantics. He takes much of Millikan’s picture on board, but tries to fit it with conceptual analysis in the form of internally available analytic knowledge of meaning. Finally, Willem deVries argues that a speaker’s grasp of the rules of language is of fundamental importance in securing genuine intentionality, semantics, and reasoning, notions that have their home only within



Sellars' "space of reasons." Millikan, he says, neglects this aspect of Sellars' thought in her exclusive concentration on the causal order.

Jesse Prinz, in "Are Millikan's Concepts Inside-Out?," introduces the label "Outerism" to refer to a bundle of views which are often held together, as they are in Millikan's case: "Millikan is a realist, an externalist, a naturalist, an opponent of Fregeanism, and a crusader against Meaning Rationalism." Why "Outerism"? Because "... it shifts attention away from inner psychological resources and explains mental content by appeal to features of the mind-independent world." Prinz argues that the inner/outer distinction has been overblown, that the truth lies somewhere in between, and that Millikan's view is less Outerist than her rhetoric would lead one to believe. Instead, she is (or should be) an "inside-outerist": one who accepts that mental content is determined partially by what is inside the head, and partially by what is outside. In particular, it must be admitted that sometimes (especially when we run into unfamiliar cases) we are forced to make conscious decisions about what is going to fall under a particular concept we are using. For example, consider someone whose head is shaved except for a single hair standing out in bas-relief. Is he bald? Is someone with a split-brain one person, or two? These are ultimately matters for more-or-less arbitrary decision. As a result, while Prinz acknowledges that conceptions are variable among people (*contra* Frege), he argues that they do in fact play a role in fixing reference, even if only within an individual's idiolect.

Millikan staunchly defends a full-blown externalism in response to Prinz's proposed compromise, an externalism to which she was committed before Putnam's and Kripke's publications defending the orientation in the 1970s. (In fact, she complains that Putnam and Kripke's methodology of considering our intuitions about how the terms "meaning" and "reference" apply in possible cases is fundamentally internalist [LBM, ch. 7].) Millikan concentrates on two mistakes she takes Prinz to be making: first, he mixes up words and concepts. Yes, features of communicative contexts influence our word choices when we're trying to get something across, something she has never denied. In that minimal sense, we do make decisions about how to use words. However, we do *not* make "decisions" about how to apply our *concepts*. What would that even mean? We would already need to have the concepts to have the decision-thought.

The second and more fundamental mistake is that, even supposing that we do make decisions about how to use both words and concepts, this would not be relevant to determining their reference. The function of both words and concepts is to identify, not to classify. Identifying a natural clump or bump in nature (including the human parts of nature) does not require clear demarcation of that clump or bump from its surround, unlike classification, which does. To be anxious about how to apply a word or concept in unfamiliar cases is to make a mistake about what words and concepts are for.

Cynthia and Graham Macdonald take up one of the strands in Millikan's work identified by Prinz: her rejection of Fregean senses. Millikan's externalist view of linguistic meaning has the consequence that, in an important sense, language-users may not know the meanings of their own words. The Macdonalds suggest that in fact teleosemantics and the existence of something like Fregean senses can be reconciled, and so this counterintuitive consequence can be avoided.



The slimmed-down version of Fregean sense that the Macdonalds favor comes from John McDowell. The Macdonalds argue that McDowell's *de re* senses have none of the features that Millikan finds objectionable in other Fregean accounts of conceptual knowledge. In particular, *de re* senses are externalist and they do not include conceptions (which the Macdonalds agree vary among speakers and therefore cannot serve as senses). Instead, they are akin to informational channels that put a thinker in direct contact with objects. The benefit of such an account is that speakers count as knowing the meaning of their words in a way that makes their linguistic behavior subject to rational explanation, a desideratum that, according to the Macdonalds, Millikan's account puts in jeopardy.

Millikan disagrees. There is no need for any kind of Fregean senses to account for successful psychological explanations of linguistic behavior, since conceptions can do the job. Better or worse knowledge of what one is thinking of or speaking of corresponds to more or less adequate conceptions of the real humps and clumps found in nature, and so better or worse capacities to reidentify them, but there is no guarantee our concepts are adequate. (And to suppose that we know what we are thinking of when our concepts are confused is itself a confusion.) In addition, Millikan has non-Fregean solutions to the other problems traditionally addressed by appeal to senses, namely indirect discourse and opacity more generally, informative identity claims, and apparent reference to non-existents. Finally, Millikan clarifies her view of "direct" perception through the sense modalities and language, distancing herself from McDowell where the Macdonalds saw an overlap.

Millikan argues in her 2010 paper "On Knowing the Meaning: With a Coda on Swampman" that purely *a priori* analysis is not the right way to go about examining the meanings of empirical terms: empirical meaning is immutably embedded in the actual world. In "Weasels and the *A Priori*," David Braddon-Mitchell tries to reconcile this view with his own commitment to there being analytic truths about the content of concepts/terms.

According to two-dimensional semantics, associated with a term in a language there is a description that is analytic: e.g., the A-intension of "water" is something like "that actual substance which is drinkable and which forms the dominant fluid found in the lakes and rivers on earth." The speakers of Twin-English have a word "water" which has the same A-intension but a different referent.

Millikan points out that there is no shared introspectable A-intension, and Braddon-Mitchell concedes that she is right. There are lots of different ways of identifying water. Braddon-Mitchell's way of putting this point is that if two-dimensional semantics and neo-descriptivism are right (considering concepts rather than words, for the moment), we really do have massively varied concepts: in one dimension of content, the contents of our water thoughts vary dramatically.

Braddon-Mitchell backs off from his concession, however: he thinks that because we have dispositions to coordinate with others and defer to experts, our mental idiolects will not be that diverse after all. But, as Millikan points out, we are not disposed to coordinate our *concepts* in this sense (our identifying descriptions), or in Millikan's terminology our conceptions, with others – how would we even know what others' conceptions are? What we coordinate are referents.



The coordination story is more plausible for words. Braddon-Mitchell says: "... the meaning of words in languages supervenes on what concepts the words are typically used to communicate, and if most speakers use 'water' to communicate that something is the same substance as the potable stuff that fills the rivers, then that's what the word means." So long as the extension of all of our concepts of water overlaps, the word "water" does its job. However, "... perhaps the community average that we would defer to after equilibration is that the extension of the word 'water' is one which is given by the same-substance relation and a wide set of identifying descriptions."

Braddon-Mitchell accepts a large part of the Millikan view, and wants to integrate it with the two-dimensionalist view. There are real substances, such as water, and our words and concepts pick them out. "Water" refers to that substance no matter what idiosyncratic handle we may have on it, or what recognitional cues we have for detecting it. And, Braddon-Mitchell suggests, the two-dimensionalist could take Millikan's story about the "same substance" relation and use it to fill out his A-intensions: the story is something we're disposed to accept upon critical reflection (at least for many words and concepts), and so helps determine their extensions.

The disagreements that remain are about whether once you have accepted this much of the Millikan view there is any reason not to simply adopt it (rather than using it as an adjunct to your two-dimensional semantics), and, relatedly, about whether or not there is any commonality in content between agents in states which are narrowly the same but in different environments (Earth and Twin-Earth). The two-dimensionalist thinks the answer to the second question is "Obviously yes," and that this is a reason to prefer the two-dimensional view over Millikan's. Millikan replies that this would be to misunderstand concepts and the words that express them as classifiers rather than reidentifiers. Since they aren't classifiers, the dispositions alluded to by the two-dimensionalist are an idle fancy with no theoretical relevance. Perhaps they reveal potential trajectories in the development of semantic mapping functions in the language if certain changes were in fact to take place, either locally or in our broader understanding of the world. But current usage cannot determine this, and such potential trajectories have nothing to do with current linguistic meaning. (This is exactly why Millikan claims the case of Swampman is not helpful for understanding the nature of meaning, as explained in her 2010 paper.)

Braddon-Mitchell's concluding suggestion is that there are two different projects that need to be distinguished: Millikan's, which is "... explaining the existence and persistence of mental or linguistic phenomena" and a second project which is about "... the dispositions we have to respond to new information and new environments." He concedes that Millikan's project is well motivated and the other needs defending, not so much as regards its coherence, but as regards its explanatory utility. Millikan agrees. The first project is the one that she and everyone else engaged in the discussion is (or should be) interested in.

Millikan's 2005 paper "The Son and the Daughter: On Sellars, Brandom and Millikan" invites reflection on her philosophical ancestry and the family resemblances (or lack thereof) between her and her surprising philosophical sibling, Robert Brandom. Two of the authors in this collection take up this theme. Charles Nussbaum suggests that Hegel is Millikan's unacknowledged philosophical grandfather, at least in her view of kinds and



of the law of noncontradiction, whereas Willem deVries, in “All in the Family,” focuses on Sellars the father. Both have in mind the contrast between the causal order and the rational order, threads that are combined in Hegel and Sellars, but split apart in Millikan and Brandom. Millikan has pursued the project entirely within the causal order, whereas Brandom has focused on the rational order; Nussbaum and deVries are united in their concern that the son and the daughter are both neglecting half their inheritance.

DeVries makes this issue the centerpiece of his contribution. There are two ways of looking at the patterns to be found in linguistic behavior: as selected, where the pattern *does not* occur “because of the rules”; and as endorsed by the community, where the pattern *does* occur “because of the rules.” In her account of language, Millikan in particular has focused exclusively on the first, thereby neglecting the importance of a speaker’s *grasping* the rules of the language, rules or community norms that cannot be defined in causal-order terms. He claims that it is only through this grasp – “reflexivity” – that linguistic behavior can be truly rule-governed, allowing for the possibility of real inference and semantics (constituted by inferential role). There is no intentionality without reasoning, and no reasoning without acts governed by community norms.

Millikan replies that community norms do not play the central role that Sellars and deVries take them to. Much of language learning proceeds without correction, and linguistic mistakes are best seen as departures from one’s own purposes (which may include purposes to speak as others do), rather than from community norms. The role Millikan envisages for the language community is not as a source of norms, but rather as a medium through which we perceive the world. Since language is just another medium for reidentification, the thesis of the diversity of reidentification methods and its consequences applies. For most language, there are no stable “rules of language use” to be learned, other than rules for how words map onto the world: these are simple correspondence rules, not complex rules of inference. As a result, the normative structures that Sellars takes to be constitutive of the semantic (community norms of inference and language entry/exit) play only a supportive causal role for Millikan in the service of the fundamental conceptual function of reidentification. The simplicity of that function ensures that the semantic is *not* holistic, *contra* Sellars. Millikan also argues that going metalinguistic should not be seen as moving out of the causal order and into the “rational order”: she has grave doubts about the “autonomy of reason,” so dear to Sellars. Millikan is content to stay within the causal order in explaining semantic facts.

This volume will hopefully serve many functions. (If it goes through enough printings, those functions may even become direct proper functions, not merely derived from the editors’ intentions.) Some of those new to Millikan will hopefully find one or two of the contributions to be just the right introduction for them, given their background; we also believe that there is much here to enrich the knowledge of Millikan experts. Other readers, familiar only with Millikan’s widely known work on biological function and teleosemantics, will hopefully discover her insights in other areas, like epistemology and philosophy of language. We hope that still others will come to realize that in dismissing Millikan’s teleosemantics on the basis of intuitions about Swampman, they have inadvertently been begging the question against her. It is impossible to subject the full compass of Millikan’s work to scrutiny in the span of a single volume, but we hope the following

essays and replies will help bring new appreciation and understanding for the work of this remarkable philosopher.

Finally, the editors would like to express our profound appreciation to all of our contributing authors and the Wiley-Blackwell editorial team for the care and attention they brought to this project, and for their patience with us through the long process from inception to completion. Most of all, we owe a great debt to Ruth Millikan, not only for the enormous amount of work she put in (with an alacrity that put us to shame), but for her generosity to us and the many other younger scholars who, whatever their philosophical orientation, have benefited immeasurably from knowing her. (Many of our contributors know exactly what we're talking about.) Here's to you, Ruth!