King's Research Portal

DOI:

[10.1002/9781118328118.ch3](https://doi.org/10.1002/9781118328118.ch3)

Document Version

Early version, also known as pre-print

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Shea, N. (2013). Millikan's Isomorphism Requirement. In D. Ryder, J. Kingsbury, & K. Williford (Eds.), *Millikan and Her Critics* (pp. 63-80). (Philosophers and Their Critics). WILEY-BLACKWELL.
<https://doi.org/10.1002/9781118328118.ch3>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

[Draft of 14 Sept 2010]

To appear in *Millikan & Her Critics*, J. Kingsbury & D. Ryder (ed.) (Blackwell)

Millikan's Isomorphism Requirement

Nicholas Shea

Abstract / Introduction

Millikan's theory of content purports to rely heavily on the existence of isomorphisms between a system of representations and the things in the world which they represent – “the mapping requirement for being intentional signs” (Millikan 2004, p. 106). This paper asks whether those isomorphisms are doing any substantive explanatory work. Millikan's isomorphism requirement is deployed for two main purposes. First, she claims that the existence of an isomorphism is the basic representing relation, with teleology playing a subsidiary role – to account for misrepresentation (the possibility of error). Second, Millikan relies on an isomorphism requirement in order to guarantee that a system of representations displays a kind of productivity. This seemingly strong reliance on isomorphism has prompted the objection that isomorphism is too liberal to be the basic representing relation: there are isomorphisms between any system of putative representations and *any* set, of the same cardinality, of items putatively represented. This paper argues that all the work in fixing content is in fact done by the teleology. Deploying Millikan's teleology-based conditions to ascribe contents will ensure that there is an isomorphism between representations and the things they represent, but the isomorphism ‘requirement’ is playing no substantive explanatory role in Millikan's account of content determination. So an objection to her theory based on the liberality of isomorphism is misplaced. The second role for isomorphism is to account for productivity. If some kind of productivity is indeed necessary for representation, then functional isomorphism will again be too liberal a constraint to account for that feature. The paper suggests an alternative way of specifying the relation between a system of representations and that which they represent which is capable of playing an explanatory role in accounting for Millikan's type of productivity. In short, the liberality of isomorphism is no objection to Millikan's teleosemantics, since the isomorphism ‘requirement’ need play no independent substantive role in Millikan's account of representation.

Contents

- (1) Isomorphism and Functional Isomorphism
- (2) Millikan's Reliance on Functional Isomorphism
- (3) Isomorphisms and Productivity
- (4) Exploiting 1-1 Maps Which Preserve Natural Relations
- (5) Conclusion

(1) Isomorphism and Functional Isomorphism

An isomorphism is a mathematical function or mapping between two sets of items. The items in these sets can be concrete or abstract. A function takes each element in the *domain* to an item in the *range*. So the mapping from people to their mothers is a function. $Mother(x)$ takes each person to the one person who is their biological mother. If every item in the domain gets mapped to a different item in the range, then the converse map, taking each element in the range back to its corresponding element in the domain, will also count as a function. The map from adults to their social security numbers is like that. $SocSec(x)$ takes adults as input and delivers for each a unique social security number. The inverse map takes as input social security numbers that are in use, and for each delivers a unique person.

We can use f^{-1} for the inverse of f . So $\text{SocSec}^{-1}(x)$, taking as its domain the social security numbers that are actually in use, is a well-defined function. Contrast $\text{Mother}(x)$. It's inverse is not a function, since it maps some people to more than one person (it maps my mother to all my brothers and sisters).

When a function has an inverse that is also a well-defined function, notice that we have a correspondence relation. Every element in the domain corresponds to a unique element in the range, and vice versa:

At its simplest, an *isomorphism* is such a 1-1 correspondence between two sets of items.

Whether there is such an isomorphism will depend on what we are taking the domain and the range to be. Defined on the domain of children in my son's parent & child music class, the *parent of* relation picks out a function onto a range consisting of adults in the class. Since it happens that each parent brings only one child, the inverse is also a function

(from adults in the class to children in the class). That sets up a 1-1 correspondence or isomorphism between the parents and the children.

These are the standard kinds of examples used to introduce the concept of an isomorphism. They can be misleading in an important way. In each, the isomorphism described corresponds to some real relation: the natural relation of being a parent, the society-based relation between people and unique social security numbers. But an isomorphism need not correspond to any natural or systematic relationship at all. Any way of lining up each element in the domain with a corresponding unique element in the range defines an isomorphism. So when there is one isomorphism between multiple items, there will always be many others too. If we were to pair up each child in my son's music class with a different adult chosen at random, that too would be an isomorphism:

The crucial thing to notice is that there *are* all these isomorphisms - all at once, we might say. If we allow ourselves to talk about the *existence* of an isomorphism (some kind of abstract mathematical object), I, I* and I** all exist (without the adults and children in the class having to change in any way). There is, indeed, something special about isomorphism I. It corresponds to the natural relation of parenthood. But as an isomorphism — a mathematical object — I is absolutely on a par with I*, I**, and all the many other isomorphisms between the two sets.

The term 'mathematical function' can mislead in a similar way. Familiar examples of mathematical functions do something regular and systematic to the elements they are defined on. They are functions like $x+10$, $7x$, x^2 and $\cos(x)$. The map from 1, 2, 3, 4 to 1, 4, 9, 16, respectively, is a mathematical function (squaring), but the map from 1, 2, 3, 4

to 76, $\sqrt{2}$, $-\frac{5}{8}$, Π^e , respectively (chosen so that there is no systematic relation – that I can see) is just as much a mathematical function as is squaring. Indeed, the mathematical concept of function equally includes the mapping from 1, 2, 3, 4 to Tibbles the cat, this orange on my desk, the sun and the Queen. That mapping, which also happens to be an isomorphism, exists alongside squaring in the abstract realm of functions. Something more than function talk is needed if we are to say why squaring is more interesting in some way. So the question of whether *there is* an isomorphism between two sets of items is rather straightforward. If the two sets have the same number of items (or the same cardinality, to include transfinite sets), then there will be an isomorphism between them. And if there's one, there will be many isomorphisms between the two sets.

Before turning to Millikan's theory, we need one final, crucial piece of mathematical machinery. So far, we have been considering just unstructured sets of elements in the domain and range, in which case any 1-1 map is an isomorphism. But the more useful and important concept of isomorphism is that of a structure-preserving map. Once we consider the sets of elements in the domain and range as being endowed with some structure, we have to ask whether those structures also correspond under I . Unlike the bare existence of a 1-1 map, if we have two structures in mind, it is a very demanding constraint that there should be a structure-preserving map between them. The structures in question may be relations between or operations on the elements, for example. Philosophy has paid particular attention to 'functional isomorphisms' – where the 'function' is for mappings between elements *within* the domain of the isomorphism, and for corresponding functions between elements within the range of the isomorphism. Consider again my son's music class, with parent-child pairs sitting around in a circle. There is a function f within the children that takes each to their left-hand neighbour (the range and domain of f are the same set – the children in the class). And there is a corresponding function g within the adults, taking each parent to the parent on their left:

So here we have not only a 1-1 correspondence between parents and children, I , but also a correspondence between the functions f and g : under I , a maps to A , and the child to the left of a , $f(a)$ (i.e. b) maps to the child to the left of A , $g(A)$ (i.e. B). Or to put it more formally:

$$I(f(x)) = g(I(x))$$

In that case, we say there is a *functional isomorphism* between function f on the children and function g on the parents.

Recalling the liberality of the mathematical concept of a function, it will be obvious that for every function defined on the elements of the domain of some isomorphism I , there is a corresponding function on the elements of the range of I . We can use f and I to generate such a function g as follows. For each y in the range of I we use I^{-1} (the inverse of I) to see which element is mapped to y (e.g. to see which child belongs to that parent), we perform the given function f on that element (e.g. f takes us to the next child to the left), and then we use I to map that element back to the corresponding element in the range of I (the parent of the child to the left of {the child of the parent y we started with}). To give a graphical example, we can see that g will take element A to element B :

To get g ,

In short, we can use f and I to define the corresponding g as $g(y) = I(f(I^{-1}(y)))$.

Given the liberality of functions, whenever there is an isomorphism between two sets, there will be very many functional isomorphisms between them. And for any given function on the domain of I , there is some function on the range of I to which it is functionally isomorphic (also if we start with a function on the range). So it is a rather trivial question to ask whether there is *some* function on a range of elements D (for example, on a set of putative representations) which is functionally isomorphic to some function of interest defined on a domain of elements R (potential representeds), because there will be such a function iff there is an isomorphism between D and R , which there will be iff they have the same cardinality. We can, however, ask a more substantive question with respect to a function f on D *and* a function g on R , namely whether there is a functional isomorphism between those particular functions f and g . So we might ask whether a particular function of interest on D (e.g. the one given by the left-of relation) is functionally isomorphic to a particular function of interest on R (e.g. the one given by the left-of relation) under some 1-1 map between D and R (as there is: the parent-of function is a 1-1 map that preserves the rotational structure - it takes f to g). It is by no means guaranteed that the left-of function on the children corresponds to the left-of function on the parents. (Notice that if there were such a correspondence for some natural reason, then it would something we could make use of, e.g. to find out which adult is to the right of A by checking to see which child is to the right of a .) Whereas the existence of an isomorphism between sets of elements bare of any structure is a very liberal requirement, the existence of an isomorphism between two sets of elements endowed with structure is a very demanding requirement. The entities in mathematics are exhaustively characterised by their roles in structures. Accordingly, proving that two sets of such elements are isomorphic is a very strong result, because it means that everything true in one system is true in the other - there are no mathematically-relevant differences between the systems. The liberality that arises when deploying isomorphism talk in the field of theories of content arises, to a first approximation, because the relevant structure on the set of elements represented is usually unspecified.

This careful walk through the mathematical concepts should give us the tools to be

able properly to assess Millikan's claims about the explanatory work that can be done by isomorphisms and functional isomorphisms. The main lesson so far is that they are exceedingly liberal relations. Paradigm cases in which isomorphisms correspond to natural relations and functional isomorphisms subsist between naturally-defined functions are prone to mislead us about how much substance there can be to the claim that there *is* an isomorphism or a functional isomorphism between two sets. In the next section we will examine ways that Millikan seeks to rely on these concepts.

(2) Millikan's Reliance on Functional Isomorphism

Millikan requires there to be an isomorphism between a system of representations and that which they represent. This section sets out the explanatory work that Millikan claims her isomorphism requirement can perform. The central question for the paper is whether the isomorphism requirement acts as a substantive constraint within Millikan's theory. Millikan relies on it in two main ways. First, it supports the claim that representational systems display productivity. Second, Millikan claims that a particular isomorphism between representations and representeds plays a role in explaining why the representations in a given system have the contents they do. This section sets out those claims.

Millikan most frequently turns to functional isomorphisms to account for productivity. She argues that representations must display productivity to be genuinely intentional. The kind of productivity she has in mind arises when representations come in systems, with transformations between elements in the system corresponding to systematic transformations of that which is represented. For example, in the honeybee's nectar dance, the angle of the dance to the vertical corresponds to the direction of nectar in relation to the sun. So a transformation of the dance, 10 degrees further away from the vertical, say, corresponds to a transformation of the represented location, 10 degrees further away from the direction of the sun. The argument is (i) that there are such functional isomorphisms; and then (ii) that they give rise to a kind of productivity. The first statement is in *Language, Thought and Other Biological Categories* ("LTOBC"):

When an indicative intentional icon has a real value, it is related to that real value as follows: (1) The real value is a Normal condition for performance of the icon's direct proper functions. (2) There are operations upon or transformations (in the mathematical sense) of the icon that correspond one-to-one to operations upon or transformations of the real value such that (3) any transform of the icon resulting from one of these operations has as a Normal condition for proper performance the corresponding transform of the real value. (Millikan 1984, p. 107)

The claim is that there will always be a 1-1 correspondence (isomorphism) between a function on the system of representations (the domain) and a function on the worldly affairs that are represented (the range). That is the basis for a kind of productivity. Millikan treats compositionality as a species of productivity. Compositionality arises from replacing parts of a representation: if I can think <John is tall> and I have the concept MARY, then I can think <Mary is tall>. Millikan sees variant aspects of the representation playing a similar role: the angle of the bee dance changes, allowing different dances to represent different directions. The representation 'is not articulated into parts but into invariant and variant aspects' (1984, p. 107). The bee dance always says *nectar at ...* – that much is invariant – but different dances say the nectar is at different distances, and in different directions. Millikan is strongly committed to the idea that this is genuine productivity, on an equal footing with the more obvious productivity of natural language sentences (Millikan 1984, p. 108). Whether she is right about that is a matter for another day. For now, the point is that an isomorphism is claimed to underpin a particular feature of systems of representations: that relations between the representations correspond to relations between the items represented (a relation-preserving isomorphism). That claim recurs throughout Millikan's work (1984, ch. 5; 1989, p. 287; 2004, ch. 6; 2005, ch. 5).

Between which items is this isomorphism supposed to subsist? In the domain, are we dealing with actual bee dances (for example) or merely potential bee dances? And in the range, do we find actual locations of nectar, or rather potential or represented locations? It is clear that we are not dealing with a correspondence between actual dances and actual locations, since Millikan has designed her theory to allow for misrepresentation, so there

will be instances when there is no nectar at the location that an actual dance is supposed to correspond to. We could restrict our attention to the types actually found in the history of the species: dances danced and nectar locations on occasions that led systematically to survival and reproduction. These are the cases that serve to fix the correctness condition of the various kinds of dance. But Millikan wants dances that happen not to have occurred at all in the evolutionary history to count as instances, provided they fall within the general overall pattern. Perhaps, by chance, no dance was ever danced at 84 degrees to the vertical. Still, that potential dance has an associated truth condition, given by the general rule that angle of dance to the vertical maps to direction of nectar with respect to the sun. For these reasons, we should think of the relevant isomorphism as subsisting between *potential* dances and *potential* locations of nectar. And the structure preserved is rotation: the operation of rotation on the angle of potential dances corresponds to the operation of rotation on the potential direction of nectar. That defines a functional isomorphism from {potential dances with rotational structure} to {potential directions with rotational structure}.

This way of understanding the domain and range of the isomorphism underpins the reason why Millikan sees variant structure in the domain of representations as a kind of productivity. The isomorphism specifies a way of giving the content of a new dance, not previously danced in the history of the species. That is, once the relevant correspondence relation between potential dances and potential locations has been established, that correspondence relation may assign contents to new representational items. A question remains about how such a correspondence is established. It looks as if a theory of content is needed to determine what is the represented location that corresponds to each of the putative representations (dances). We will see below that relational proper functions may establish correspondence relations that extend to genuinely new cases. What Millikan's isomorphism requirement adds, then, to the mere existence of a correspondence relation established by the theory of content, is that there should be some structure over the representations that recapitulates some structure over the items that are represented. In examples like the bee dance, a natural relation on the system of representations (rotations of the dance) is mapped to a natural relation on the items represented (rotations of the

direction of nectar). But, as we saw above, the mathematical concept of isomorphism does not bring with it a restriction to natural functions or relations. Once a theory of content delivers a correctness or satisfaction condition for each representation, there will be lots of functions (in the mathematical sense) on the representations that correspond to functions on the items represented, and vice versa.¹ So it doesn't look like the requirement that there be a functional isomorphism can be doing any explanatory work. Any adequate theory of content delivers a 1-1 map between representations and the potential states of affairs they represent, and that map will preserve *some* structure on the elements, and so count as a functional isomorphism.

The existence of such isomorphisms may have more bite if we restrict our focus to certain kinds of structure on the potential representations / representeds. Many of Millikan's examples are ones where natural relations on the representations correspond to natural relations amongst the individuals or properties represented. I take up below the question of the explanatory work that may be performed by such a restriction (section 4).

So far, we have raised the suspicion that the isomorphisms that Millikan says underpin a kind of productivity are no more than a necessary consequence of the fact that a theory of content sets up a correspondence between potential representations and truth conditions. This is where Millikan's second role for isomorphism comes in, because she claims that mapping functions enter into an explanation of why representations have the content they do (1984, pp. 99-100, 246; 2000b, p. 6;² 2004, ch. 8; 2005, p. 102). That is, they are prior to the theory of content, rather than the result of it.

Recall the structure of Millikan's teleosemantics (simplifying considerably). A behavioural system separates into co-operating producer and consumer subsystems, with the producer giving rise to a range of intermediates, which cause the consumer to behave in a variety of ways. Amongst all the ways in which a consumer might behave in response to a given intermediate R we focus on the types of behavioural response B that, in the

---

¹ Strictly, the relation picked out by the theory of content will not be an isomorphism unless the theory assigns a unique content to each different representation (otherwise, the inverse would not be a well-defined function). To deal with cases of redundancy – different representations with the same correctness or satisfaction condition – we need to generalise to include *homomorphisms*.

² Page references to Millikan 2000b refer to the readily-available version available in pdf format from Ruth Millikan's website at uconn.edu.

evolutionary past, led systematically to the survival and reproduction of the system (i.e. that are part of a Normal explanation). Doing B is then a candidate for the imperative content of R, and worldly conditions that enter into a Normal explanation of why B leads to survival and reproduction are candidates for the indicative content of R. Millikan claims that isomorphisms (mapping rules) enter into that Normal explanation. One of the most concise statements in LTOBC concerns beliefs (although the point applies to representations in general):

In order to [be] an intentional icon, the belief token ... must have as a proper function to adapt an interpreter device to conditions in the world so that this device can perform proper functions, and it must be part of the most proximate Normal explanation for the interpreter's proper performance that the belief – the inner sentence – maps conditions in the world in accordance with some definite mapping rule. (Millikan 1984, p. 146)

Here we have the idea of a mapping rule entering into a Normal explanation of the operation of the system. That is supposed to be prior to content determination. What makes it the case that representations have the content they do is that there is a Normal explanation that mentions a particular mapping rule. The consumer system acted on intermediates in a way that led historically to successful behaviour, a mapping rule is part of an explanation of that success (1984, pp. 99-100), and in virtue of all of that, that mapping rule becomes the rule for delivering truth conditions or satisfaction conditions. There is an important distinction here between a representation actually *mapping* onto a condition and a *mapping rule* (a distinction that sometimes turns into an ambiguity in the use of the term). A *mapping rule* is an isomorphism between two sets of items. There being an isomorphism between two sets of items does not depend upon them being concurrent or co-present, or upon their standing in any other natural relation. (Indeed, we remarked above that the relevant isomorphism was between types of entities rather than tokens: between potential dances and potentially represented locations, for example.) By contrast, a representation *maps* some condition, in Millikan's terminology, when the

representation is actually tokened and the condition actually obtains. For example:

Intuitively it is clear that in some sense of “mapping,” the bee dance that causes watching bees to find nectar in accordance with a historically Normal explanation is one that maps in accordance with certain rules onto a real configuration involving nectar, sun, and hive. As such it is an indicative intentional icon. The bee dance also maps onto a configuration that it is supposed to produce, namely, bees being (later) in a certain relation to hive and sun — that is, where the nectar is. So the bee dance is also an imperative intentional icon. (1984, p. 99, emphasis added)

Millikan here uses ‘maps’ for a relation obtaining between a particular dance and real nectar at an actual location. The properties that are connected by a mapping rule can be instantiated: a dance is performed in a particular way and there is nectar at a particular direction and distance from the hive where it is performed. When the types so instantiated are indeed connected by a mapping rule, Millikan says the dance ‘maps’ the location. Whether a particular dance maps a particular location depends upon which mapping rule is in question. Is Millikan right to claim that the existence of a mapping rule is a substantive constraint on the theory of content?

That there should be such an isomorphism does not act, in Millikan’s theory, as a substantive constraint on content. On occasions that led to survival and reproduction, particular dances were performed and nectar was found at particular locations. We can indeed generalise across those instances: for a dance at θ to the vertical, there was nectar at θ degrees from the direction of the sun. That generalisation does enter into a Normal explanation of the success of the behaviour: operation of the producer-consumer system led to survival and reproduction when producer bees produced dances at θ to the vertical and there was nectar at θ to the hive. It is this natural relation between dance-types and nectar-location-types, not the isomorphism as such, that explains success (and hence gives the content).

There are two points here. The first is that it is not the isomorphism — the 1-1 mathematical function — that enters into the Normal explanation, but instead some natural

relation between dance types and locations. The second is that it is the Normal explanation which picks out a particular isomorphism. It would get things precisely the wrong way round to claim that there being an isomorphism is a substantive constraint on the theory of content. Instead, Millikan's account of content delivers an isomorphism. It may be relatively benign to talk as if it is the isomorphism as such, rather than a corresponding natural relation, which explains successful behaviour. But that leads to the second, more serious mistake, which is to claim that the existence of isomorphisms is part of what makes it the case that representations represent as they do – that they are prior to the theory of content.

Millikan sometimes writes as if isomorphism (or picturing) is the basic representing relation, and that teleology only comes in to account for error:

... naturalist theories of the content of mental representation are often divided into, say, picture theories, causal or covariation theories, information theories, functionalist or causal role theories, and teleological theories, as though these divisions all fell on the same plane. That is a fairly serious mistake, for what teleological theories have in common is not any view about the nature of representational content. "Teleosemantics," as it is sometimes called, is a theory only of how representations can be false or mistaken, which is a different thing entirely. Intentionality, if understood as the property of "ofness" or "aboutness," is not explained by a teleological theory. ...

... What teleological theories do not have in common is an agreed on description of what representing – what "ofness" or "aboutness" – is. They are not agreed on what an organism that is representing things correctly, actually representing things, is doing, hence on what it is that an organism that is misrepresenting is failing to do. To the shell that is "teleosemantics" one must add a description of what actual representing is like. When the bare teleosemantic theory has been spent, the central task for a theory of intentional representation has not yet begun. Teleosemantic theories are piggyback theories. They must ride on more basic theories of representation, perhaps causal theories, or picture theories, or informational

theories, or some combination of these. (Millikan 2004, pp. 63, 66)

[Teleological] theories generally begin with some more basic theory of the relation between a true thought, taken as embodied in some kind of brain state, and what it represents, for example, with the theory that true mental representations covary with or are lawfully caused by what they represent, or that they are reliable indicators of what they represent, or that they "picture" or are abstractly isomorphic, in accordance with semantic rules of a certain kind, with what they represent. The teleological part of the theory then adds that the favored relation holds between the mental representation and its represented when the biological system harboring the mental representation is functioning properly ... . (Millikan 2003, [p. 4 of m/s])

Millikan relies on isomorphism or picturing as the basic representing relation on which her teleosemantic theory is based (1984, pp. 9, 11; 2003, [p. 9 of m/s]; 2004, p. 79). That has led to the criticism that the concept of isomorphism is far too liberal for it to be able to do any substantive explanatory work. We saw the reasons for that liberality in section 1. Godfrey-Smith provides the clearest statement of the objection (Godfrey-Smith 1996, pp. 184-187).

However, a better interpretation is that Millikan is not taking the existence of a pre-existing isomorphism to be a substantive constraint on the theory of content. Instead, the facts about Normal explanations pick out an explanatorily-significant natural relation (eg: between dances at θ to the vertical and nectar at θ to the sun), and this relation specifies a particular isomorphism I between representation types and properties represented. The fact that there is such an isomorphism does no explanatory work. (After all, there is an isomorphism between the system of representations and very many other ways of assigning content to them.) What enters into the Normal explanation is that there was nectar at θ to the sun when a dance was produced at θ to the vertical. That reading of Millikan is clearest from the following:

Papineau and Millikan claim that it is only the uses to which mental representations are put that is relevant to their content. Millikan claims that a true representation maps onto its represented in accordance with semantic rules *determined by* the way the systems using the representation are designed to react to it in guiding, perhaps first inference processes, but ultimately behavior. (Millikan 2003, [p. 8 of m/s], italics added)

That passage makes it clear that the relevant isomorphism ('semantic rule') is determined by the way consumer systems are designed to react to the representation. So the *existence* of isomorphism I plays no explanatory role. The existence of I is not part of what makes it the case that the representations have the content they do. But the theory of content will deliver a correspondence, giving a correctness or satisfaction condition for each representation in the system to which the theory applies. So this content assignment will indeed pick out an isomorphism, which Millikan calls the "semantic rules" or "mapping rules" for the system of representations. There is indeed an isomorphism or correspondence relation that gives the content of a system of representations. But it is wrong to suggest that the isomorphism is basic and the teleology is needed only to account for error. The bare existence of an isomorphism between a system of representations and potential contents they could represent plays no explanatory role.

Once we have dismissed the idea that the existence of an isomorphism is prior to or explains content determination, the only substantive role left for an isomorphism requirement is the first one identified above – to account for productivity. In the next section we examine that role.

(3) Isomorphisms and Productivity

So far, we have seen that Millikan's theory of content will serve to identify a particular isomorphism between representation types and represented properties (/objects/states of affairs) as privileged – as giving the content of the representations (as will any theory of content, subject to the point about one-many mappings / homomorphisms). We have also

seen that it is an automatic consequence that there will be operations on the domain of representations that correspond to operations on the range of representeds (indeed, there will be many such operations). We noted that the bee dance has an additional feature. There, the operations which are functionally isomorphic as between representations and represented correspond to natural relations (rotation of dances and rotation of direction). Is it a further requirement on intentionality that the functions that correspond under the relevant functional isomorphism be ones that correspond to natural relations? Millikan is clear that it is not:

Isomorphisms can be defined by functions that are as bizarre, as gruelike, as you please. (Varieties, ch. 6, p. 84)

That is what she should say. It follows from the fact that the relevant isomorphism (“semantic rule”) is determined by the way consumer systems are designed to react to intermediate representations in guiding behaviour. Suppose bee dances represent nectar at the distances given by I (on the left below):

I		I*	
1 waggle	300m	1 waggle	75m
2 waggles	150m	2 waggles	300m
3 waggles	100m	3 waggles	60m
4 waggles	75m	4 waggles	150m
5 waggles	60m	5 waggles	100m

An explanation of how the bee dance system led to survival and reproduction appeals to the fact that, in the historical past, consumers of dances of 3 waggles were disposed to fly off 100m before searching for nectar, and such dances contributed systematically to survival and reproduction when there was nectar 100m from the hive; similarly for each of the other variants. But had the consumer bees had the disposition to fly the distances given

under I^* above, then the bee dance system would have contributed systematically to survival and reproduction when there was nectar at the distances given by I^* (provided producer bees ensured the number of waggles correlated with nectar at those locations often enough to be worth acting on). So I^* would give the content. That is just to reiterate Millikan's point that the relevant isomorphism is determined by the way consuming systems are designed to react.

Of course, the set of dispositions given by I can be more compactly described: consumer bees were disposed to fly off a distance given by $300\text{m}/\#\text{waggles}$. A natural relation amongst the dances ($\#$ waggles) corresponds to a natural relation amongst the locations (distance^{-1}). But I^* is just as much a functional isomorphism as I is. The function f on the number of waggles corresponds to the function g on the distances, where f and g are given by:

f :	g :
1 waggle $\rightarrow$ 2 waggles	75m $\rightarrow$ 300m
2 waggles $\rightarrow$ 3 waggles	300m $\rightarrow$ 60m
3 waggles $\rightarrow$ 4 waggles	60m $\rightarrow$ 150m
4 waggles $\rightarrow$ 5 waggles	150m $\rightarrow$ 100m

Of course, g is not a very natural function; but it is just as much a function for all that. The example demonstrates why Millikan has to accept bizarre or gruelike isomorphisms. So the functional isomorphism 'requirement' — that there must be operations on the representations that correspond to operations on the representeds — is not a substantive constraint on the class of isomorphisms that are admissible as giving the content of a system of representations.

When it comes to relying on an isomorphism to account for productivity, it seems that isomorphisms that have a natural structure like I have an advantage over more arbitrary mappings like I^* . Cases like I arise because there is some mechanism that acts in a systematic way on the intermediate representations. So the function that describes I compactly — $\text{distance} = 300\text{m}/\#\text{waggles}$ — may correspond to a single mechanism by which the system operates.³ If so, the facts about selectional history will establish that the dance

---

³ It has been suggested that there is a straightforward explanation of why *incoming* bees first behaved in the corresponding way. The further they have come from the nectar, the less energy they have left when they

producing mechanism has a *relational proper function* - to direct the consuming bees to a place at the distance given by the rule. Similarly for direction, dance producers have the relational proper function of causing consumer bees to fly in a direction corresponding to the angle of the dance. A particular dance has its own proper function, deriving from this relation - a *derived proper function*.⁴ For example, a dance at 84 degrees to the vertical has the derived proper function of sending consumer bees in the direction 84 degrees from the direction of the sun. The particular dance has that derived proper function, whether or not a dance of 84 degrees ever occurred in the evolutionary history. The upshot is that when a system of representations have a proper function arising from natural properties of the system during episodes of selection, then the contents of those representations may be expressed in terms of that relation. And that does indeed give rise to a kind of productivity. Current instances which have not occurred in the history of selection may nevertheless fall under that relation, and so have derived proper functions, hence content. Notice that the list-type cases will give rise to relational proper functions, too. In the case of I^* above, the relevant relational proper function is given by I^* . But such list-type isomorphisms will not apply to any new cases. The application to new cases arises from there being something systematic in the natural operation of the system of producer-representation-consumer during the history of selection.

The bee dance producer has two relational proper functions (the ones for direction and for distance). That results in further opportunities for productivity. Even if nectar at 84 degrees is part of the selectional history, and nectar at 300m is part of the selectional history, it may be that nectar 300m away and at 84 degrees happens not to be part of the selectional history. If so, such a dance will have a derived proper function, and a content, even though no instance of that dance was ever the basis of selection. That is a kind of productivity.

It also looks like a kind of compositionality. But it is important to notice that it is not compositionality in the regular sense. In conceptual thoughts and natural language sentences, the constituents make no claims on their own. They are unsaturated. For

---

reach the hive, so the fewer waggles they perform.

⁴ It also has an *adapted proper function*, but we do not need that part of Millikan's terminology in what follows.

example, the predicate ‘is tall’ is perfectly meaningful, but it has no truth condition or satisfaction condition. Contrast the bee dance case. There, each relevant dimension of variation has its own truth condition. A dance of 3 waggles says that there is nectar 100m away, irrespective of the direction in which it is performed. There are no separate unsaturated terms, like a subject and a predicate. For some purposes it will be important to distinguish between these cases.⁵ Sometimes a system of representations contains two or more dimensions of natural variation, but each dimension has its own fully-formed representational significance (truth, satisfaction, etc.). Multiple variant aspects build up a list of truth conditions, roughly as if they were terms in the propositional calculus connected by conjunction. In other cases like the conceptual constituents of a belief, the variant aspects Millikan describes are each unsaturated, and build up only together into a truth-evaluable thought, like terms in the predicate calculus.

To recap, we saw in section 2 that it would be a mistake to think that some isomorphism between representations and represented is part of an explanation of why representations have the content they do. Rather, a Normal explanation will advert to physical properties of the system of producer and consumer which make it the case that representations have the content they do – the theory picks out a special isomorphism, rather than being based on the existence of an isomorphism. In this section we saw that, where the selectional story supports a mapping rule that extends to new cases, there is a kind of productivity. But that works via relational proper functions arising from natural relations, rather than being based on the existence of some pre-existing isomorphism. Finally, we saw that it is not built into the theory that bizarre or grue-like isomorphisms are excluded. The concluding section suggests that more natural isomorphisms may have a further significance.

(4) Exploiting 1-1 Maps Which Preserve Natural Relations

Recall the contrast between isomorphisms I and I^* above. Under I , a natural relation on the

---

⁵ For example, the idea of separate relational proper functions applies unproblematically to the first case. It is less clear how it applies to unsaturated representational constituents.

system of representations corresponds to a natural relation on the properties represented. Rotation of the dance away from the vertical corresponds to rotation of the direction of nectar away from the sun. Under I^* , there is no such correspondence between natural operations. Given an appropriate consumer, I^* would qualify as giving the content of the range of representations. But surely there is some natural advantage of I over I^* ? It looks less likely that there would be a consumer that reacts in accordance with I^* . This section sketches a reason why that reaction may be justified.

A first reason is just that smooth or natural correspondences are more likely in nature. The consumer system has to react to a range of intermediate representations. Natural relations between representation-type and appropriate behaviour may be easier to implement, just from an engineering standpoint. So although a consumer-based theory of content tells us what the contents would be if there were some grue-like connection between representations and behaviours of the consumer that figure in a Normal explanation, in natural cases we are more likely to come across natural relations. So, at the very least, there may be natural reasons to expect many cases to be like the bee dance.

Does the point run deeper? Isn't it some kind of achievement to produce a structure (a system of representations) that is isomorphic to things in the world? It took years of work involving many pieces of technology for the Survey of India to produce systems of representations (maps) that bear rich isomorphisms to portions of the Indian countryside. And having these artefacts was immensely useful in controlling and defending the territory. If isomorphisms are so cheap, why is it such an achievement to produce items that stand in such relations?

Although Millikan has emphasised the role of consumers in her theory of content, she also writes in places as if it is an achievement of representation producers that they give rise to a system of representations that are actually isomorphic to some set of things in the world (1984, p. 146; 1995, [p. 4 of m/s]). We saw in section 2 that such isomorphisms are found in specifying the normal conditions for the producer to perform its function. That does not imply that producers actually do produce a system of representations that bear any interesting isomorphisms to properties of interest. From the discussion in section 1 it

should be obvious that to produce a system of representations for which there was some isomorphism to relevant objects and properties in the world is utterly trivial. What more substantial thing might producers be doing? Godfrey-Smith talks about the producer-consumer framework as being vindicated when the producer system produces a range of states that bear some *exploitable relation* to features of the world that are relevant to the system (Godfrey-Smith 2006a). The mere fact that there is *some* isomorphism between the representations produced by the producer system and relevant properties in the world is not by itself an exploitable relation. Isomorphisms are cheap, but most are not exploitable. Is there some special class of isomorphisms such that, to produce a system of representations bearing such an isomorphism to relevant objects and properties really would be an achievement, and the structure that resulted would thereby bear an exploitable relation to the world? The importance and value of cartography suggests that there is.

In the bee dance case, the consumer system is exploiting the fact that producer bees produce dances whose angle carries information about the direction of nectar (and about distance). Millikan sometimes claims that consumers are making use of the fact that representations are produced to carry information in something like Dretske's sense (1984, p. 146; 2004, p. 79).⁶ Subsequently, Millikan has developed the notion of 'soft natural information' (2000a, app. B) or 'local natural information' (2004, ch. 3), generalising Dretske's notion; and there are closely-related notions in the mathematical theory of information, like mutual information based on Kullback-Leibler distance (Cover & Thomas 2006, ch. 2). I use the general term 'correlational information' for all such relations, which trace back to Shannon's seminal work (Shannon 1949). Carrying correlational information is an exploitable relation in Godfrey-Smith's sense.

In the bee dance case the consumer is exploiting the correlational information carried by the dances it observes. The functional isomorphism between {rotation on dances} and {rotation on directions of nectar} is not being exploited as such. The grue-like isomorphism I* above would be just as good for consumers, if the correlational information carried by each dance was as reliable. So this is a case where, although there is an interesting

---

⁶ But it is no part of the theory that representations must carry correlational information (1995; 2004, pp. 67-68; 2007).

isomorphism between natural relations on the representations and some relevant natural properties, that natural relation is not being exploited as part of what allows the consumer system to perform its functions.

In other cases, it may be important to see that an isomorphism really is being exploited, rather than just being a side-effect of the natural operation of the system. There is not space here to do justice to this idea, so a sketch will have to suffice. There are two elements to the picture. (1) what distinguishes the exploitable isomorphisms from those that are ubiquitous and cheap? (2) what is it for a consumer mechanism to exploit an isomorphism from this limited class? I have already hinted at an answer to the first question. When natural relations on a set of putative representations correspond to natural relations on a set of things in the world, then that is indeed something that a natural system can make use of. Any old isomorphism would be useful if the consumer knew the code. But correspondences between natural relations are easier codes to crack. Such a correspondence will be useful to the consumer if a natural relation on the representations is something that it can detect and respond to differentially with different output behaviours. And the natural items in the world to which the representations are isomorphic must be *relevant* to the consumer, in the sense of potentially making a difference to whether its behaviours are successful (ultimately a matter of survival and reproduction, in Millikan's teleosemantic framework).

This is the sense in which my map of Oxford bears an exploitable relation to the city of Oxford. The natural relation of 2D spatial position on the map is one which I can easily read and make use of. And this readable relation corresponds to properties of the world which are relevant to the success of my behaviour, namely the spatial relations between places in the city. I exploit those relations whenever I use the map to calculate a route. Suppose I am inclined to trace routes on the map, measure them, and take the shortest. Then I am making use of the fact that the shortest route in centimetres on the paper will correspond to the shortest route in metres on the ground. A map can also be rigged up to carry correlational information, of course. That happens when a car satellite navigation system uses GPS to plot your current position on a map. But online correlational information cannot be the whole story of how a map is used to calculate a new route.

In short, an isomorphism between readable natural relations on a system of representations and real world items that are relevant to the consumer of those representations (a “natural isomorphism”) may constitute an exploitable relation, and may actually be exploited by consumer systems in some cases. If that is right, then there are indeed some representing systems for which the existence of an isomorphism – of this much more tightly constrained sort – plays a substantive role in the theory of content.

A final caveat is that these natural isomorphisms need not be exact. It does not make a map useless that one of the points on it is slightly out of place. And any map will be imperfect at some level of detail. Some of these inaccuracies may be accounted for by individuating the isomorphisms at an appropriately coarse level of grain. But there may be other cases where we need a notion of the isomorphism itself only holding approximately. That suggests that the mathematical concept of isomorphism may not be the most useful explanatory tool here. It may be more useful to treat the system of representations as a model of its target (Godfrey-Smith 2006b, Weisberg 2007). Notice that this point about replacing isomorphisms with models does not arise in the discussion above, where the consumer fixes the isomorphism, since it is facts of the matter about the evolutionary story that deliver the appropriate isomorphism. There, the theory of content works so as to deliver a precise isomorphism, to which operation of the system will have been at least approximately true. We are not thinking of the isomorphism as some pre-existing relation that the consumer makes use of. Once we do look for pre-existing exploitable relations, as suggested here, it is clear that the correspondence need not be exact to be exploitable.

(5) Conclusion

The mathematical concept of isomorphism or functional isomorphism is unsuited to playing a substantive role in Millikan’s theory of content. Instead, various natural relations adverted to by the theory (in a Normal explanation of the operation of the system) make it the case that each representation has a particular content. That picks out a particular isomorphism as suitable for giving the content of the representations. But we saw two ways in which more constrained notions do have a role to play. Firstly, where natural

relations between a system of representations and things in the world enter into an explanation of how the producer-consumer system managed to survive and reproduce, then that evolutionary story may give rise to relational proper functions. In such cases, the theory will assign contents to new representations that fall within the relational type, which is a kind of productivity. Secondly, natural isomorphisms are exploitable relations in their own right, and there are probably cases where consumer systems make use of such relations, in a way that cannot be fully explained by appeal to correlational information. These points do nothing to undermine Millikan's theory of content. Rather, they clarify the theoretical machinery needed to deliver her results.

Nicholas Shea

nicholas.shea@philosophy.ox.ac.uk

14 September 2010

Acknowledgements

Many thanks to Peter Godfrey-Smith and Ruth Millikan for very helpful comments on earlier drafts of this paper.

References

- Cover, T.M. and Thomas, J.A. 2006: Elements of Information Theory, second edition. Hoboken, NJ: John Wiley & Sons.
- Godfrey-Smith, P. 1996: Complexity and the Function of Mind in Nature. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 2006a: Mental representation, naturalism and teleosemantics. In Papineau, David and Macdonald, Graham (eds.) New Essays on Teleosemantics. Oxford: OUP.
- Godfrey-Smith, P. 2006b: The strategy of model-based science. *Biology and Philosophy*, 21,

725-40.

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

Millikan, R.G. 1989: Biosemantics. *Journal of Philosophy*, 86, 281-97.

Millikan, R.G. 1995: A bet with Peacocke. In Macdonald, Cynthia and Macdonald, Graham (eds.) *Philosophy of Psychology: debates on psychological explanation*. Oxford: Blackwell pp. 285-92.

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

Millikan, R.G. 2000b: Naturalizing Intentionality. In Edevitch, Bernard (ed.) *Philosophy of Mind, Proceedings of the Twentieth World Congress of Philosophy, Volume 9*. Philosophy Documentation Center.

Millikan, R.G. 2003: Teleological Theories of Mental Content. In Nadel, L (ed.) *The Encyclopedia of Cognitive Science*. Macmillan.

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

Millikan, R.G. 2005: *Language: A Biological Model*. Oxford: OUP.

Millikan, R.G. 2007: An Input Condition for Teleosemantics? Reply to Shea (and Godfrey-Smith). *Philosophy and Phenomenological Research*, 65, 436.

Shannon, C.E. 1949: The mathematical theory of communication. In Shannon, Claude E and Weaver, Warren (eds.) *The Mathematical Theory of Communication*. Urbana: University of Illinois Press.

Weisberg, M. 2007: Who is a modeler? *British Journal for the Philosophy of Science*, 58(2), 207-233.