

Millikan, meaning and minimalism

Article

Accepted Version

Borg, E. ORCID: <https://orcid.org/0000-0003-2725-9568>
(2018) Millikan, meaning and minimalism. *Theoria*, 84 (3). pp.
242-258. ISSN 1755-2567 doi:
<https://doi.org/10.1111/theo.12164> Available at
<https://centaur.reading.ac.uk/78612/>

It is advisable to refer to the publisher's version if you intend to cite from the work. See [Guidance on citing](#).

To link to this article DOI: <http://dx.doi.org/10.1111/theo.12164>

Publisher: Wiley-Blackwell

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the [End User Agreement](#).

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Millikan, Meaning, and Minimalism*

Abstract:

Across a series of seminal works, Ruth Millikan has produced a compelling and comprehensive naturalised account of content. With respect to linguistic meaning, her ground breaking approach has been to analyse the meaning of a linguistic term via the function it performs which has been responsible for securing the term's survival. This way of looking at things has significant repercussions for a number of recent debates in philosophy of language. This paper explores these repercussions through the lens of what is known as semantic minimalism, using the tenets of minimalism to draw out some questions for Millikan's approach to the semantics/pragmatics divide.

KEYWORDS: Millikan, Minimalism, representation, semantics, pragmatic enrichment.

Over a series of seminal works, Ruth Millikan has set out a comprehensive and compelling naturalised account of representation, which explains the representational content of linguistic signs via the same basic biological (teleological) mechanisms that she holds explain all forms of intentionality. The sheer breadth and scale of Millikan's account is extremely impressive and her way of looking at linguistic signs is appealing, however getting clear on exactly what her account predicts for core issues in the philosophy of language is not always easy. Thus in what follows I want to try to provide a route through some of these issues, using an alternative approach (that of so-called 'minimal semantics', Borg 2004, 2012) to help orient discussion. The suggestion is that approaching Millikan's work from this perspective will help to highlight how her approach relates to some much-discussed issues around the semantics/pragmatics boundary, while also bringing into focus a potential lacuna in the account. For, as we will see, approaching things in this way highlights the close explanatory circle between three key notions in Millikan's work: linguistic signs, linguistic functions, and domains, and raises the question of whether there is really sufficient independent individuation of these elements to allow the account to make a robust ruling on the kinds of cases that have exercised so many in recent philosophy of language.

The structure of the paper is as follows: §1 provides a brief overview of Millikan's approach to linguistic meaning. Then, in §2, I introduce minimalism via four core tenets. The rest of the paper then uses those four tenets to explore Millikan's account of linguistic meaning: §2.1 deals with meaning and intentions, §2.2 explores the nature and prevalence of context-sensitive expressions, §2.3 turns to recent issues at the semantics/pragmatics border, and §2.4 looks at propositions and compositionality. I conclude by suggesting that while there are several points of affinity between Millikan's view and minimalism, there is a substantive difference when it comes to

* This paper was written for a symposium in honour of Ruth Millikan, on the occasion of her receipt from the Royal Swedish Academy of Sciences of the Rolf Schock Prize for Logic and Philosophy 2017. I'm grateful to Peter Pagin and others in the Philosophy Department at Stockholm for organising such a great event, to those at the symposium for insightful comments and discussion, and, in particular, to Ruth (for help with this piece but also of course for her insightful and inspirational work).

questions about where to draw the semantics/pragmatics border and, related to this, what the bearers of semantic content might be.¹

1) Millikan's naturalistic view of linguistic meaning

Millikan aims to provide a thoroughly naturalised account of meaning in general.² Hence her account is designed to apply to all forms of representation, whether in thought or in language. So her starting point is that language use is to be explained by the same kind of teleological mechanisms that explain the rest of the biological life of the organism. This starting point immediately puts her at odds with other influential ways of approaching language (most obviously, it is in tension with a Chomskian approach). To see this, consider some of common philosophical assumptions about the fundamental nature of language:

- Language is governed by rules.
- Meaning is in the mind.
- A univocal term in a public language is associated with a psychological state common to all competent users.

Millikan, however, denies all of these common assumptions. Instead, according to her:

- Language is governed by probabilistic, conventional correlations.
- Contra Chomsky there are public languages and the function of public language is communication about worldly states (like perception, language provides a route to knowledge of the external world).³
- There is no guarantee of psychological homogeneity. External reference is what surfaces in meaning, not thoughts.

So, instead of starting an exploration of language by focusing on what we might have in our heads that enables language use (making philosophy of language a branch of individual psychology), Millikan starts from the perspective of public signs and the practices that have led to those signs' survival and proliferation. According to Millikan, language is governed not by strict deterministic laws but by conventional correlations between signs and what they signify. The central norms of language are like the norms of function and behaviour that account for the survival and proliferation of biological species. Thus what makes a linguistic sign meaningful is that it has a purpose that is (sometimes) realised by pairs of speakers and hearers, and that uses of that sign re-

¹ I suspect that at the heart of the worry I am going to raise is a deeper worry that the core explanatory notion that Millikan relies on – that of 'what something is selected for', or 'what function is responsible for something's survival' – doesn't extend properly from the biological realm to the linguistic. However, rather than levelling this charge directly, I'll approach the issue somewhat obliquely, through discussion of the relationship between signs, functions and domains.

² See, for instance, Millikan 2005.

³ Typically but not universally, for we can and do talk about psychological states.

occur due to the precedent set by past uses. For a sign to have a conventional purpose, then, is just for that sign typically to recur with that use on account of precedent (2004: 141).

The conventions in question are not (contra Dretske) strict regularities since language forms are often used by speakers for irregular purposes and hearers often fail to complete conventional patterns (they don't believe what they are told, don't do as they are directed). Specific linguistic forms survive and are reproduced together with cooperative hearer responses because, enough of the time, these patterns of production and response benefit speakers and hearers. To make this more concrete, we might ask what it is that makes the English language sign 'dog' about *dogs*? The answer is that there is a conventional practice of using this kind of sign to communicate things about dogs, and it is this practice that is responsible for the continued use of the sign. For, as Millikan notes (2004: 25):

Consider [for example] a speaker whose purpose in using the word 'dog' is to communicate about or call attention to facts that concern dogs...Such a speaker will eventually stop trying to use the word 'dog' for this purpose if there is no evidence that it ever has this effect on hearers. Similarly a hearer whose language-understanding faculties turn his mind to dogs with the purpose of collecting info about dogs whenever speakers use the word 'dog' will soon unlearn this response if speakers never use the word 'dog' such that it carries information about dogs.

What we need, then, to explain the intensional content of a linguistic sign is not some unexplained meaning entity or some kind of appeal to the thoughts that might accompany production of that sign, rather we need to look to the shared, community practices of using certain signs to communicate information about the world. On Millikan's naturalised model, we should replace the intensional schema ['X' means Y] with talk of the purpose (i.e. linguistic function, or what Millikan sometimes terms 'stabilizing cooperative function') of X/Y, where linguistic functions are functions that have been of interest to both speakers and hearers on enough occasions that a practice has been sustained in which speakers sometimes use the forms for that purpose and hearers sometimes respond to the forms in the required way. Importantly, we should note that functions are tied to domains, i.e. to parts of the world where the signs perform the function in question. This is easiest to see with respect to Millikan's discussion of the use of bird tracks as a sign of the presence of that kind of bird (what she terms a 'natural sign'). So a particular set of tracks might mean (i.e. have the purpose of indicating) pheasants relative to one forest, A, where the tracks are reliable indicators of the presence of pheasants. However, in another forest, B, where there are both pheasants and partridges making tracks which viewers happen to find indistinguishable, then, relative to domain B, the very same track will have the purpose of indicating [pheasant or partridge]. And as for bird tracks, so for linguistic signs. Thus one and the same type of sign can have different purposes, and hence mean different things, relative to different domains.

This realisation has important repercussions for what we take the bearers of semantic content to be, for, perhaps contrary to an unreflective assumption, it is clear that these bearers cannot be type words (where these are individuated by phonetic or typographical features). This has to be the case since two tokens with the same surface form can have different stabilizing functions (e.g. homonyms, or proper names, see 2004: 134), while two tokens with different surface forms can have the same stabilizing functions. An example of the latter kind surfaces with respect to referential uses of descriptions, as Millikan 2004: 153 notes:

Compare ‘that book over there with the blue cover’ with ‘the book over there with the blue cover’, placed in the same context. These two do not seem to differ in meaning or way of functioning.

So, instead of type level phonetic or typographic objects, instead it seems that stabilizing functions (and thus semantic content) will attach to phonetic/orthographic types relativized to a domain, for it is within a domain that a given expression type acquires a specific purpose. As she writes: “the central idea is that there is a historically positioned domain to which the sign is bound” (2004: 35).

So for example, with proper names, we might hold that:

- i. ‘Jeremy’ is a locally recurrent sign for Jeremy Bentham within the domain [discussions of Utilitarianism].
- ii. ‘Jeremy’ is a locally recurrent sign for Jeremy Corbyn within the domain [UK politics in 2018].

One and the same surface type expression (‘Jeremy’) but two different purposes relative to two different domains. Or again:

- i. ‘the dog’ is a conventional sign for dogs within the domain [English speakers].
- ii. ‘the dog’ is a conventional sign for Fido within the domain [my pets].⁴

Again, then, working out what purpose attaches to a surface-level sign (and hence working out exactly what sign we are dealing with) requires an appeal to a specific domain within which a specific purpose is realised.

Obviously there is much more that could be said about Millikan’s biosemantics in general and its application to linguistic signs in particular. However I hope this gives enough of an outline of her approach to allow us to begin to explore its repercussions. As noted at the outset, I want to approach Millikan’s view through the lens of an alternative approach to semantic theorising – that of so-called ‘semantic minimalism’ (argued for in Borg 2004, 2012, amongst other places). There are three reasons for doing this. First, I think that approaching things in this way may help to draw out some of the finer points of Millikan’s view of linguistic expressions, that might otherwise get

⁴ See Millikan 2004: 134-5, also Millikan: 2017 (p.173 in draft): “The ‘Hit me!’ said to request a card while playing blackjack is not the same construction as the ‘Hit me’ used to request an assault”.

swamped by issues about representation in general. Secondly, I take it that a robust, naturalised account of content would be a nice thing for a semanticist to have, so it is interesting to ask whether someone who signed up to a minimal approach to semantics could also embrace a Millikan-style teleological approach to linguistic meaning in general. Finally, however, I'm going to argue that approaching Millikan's account from the perspective of minimalism highlights a potential lacuna in the account, concerning the relationship between Millikan's core explanatory notions of sign, function, and domain. So, in the next section, I want to introduce four tenets that I've used elsewhere (Borg 2012) to characterise minimal semantics and then use these tenets to shape the exploration of Millikan's approach.

2) Millikan and Minimalism

In other works, I have defended an approach to semantic theorising I've called 'minimal semantics'. (Minimal, in part, because of the minimal job description it gives to semantics – specifically while a semantic theory should explain things like compositionality and systematicity, and certain limited judgements about expressed content, it should not attempt to explain the vast majority of our judgements about what a speaker communicates by her utterance.) Minimalism, as I construe it, is defined by four claims:

- (MS1) Semantic content for well-formed declarative sentences is propositional content.
- (MS2) Semantic content is fully realized by lexico-syntactic content; there is nothing got 'for free' at the semantic level.
- (MS3) There are only a limited number of lexicalized context-sensitive expressions in natural language.
- (MS4) Recovery of semantic content is possible without access to current speaker intentions (crudely, grasp of semantic content involves 'word reading' not 'mind reading').

Minimalism thus belongs with other formal approaches to semantics, where sentences, relativized to contexts (to accommodate the contribution of unarguably context-sensitive elements like tense markers, demonstratives and indexicals), are the bearers of semantic content (rather than token utterances) and where the contextual contribution to this semantic content is kept to a minimum (for discussion of exactly what 'a minimum' here amounts to, see Borg 2012). The natural enemies of minimalism, then, are a range of contemporary views (such as indexicalism, contextualism, relativism, and Travis-style occasion-sensitivity; for an elaboration of these views, see Borg 2012) which predict a much more pervasive role for context. According to these latter views the central task of semantic theorising is to capture our intuitions about communicated content (often couched in terms of our intuitions about 'what is said'), and this intuitively communicated content often (perhaps always) requires a rich appeal to the context of utterance. So, for instance, the

speaker who says "Jill got married and had children" intuitively (usually) communicates that [Jill got married and *then* had children], while the utterance "Naoki is tall" might communicate that [Naoki is tall *for a 12 year old*] in one context but some quite different proposition (e.g. [Naoki is tall *for a jockey*]) in another. To capture these judgements about communicated content we clearly need to appeal to a context of utterance (to Naoki's age, or gender, or occupation, etc., depending on the circumstances of the utterance) and thus non-minimalist views suggest that the contribution of context to content must be allowed to range much wider than things like reference determination for overt indexicals and demonstratives. If we are to properly capture semantic content contextualists maintain that we must allow that context contributes to semantic content in ways that go far beyond that allowed for in (MS1-4). In particular, most non-minimalist accounts suggest that we must countenance 'free pragmatic effects': modulations of content which occur due to features of the context of utterance but which are not demanded by anything in the lexico-syntactic form of the sentence produced.

So, where does Millikan's view stand with respect to this kind of debate? Does she side with the more minimalist picture or is her position better viewed as aligned with non-minimalist approaches? I want to suggest that the picture here is complicated as in fact her approach cross-cuts some of the core issues in play (and hence that it is instructive about the ways in which the various elements that are debated can hang together). However, on perhaps the most fundamental issue – the possibility of free pragmatic enrichment or modulation – Millikan's view *seems* to be non-minimal.⁵ However I'm going to suggest that (assuming that this is right) it opens up a potential worry with the account, for allowing free pragmatic effects makes pressing the question of how we identify the stabilising function of an expression and (related to this) how finely or otherwise we carve the domains within which expressions have their functions. For without a robust, independent answer to these questions there is a worry that the account will lack the resources required to give a decisive answer on some of the core examples highlighted in the minimalism/contextualism debate.

To see this, in the next section I want to work through each of (MS1-4), exploring how Millikan's approach stands with regard to it. I'll work through the commitments in reverse order, beginning with MS4, as that is where there is, I think, most optimism for consensus between Millikan and minimalism, and I'll end by looking at MS1, which seems to be a genuine anathema to Millikan's approach.

2.1) Meaning and intentions (MS4)

⁵ At least, that seems to be the view in her 2004, but it is much less clear that it is the view in her most recent book, *Beyond Concepts*. One question then might be whether the view is simply somewhat unclear in earlier work or whether Millikan has changed her mind across the two works.

For Millikan understanding language is a form of direct perception of the world: just as I can get the information that the dog is in the garden by looking out the window, so I can get this information by being told it. What language gives us is another informational route to the world and it is worldly reference that figures in specifications of meaning, not the thoughts occurring in another person's head. (MS4) chimes with this (anti-psychologistic) approach to semantic content, maintaining that semantic content is about world-invoking truth-conditions. This claim – that the mental states of the current speaker are not constitutive of semantic content – seems to be a strong point of consensus between Millikan-ism and minimalism, and a point where both the views perhaps diverge from more contextualist approaches (though it is hard to generalise about the relationship between non-minimalist approaches and speaker intentions; Travis, for instance, although he thinks a rich appeal to context, incorporating free pragmatic effects, is necessary, doesn't think this will go via speaker intentions).⁶

We should note, however, that a serious repercussion of adopting a general anti-psychological stance for semantic content is that it makes an account of demonstrative reference difficult: to work out what a token utterance of 'this' or 'that' refers to (and hence what contribution it makes to the truth-conditions of the utterance in which it occurs) it seems natural to think that we need to appeal to what the speaker intends. What makes the dog's expression, rather than the dog's collar or colour, the referent of my token utterance of 'That's pretty' is that I intended to refer to the dog's expression rather than its collar. Yet this obvious move seems to be ruled out by any account that treats speaker intentions as otiose when determining semantic content. Millikan's response to this worry (2004: 134) is, I think, to appeal to the domain on which a speaker is focusing (where obviously this notion of 'focus' must itself be non-intentional, cashed out instead in terms of things like eye-direction, topic of conversation, objects in the environment, etc). However opponents might wonder whether this approach can really meet the challenges posed by those, like Predelli 1998, who have argued so vociferously against the possibility that non-intentional criteria decide demonstrative reference. I won't pursue this point here, beyond noting that both Millikan-ism and minimalism will owe a satisfactory account of apparently intentional aspects of linguistic meaning, such as determining the reference of a demonstrative.

2.2) Context-sensitive expressions? (MS3)

⁶ See Travis 2008. The claim that current intentions are otiose in fixing semantic content also surfaces in recent arguments that we cannot fix so-called 'explicature' content – the level of somewhat pragmatically enhanced content that is commonly held to yield what a speaker says or asserts – by appeal to the mental states of speakers or hearers; see Borg 2016, 2017. Note that Millikan 2017 denies that there is a special level of 'what is said' content, making do with just two kinds of content – semantic content and implicature content.

Minimalism is committed to the idea that not all, or even almost all, expressions in a natural language are semantically context-sensitive (this claim serves to differentiate minimalism from what I've termed 'indexicalist' approaches, like that of Stanley 2002, which maintain that all or almost all expressions have hidden, syntactically marked contextual variables as part of their logical form). Millikan, I think, holds the contrary view that all terms are context-sensitive. She writes "There is no way of adding to a local sign so as to completely free it from context", 2004: 43. However the kind of context-sensitivity she has in mind is not more indexicality per se, but things like: identifying to which language a token belongs, what purpose a token is being put to in context, and making implicit reference to features of the domain (e.g. she holds that location is often implicitly represented by signs). This is just to emphasise, however, that it is conventional use within a domain that fixes content, and I think that this shifts the main point of conflict to MS2, rather than MS3.

Before turning to MS3 however, one additional point on the nature of unarguably context-sensitive terms is perhaps in order, for we should note that Millikan imposes a 'knowing which' condition on demonstratives (2004: 57), whereby speakers and hearers have to know which object they are talking about to use/understand expressions like 'that dog' (used referentially). Arguably, however (see, e.g., Borg 2002), not all referential uses of demonstrative terms (where we are individuating what it is to be a demonstrative by surface form) meet such a constraint. For uses of so-called 'deferred demonstratives' seem acceptable even where speaker and hearer are not in a position to identify the particular object of conversation in any substantive way. So, for instance, if I point at a book and say 'That author is my favourite' it may be that I have no way to identify the person picked out by the demonstrative expression except via the description 'the author of that'. Or again, if we are having a discussion about the fathers of children in a school year group it seems I might point at a child and say 'I've never met that dad', where there is no requirement that speaker or hearer have non-descriptive identifying knowledge of the father in question. Perhaps, then, when used in these kinds of ways Millikan would prefer to treat tokens of 'this F' or 'that G' as having a descriptive semantic content (since there seems to be no difference in function in these contexts between saying 'that author is my favourite' and saying 'the person who wrote that book is my favourite'). If this is right, then we would have another instance of a surface form type ('this'/'that') which embodies (at least) two different functions (a referential and a descriptive function) and thus which is ambiguous between two distinct semantic types.

2.3) The semantics/pragmatics border (MS2)

Millikan acknowledges that not all uses and not all meanings qualify as part of the semantic content of a sign. As she writes 2004: 26-7:

The functions of conventional language devices considered as such are memetic purposes. But when language parts are used in figures of speech or used as bases for Gricean implicatures, the underlying memetic purposes of these expressions are crossed by the speaker's purposes. Then what the speaker means may not be what the words mean, or it may be more than what the words mean. The very same expression token then has two purposes derived from two different sources that cross, a literal meaning deriving from its function in the public language and a pragmatic meaning deriving from the speaker's purposes. Public language meaning and speaker meaning often diverge in this way.

The study of semantics, then, is the study of what is conventional in language use (public language meaning), while pragmatics involves the study of what is conveyed in ways not yet hardened into convention (speaker meaning). This makes the semantics/pragmatics divide one of degree:

Once you grasp that for a usage to be conventional is just for it typically to recur on account of precedent, the debate about what is 'said' (conventionally signified) versus only pragmatically 'implicated' takes on a clearer meaning. But it also becomes clear why the line is vague. Transition from nonconventional use to conventional is a gradual, largely statistical matter (2004: 141)...Figures of speech, implicatures, and other forms of usage [that] are slowly moving from being entirely innovative, through being somewhat familiar, to being handled automatically without parsing or derivation of meaning from compositional structure (2004: 145).⁷

Yet how, we might wonder, does this model of the semantics/pragmatics divide (where what is semantic concerns regular function, relative to a domain, while what is pragmatic concerns what is innovative and original) play out with respect to the particular kinds of cases that have so exercised recent writers in this area? There seem to be two worries with the proposal. First, we might worry that characterising semantic content purely in terms of 'content regularly conveyed on account of precedent' moves wholesale a range of cases often treated as pragmatic into the semantic realm (where this move may be thought to fit uncomfortably with other kinds of evidence). Second, there is a worry that the ruling the account will give for certain specific cases of potential pragmatic enrichment of semantic content remains unclear. I'll consider these two points in order.

The first worry is that there are many statistically very common, regular uses that theorists nevertheless have often claimed to be instances of pragmatic, speaker meaning. So, for instance, take Grice's infamous category of conventional implicatures (see Grice 1989), where, e.g., an utterance of 'A but B' is held to be semantically equivalent to something like '((A&B) & (this combination of properties is unusual))'. On Millikan's preferred way of carving things up though, the very idea of a conventional implicature – of a pragmatically conveyed content which is reliably communicated whenever the expression is used – is simply impossible. If an utterance of 'but' conventionally conveys an element of comparison or tension then this element must constitute part of the signs semantic content. Of course, we might think that ruling out conventional implicatures is reasonable, given their tendentious nature (see Bach 1999), but there are also other, potentially more worrying cases. So consider Grice's category of 'generalised conversational implicatures', where the most famous instances involve 'scalar implicatures' (such as using the term 'some' to

⁷ See also 2017 (draft: 181): "The semantics/pragmatics distinction [may be] indefinite, in various degrees, over a significant portion of any language".

convey not the logical content ‘some and possibly all’ but the more refined content ‘some and not all’).⁸ Generalised conversational implicatures do, it seems, ‘typically recur on account of precedent’ (as Millikan’s definition of conventional usage above demands), but is this sufficient to make them semantic?

Well, it seems there are at least some reasons to think not, since locating these effects as part of the semantic content of the sign seems problematic in light of other kinds of evidence. For instance, amongst young children (even those who are in the process of acquiring a language from speakers who typically do use scalar terms to convey their enriched readings) the tendency is to give scalar expressions an unenriched, logical reading. So, for instance, children hearing a sentence like ‘Emile is wearing a hat or a scarf’ tend to judge the sentence true when Emile is shown wearing both a hat *and* a scarf (showing that children tend to interpret the scalar term ‘or’ as having its logical reading ‘A or B or both’), whereas adult English speakers exposed to the same prompt in the same context tend to judge the utterance as false (revealing that they are interpreting ‘or’ with the enriched scalar reading ‘A or B and not both’).⁹ So, treating scalar terms as having a semantic content that incorporates their enriched readings seems to fit poorly with acquisition data. Looking at conventional use amongst children, it seems that a scalar terms like ‘some’ should be taken to mean ‘some and possibly all’, whereas looking at adult English speakers, it seems that ‘some’ means ‘some and not all’. However, given Millikan’s view that functions attach to surface forms relativized to domains, we might think that her theory has the resources to cope with these kinds of findings. For perhaps we should simply allow that English has two different signs, with two different functions and thus two different meanings: there is the child-English ‘some’ used to convey the logical meaning, and the adult-English ‘some’ used to convey the enriched meaning. Whether or not this kind of move would be admissible within Millikan’s framework is, I want to suggest, not entirely clear and the reason for this is that it is unclear whether this is a legitimate way of carving up domains. Furthermore, this question – about how fine-grained domains are allowed to be – becomes more pressing when we turn to consider the category of so called ‘free pragmatic effects’.

‘Free pragmatic effects’, as noted above, are contextual contributions to content which are unmarked in the lexico-syntactic form of the expression and (MS2) was designed specifically to rule out such effects from contributing to semantic content. According to MS2, the route to semantic content runs along exclusively lexico-syntactic trails, so context is only relevant to semantic theorising where it is brought in to play by a formal element of the sentence, such as a demonstrative or tense marker. Contrary to MS2, however, it seems that Millikan’s approach will allow such effects to be semantically relevant (2004: 152):

Suppose that you are a surgeon and I am your assistant, and during an operation you direct ‘scalpell!’, then ‘scissor!’, then ‘Suture!’ Who is to do what with the scalpel, then the scissors, then the suture, and when and

⁸ For an introduction to scalar implicatures see e.g. Horn 1984.

⁹ See Papafragou & Musolino 2003, Pouscolous & Noveck 2009, Paltiel-Gedalyovich 2011.

where is not lexicalized, but it is determined by the context in an entirely conventional way. Similarly if I say ‘it’s raining’ the place at which I say this conventionally determines the place of the intentionally signified rain. For example, the following dialogue is not possible *within* the conventions of language (which is why it could be a joke).

‘It’s raining!’ ‘Where?’ ‘In Tahiti.’

‘It’s raining’, standing alone, simply is not a way you can conventionally say in English that it is raining somewhere or other...Adding lexical items that hold places showing where extralinguistic context is part of a sign surely simplifies matters for the interpreter rather than complicating matters.¹⁰

While this doesn’t seem quite right as a claim about the conventions surrounding ‘It’s raining’ in English (consider a context in which we are currently in the UK but are going to Tahiti on holiday soon and you are reading a newspaper which contains global weather reports – in this context the dialogue seems perfectly possible within the conventions of English), the quote does suggest that in general Millikan is willing to allow that at least some free (unmarked) contextual effects are semantically relevant (that ‘extralinguistic context’ can be ‘part of the sign’ as she puts it). Yet, as with any account of this form, we now need an answer to the question of which such effects are semantically relevant and when. Consider some examples:

Kind	Utterance	Conveyed content
Completions	Jill’s ready.	Jill’s ready <i>for the exam</i> .
Enrichments	This steak is raw.	This steak is undercooked.
Loose talk	It’s 2pm	It’s <i>roughly</i> 2pm
Particularised implicatures	There is a garage around the corner.	There is an <i>open</i> garage <i>just</i> around the corner.

In each of these cases, it seems that deciding whether a use ‘typically recurs on account of precedent’ will depend on the question of how thinly or thickly we have sliced domains. To see this consider, as an example, Grice’s famous example of a particularized implicature ‘There is a garage around the corner’: as Grice discusses, in a context where we are dealing with a driver who has run out of petrol, this utterance will typically convey that there is an open garage around the corner, but is that then the semantic content of the sign? Well, if the relevant domain is indeed ‘out of petrol’ events then the answer is ‘yes’, but then we will need to allow that the surface level sign can realise

¹⁰ See also 2004: 134-5 for a similar claim about quantifier restriction, comparative adjectives and possessives. Also *Beyond Concepts* (2017): “The usage-based picture of language suggests that much more can be said in purely semantic, conventional, although idiomatic ways, hence much less may need to be improvised through pragmatics than classical Gricean pragmatics supposed”. (Beyond concepts, draft version p. 181). However for at least some expressions, Millikan’s recent line seems potentially more minimalist than the above quotes suggest, thus in 2017 (draft p.217) she writes: “Possessives are sometimes proffered as cases where determinate meanings must be supplied by the speaker’s intention. ‘Peter’s book,’ for example, might refer to a book Peter owns or carries or is reading, to one he wrote or bought or brought, to one he is recommending or balancing on his head, whereas its semantic content seems to designate merely some kind of pairing relation between Peter and a book. Following Dretske’s comment again (§15.4), there is nothing in what the possessive says that makes it about one of these relations rather than another....Think of ‘Peter’s book,’ then, as meaning book bearing the pairing relation to Peter where the definite article functions just like ‘the’ in ‘I fed the dog.’ The possessive, like the definite article, prompts the hearer to identify the more specific pairing or associating relation that is at its ground, a pairing relation being talked about but not semantically identified”. (See also 2017: 220 for similar comments on incomplete definite descriptions and quantifier restriction.)

a myriad of different functions across a myriad of different domains. For instance, we will need to recognise at least the following domains:

Utterance	There is a garage around the corner.
Domain 1: English	There is a garage around the corner.
Domain 2: Out of petrol	There is an <i>open</i> garage <i>relatively near</i> around the corner <i>which sells petrol</i> .
Domain 3: Village facilities	There is garage around the corner <i>within the village limits</i> .

As with the question above of whether ‘some’ means ‘some and possibly all’ or ‘some and not all’, it seems we will get different answers about what counts as semantic and what counts as pragmatic content depending on how we individuate signs (and the functions they perform), but to do this we need an independent individuation of domains. Is the domain for an utterance of ‘There is a garage around the corner’ [English], [petrol-seeking events], [unleaded-petrol-seeking-events-amongst-UK-drivers], [reasonably-priced-unleaded-petrol-seeking-events], or what? Millikan’s model is, I think, that functions individuate signs, and domains individuate functions, but if this is right then we need to know what individuates domains (and it better not be signs or functions). Without an answer to this question it is not possible to predict what ruling the account will make on the kinds of controversial cases of potential pragmatic enrichment which have so exercised theorists in this area. The concern here is, at heart I think, a worry about whether or not the notion of ‘conventional function’ is one which is robust enough to yield determinate semantic content. Type level expressions will have a range of functions which might count as their conventional function (thus individuating their semantic content) depending on how thickly or thinly we carve out the domain to which the expression is bound: tied to English as a whole, ‘some’ conventionally seems to express ‘some and possibly all’, tied to adult-English it seems to conventionally express ‘some and not all’; tied to petrol-seeking events ‘there is a garage around the corner’ conventionally conveys that the garage is open, tied to village-facility domains, it doesn’t. So we can only answer the question ‘what is the semantic content of ‘there is a garage around the corner?’ once we know what the domain for any particular token of this type of sign is, but by whom, or what, and how, is this domain fixed?¹¹

¹¹ We might think that Millikan’s avowed commitment to the vagueness of the semantics/pragmatics divide can help here. For it seems that one might claim that, if there is no fixed answer to the question of the precise conventional function that an expression is playing, then there is simply no fixed answer to the question of its semantic content. So whether this given token of ‘there is a garage around the corner’ semantically conveys that the garage is open or only pragmatically implies this is simply underdetermined. However, allowing this kind of move seems to threaten to undermine some of our basic linguistic distinctions, such as that between lying or misleading, and may cause problems upstream, e.g. concerning how speakers acquire a language in the first place (see Borg 2017 for further discussion of these kinds of issues).

Finally, it seems that the questions we are grappling with here also resurface (in a slightly different form) with respect to MS1, so let's turn to this final tenet from minimalism now.

2.4) Propositions and compositionality (MS1)

Millikan advocates semantic holism, whereby the meaning of an expression is fixed by its place in a whole sign system. However, if we combine this holism with her stated commitment to:

- (a) meaning is linguistic cooperative function

And combine this with the view that:

- (b) functions attach to signs relativized to domains

Then I think we are likely to run once again into issues concerning how to deliver determinate content for signs. To see this let's look at an example representation:

(S) The mother bird is stimulated to release food by the sight of the hungry baby's mouth.

To determine the semantic content of S we need to know what its the stabilizing function is (i.e. the function performed by the expression which is responsible for its survival and proliferation), but here, at least *prima facie*, there seem to be multiple possibilities for the content (S) could represent or conventionally carry, such as:

- i. The mother bird is stimulated to release food by the sight of *the hungry baby's mouth*.
- ii. The mother bird is stimulated to release food by the sight of the *place where food needs to be put to maximise chances of off-spring survival*.
- iii. The mother bird is stimulated to release food by the sight of a *concave space in a chick's face*.

This kind of indeterminacy of content seems fine with genuinely biological signs (whether the mother bird is being stimulated by the baby's mouth or by the concave space in a chick's face seems unimportant from the point of view of survival of the species), but more troubling with fully intensional signs. For the sentence (S) it seems clear that we want to deliver (i), rather than (ii) or (iii), as the content of (S). Yet for a holistic theory to yield (i), instead of (ii) or (iii), we need the sign system to which S belongs to include other instances of the same sign elements that (S) contains. For in that case we can use familiar holistic methods to constrain the meanings of the parts of the sentence. So, for instance, it would turn out to be a mistake to translate 'baby's mouth' as 'concave space in a chick's face' as, although this translation would work for (S), it would yield the wrong results for other sentences which also contain this sign, such as (S*):

(S*) The mother put the dummy back in her baby's mouth.

Yet now the worries of the previous section resurface again, for we need to know what it is for two signs to be the same. We know that on Millikan's account it is not surface level sameness that matters (i.e. not orthographic or phonological form), so what is to stop someone holding that, say, the token of 'baby' in S, S*, or S** (e.g. "The baby wants a nap") are just different signs in different domains and as such can simply have different functions? For holistic constraints to get a purchase we *have* to be able to hold signs constant across different occasions of use involving different contexts, but the worry is that nothing in Millikan's account yet shows how we can type signs in a way that allows holistic constraints to get a purchase. We need an objective account of domain, or function, to stop things being carved up too finely, but again it is not immediately clear where this objective notion of proper function or actual domain will come from.

3) Conclusion

To conclude, there can be no question about the power or scope of Millikan's philosophical vision. Her ground-breaking work holds out the promise of a unified, naturalistic account of content covering representation in all its forms. This is surely something we'd very much like and which it is clearly very hard to deliver. In undertaking this vertiginous work, Millikan's account promises a novel take on many core issues in philosophy of language, and it is clear that much is to be learnt from trying to connect debates which often reside solely in philosophy of language with a larger, overarching account such as Millikan's. However, on the flip side, getting clear on exactly how her approach maps to certain contemporary views is not always easy. Thus in this paper I have tried to spell out her view with respect to the on-going debate between minimalists and contextualists. Doing so, I have argued, highlights a potential question about how we are to individuate the bearers of semantic content on Millikan's view. For Millikan, they are not words (typed by phonetic or typographic features) but rather expressions as they realise conventional functions. However, as reflection on Grice's example of 'there is a garage around the corner' showed, a communicative function is conventional or not only relative to a domain. So ascertaining what function an expression conventionally performs (and hence its semantic content) will require a robust account of how the domain to which a token of this expression is bound is individuated. Yet, while the individuation of a domain seems relatively easy for natural signs, I have suggested that it is much less transparent for linguistic signs. What approaching Millikan's philosophy from the perspective of minimalism highlights, I think, is that, if we want a ruling on issues around the semantics/pragmatics divide, we may need to know more from Millikan about how the individuation of a triumvirate of crucial notions – sign, function and domain – actually comes about for linguistic expressions.

Bibliography

- Bach, K. 1999. The Myth of conventional implicature. *Linguistics and Philosophy* 22: 327-366.
- Borg, E. 2002. Pointing at Jack, talking about Jill: understanding deferred uses of demonstratives and pronouns. *Mind and Language* 17: 489-512.
- Borg, E. 2004. *Minimal Semantics*. Oxford: Oxford University Press.
- Borg, E. 2012. *Pursuing Meaning*. Oxford: Oxford University Press.
- Borg, E. 2016. Exploding explicatures. *Mind and Language* 31: 335-355
- Borg, E. 2017. Explanatory roles for minimal content. *Nous* 51:3. Online first: doi 10.1111/nous.12217
- Grice, P. 1989. *Studies in the Way of Words*. Cambridge, MA: Harvard University
- Horn, L. 1984. Toward a new taxonomy for pragmatic inference: Q-based and R-based implicature. In Deborah Schiffrin, ed., *Meaning, Form, and Use in Context: Linguistic Applications*. 11–42. Washington: Georgetown University Press.
- Millikan, R. 2004. *Varieties of Meaning*. Cambridge, Mass.: MIT Press.
- Millikan, R. 2005. *Language: a biological model*. Oxford: Oxford University Press.
- Millikan, R. 2017. *Beyond Concepts*. Oxford: Oxford University Press.
- Paltiel-Gedalyovich, L. 2011. Adult response uniformly distinguishes semantics from pragmatics. In *Experimental Pragmatics/Semantics*, ed. J. Meibauer & M. Steinbach. Amsterdam: John Benjamins. 101-127
- Papafragou A., & Musolino, J., 2003. Scalar implicatures: experiments at the semantics–pragmatics interface. *Cognition* 86:253-282.
- Pouscoulous, N., & Noveck, I. 2009. Going beyond semantics: The development of pragmatic enrichment. In *Language Acquisition*. S. Foster-Cohen (ed) Palgrave. 196-215
- Predelli, S. 1998. Utterance, interpretation, and the logic of indexicals. *Mind and Language* 13: 400-414.
- Travis, C. 2008. *Occasion-Sensitivity*. Oxford: Oxford University Press