OXFORD HANDBOOK IN THE PHILOSOPHY OF MIND
Brian McLaughlin, editor

BIOSEMANTICS

"Biosemantics" was the title of a paper on mental representation originally printed in The Journal of Philosophy in 1989. It contained a much abbreviated version of the work on mental representation in Language Thought and Other Biological Categories (1984). There I had presented a naturalist theory of intentional signs generally, including linguistic representations, graphs, charts and diagrams, road sign symbols, animal communications, the "chemical signals" that regulate the function of glands, and so forth. But the term "biosemantics" has usually been applied only to the theory of mental representation. Let me first characterize a more general class of theories called "teleological theories of mental content" of which biosemantics is an example. Then I will discuss the details that distinguish biosemantics from other naturalistic teleological theories.

Naturalist theories of mental representation attempt to explain, in terms designed to fit within the natural sciences, what it is about a mental representation that makes it represent something. Frequently these theories have been classified as either picture theories, causal or covariation theories, information theories, functionalist or causal role theories, or teleological theories, the assumption being that these various categories are side by side one another. But they are not. Teleological theories are specific forms of one or another, or of some combination, of the other kinds of theories. What teleological theories have in common is not any view about the nature of representational content, that is, about what makes a mental representation represent something. What they have in common is only a view about how falseness in representations is possible.

Roughly the idea is this. You tell the teleologist what you think makes some item in the head, some facet or activity of the brain, into a representation of some facet of the world, say, into a belief that it is raining. The teleologist may well agree with your theory about this. But then she will go on to point out (typically this is so) that your theory is really, at root, a story only about what it is for a mental state or activity to represent truly or correctly. You need to add a story about what a representation is like that represents falsely. And she will claim that this is very easy to do. Assume that the brain was designed, by evolution or learning, to make or to learn to make representations of the kind you have described. But what it was designed to do will not always be what it in fact does. Difficult environmental circumstances, even circumstances that merely fail to be ideal, often cause temporary failures for biological systems. System designed to produce representations will sometimes fail to produce them correctly. Sometimes they will produce items that behave in the mind/brain as though they represented something, but that in fact do not represent anything. These are false representations. They are "false" in the dictionary sense of "not genuine or real," "resembling but not accurately or properly designated as such," the sense in which false faces and false fronts are false. That is, teleological theories are best

understood as denying that there IS any state of affairs or occurrence being represented when one thinks falsely or that there IS any object at all that is being represented when one thinks emptily, say, when seeming to think about "phlogiston" or "the ether."

Similarly, there IS no object, not even an inner one, being seen when one has an hallucination. Mistaken representations, rather than representing peculiar objects, things called "intentional contents," are just representations that are failing to represent. False representations are representations yet fail to represent in the same way that something can be a can opener but be too dull, hence fail to open cans or --and this is a better analogy-- something can be a coffee maker yet fail to make coffee because the right ingredients were not put in or it was not turned on. They are "representations" only in the sense that the biological function of the cognitive systems that produced them was to make representations. Thus falsehood is explained by the simple fact that biological purposes often go unfulfilled, and the ghostly realm of intensions, reified meanings, nonexistent objects of thought, intentional objects, propositional contents, and so forth, is cleanly swept away. Strictly speaking, you can't represent something that doesn't exist.

But much work still remains. The teleologist must give an account of biological functions and of functions derived from learning that can support the view that the mature brain has the production of representations, representations of quite specific kinds, as one of its functions. The account I have given ultimately rests these functions on a variety of different kinds of selection, the most fundamental being natural selection. That story can be found in (Millikan 1984 chapters 1 and 2; 1993 chapters 1 and 2; 2002). In (Millikan 2004 chapter 1, n.2) it is defended against the claim that human intelligence could have been a genetic accident rather than a trait selected for.

The other main work that remains is to explain what it is for some facet or activity of the brain to be a representation of some affair in the world. What teleological theories do not have in common is any agreed on description of what representing is. They do not agree 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. Teleological theories, just as such, are not theories of mental content. Failure to grasp this last point has lead many to take a dismissive attitude toward teleological theories. How, they ask, could the question whether my current thought is the thought that grass is green rather than the thought that Aardvarks bark be a matter that is settled in part by reference to evolutionary history or to my past learning history? But a teleological theory, just as such, makes no attempt to explain what makes your thought be a representation that grass is green or that aardvarks bark. A prior theory of what (correct) representation is is needed for that. To the shell that is

"teleosemantics," then, it is necessary to add a description of what successful representing, actual representing, is like. The main work of this article is to explain what representing is according to "biosemantics."

Above I suggested that the teleologist can take any naturalistic theory of representation at all and turn it into a teleological theory of content. But there is a catch. If the teleologist anchors the notion of function in selection, the theory of representation adopted must allow us to explain how producing inner representations might sometimes benefit an organism. Otherwise it will be a mystery why any organism would contain systems designed by selection to make representations. Surely such a requirement is reasonable, but on careful consideration it turns out to be quite confining.

Many naturalistic accounts describe the relation of a representation to what it represents as a simple diadic relation. This is true, for example, of classical causal or covariational theories, of classical informational theories and of classical picture theories. C.S.S. Peirce, on the other hand, claimed that the representing relation is essentially triadic, involving first the representation (a "sign"), second something represented, and third an "interpretant." If producing inner representations benefits an organism, presumably this will be because the organism uses them in some way. There must be a part of the organism, or some activity of the organism, that understands or interprets these representations. Peirce spoke of the interpretant of a sign as being another sign, but taking this at face value would produce a regress. The interpreter of an inner sign cannot be supposed merely to translate the sign into another inner sign which is again translated, and so forth. "Interpreting" a sign must ultimately consist in some independently useful activity.

According to biosemantics there are several different kinds of processes that use representations. Most theories of representation deal with descriptive representations only --with representations that are designed to represent facts. But directive representations are certainly equally important --representations that tell what to do. And the most primitive and fundamental kind of representation, I believe, faces both ways at once, saying at the same time what the case is and what to do about it. For example, the dance of the honey bee tells where the nectar is and at the same time where the watching bees are to go. I call this last a "pushmi-pullyu" representation after Hugh Lofting's charming two headed creature by that name (Millikan 1996). Pushmi-pullyu representations simultaneously describe and direct. Better, since the term "representation" suggests to many people something more fancy than the simplest examples I have in mind, in particular, it may suggests symbolic forms that are "calculated over," I prefer to call the inner signs that I describe "intentional icons" (Millikan 1984) -- I will explain why in a minute. Representations that are calculated over --that participate in inference processes-- are just one very fancy kind of

intentional icon. Below are diagrams of each of the three basic kinds of intentional icons. I will discuss them in turn.

[Figures about here: 1. Descriptive; 2. Directive; 3. Pushmi-Pullyu]

In each of the diagrams there is a producer and an interpreter or "consumer." These have been designed by natural selection or by learning to cooperate with one another. Perhaps each resides in a separate organism, for example, one is part of a dancing bee, the other a part of a sister watching bee. Or perhaps they correspond merely to two different functions performed within the same brain. What the producer does helps the consumer to perform functions that loop back to make both the producer and the consumer more likely to survive or to maintain their current settings (selection through learning) or to proliferate. The presence of each is part of the normal mechanism by which the other helps itself to survive or proliferate, and this cooperation is no accident but the result of past selection or learning that operated on both together.

In each diagram the producer produces a sign that will be true or satisfied only if it maps onto some affair or affairs (the plural is for pushmi-pullyu icons) in the world in accordance with certain "semantic" rules. These are rules of correspondence between signs and world affairs that have been instantiated in the past when the consumer and producer or their ancestors have succeeded in performing their cooperative function(s). Consider the bee dance. Suppose that the activities of the dancing and the watching bees fulfill their cooperative function of directing the watching bee to a supply of nectar in the normal way, that is, through the characteristic mechanisms that have accounted in the past for success and subsequent selection of the dance-making and dance-using bee apparatuses. In such cases, the dancing bee performs a dance that maps by a certain rule, first, onto the location of nectar. Second, it maps onto the direction and duration of the flight produced in watching bees. In each case, had certain variables in the dance been different --the angle, the speed-- the nectar would have needed to be in a different place and the watching bees would have needed to go in a different direction or to a different distance for the dance to serve its purpose in the normal way. The semantic rules of beemese are defined by the way in which the set of possible well-formed bee dances maps onto the set of possible nectar locations to determine an isomorphism between these domains, an isomorphism that holds when the bee dance works properly through historically normal mechanisms. Similarly, another isomorphism holds between the set of possible dances and the set of destinations that will result if watching bees find nectar through historically normal mechanisms. The value of a system of representation lies in its productivity, which always depends on an isomorphism between the domain of the signs in the system and the domain of their signifieds. Mappings or isomorphisms are an essential feature of each kind of sign system

diagramed. The signs are like abstract pictures. This is why (following Peirce) I call them "icons." Intentional icons are produced by systems designed to make abstract pictures or icons that will map coincident with predetermined mapping rules to which their consumers are adjusted. When the systems in which they are embedded are functioning normally, they will picture in accordance with these rules, and they are then said to be "true" or to be "satisfied."

I had best add to this that no limit need be placed on the complexity of the semantic mapping functions that might map intentional icons onto the states of affairs they represent. Isomorphisms can be defined by functions that are as bizarre, as gruelike, as you please. A bizarrely coded secret message from a CIA agent can be as much an "icon" or "picture" that maps onto a certain world affair in accordance with a definite semantic mapping function as any bee dance, sentence or diagram. Intentional icons must be things apt for use by icon users, but icon users can be very idiosyncratic in their habits. For example, if mental representations are systems of brain happenings or brain states that map onto represented world affairs, no a priori limitation on the kinds of brain happenings or states involved or on the complexity of the mappings employed is implied. Every representation is in some kind of code. The complexity of the code is irrelevant. On the other hand, any intentional icons in the brain would of course have to come with inner interpreters that knew how to read them, that is, interpreters that could be guided by them reliably to fulfill their own functions. Simple codes relying on only a few principles, if they were also highly productive, tapping into rich natural isomorphisms between the domains of the signs and the signifieds, would seem much the most likely to be preferred by natural selection.

The semantic rule associated with a descriptive intentional icon determines a condition or state of affairs that must obtain if the consumer is to perform its tasks, whatever they may be, in the normal way. The consumer varies its activities systematically according to variations in the icons presented to it. The result is that the consumer's activity conforms or is is adapted to the condition or state of affairs represented by the icon so that it can perform its functions properly given that condition. Descriptive intentional icons are designed to stand in for world affairs, typically affairs outside the organism or organisms involved, and to vary according to these world affairs, controlling internal or external behavior as needed to adapt to these affairs. Thus the bee dance is, in part, a descriptive intentional icon because if the watching bees are to achieve their function of finding nectar by reacting to the dance in the normal way, it needs to correspond by a certain rule to a fact about nectar location. (We need not assume it takes any thought on a watching bee's part to react appropriately, so we need not assume that the dance is interpreted by translating it into another sign.) Similarly, consider my belief that there is yogurt in the refrigerator. Which are my belief-using systems? Some, at least, are the systems that make practical inferences

and turn the conclusions into practical action. My belief that there is yogurt in the refrigerator can help these systems to perform their function of guiding me into activities helping to fulfill my desires and plans (to eat yogurt, to make breakfast for a guest, etc.) in a normal way (not by serendipitous accident) only if there is yogurt in the refrigerator.

Notice that the content of the descriptive icon is not determined by reference to the kind of tasks that it (with the help of its consumer(s)) normally performs. The icon may be called on to promote a whole series of functions, but each of these will be fulfilled normally only if the icon is "true." If the consumer works in part by making inferences, the content is not determined by any particular set of inferences its consumer is disposed to make. It is determined by the fact that whatever inferences its consumer makes when functioning properly, the result will be (nonaccidental) true belief or successful action only if it actually represents according to a certain correspondence rule. This means that its content is not determined by its conceptual role.

Note also that in the case of descriptive icons the producer's job is primarily to make an icon that corresponds by the right rule to a state of affairs. If the producer succeeds in this task, whether through the normal mechanisms of icon production or whether by freakish accident, the intentional icon is still true. Truth does not rest on whether the means of production was normal. For example, a belief can be true without being knowledge. Taking a classic case first introduced by Dretske (1986), consider the magentosome of a northern hemispheric ocean bacterium that normally works by being pulled toward magnetic north hence toward geomagnetic north hence away from aerated surface water which is lethal to these bacteria. The magnetosome's pointing in a certain direction is an intentional icon of the direction of lesser oxygen because that is what it needs to correspond to for the bacterium's resultant motion to serve its proper function. If the magnetosome points in the direction of lesser oxygen because, serendipitously, a bar magnet under it is pointed in the right direction, then it tells the truth by accident rather than through normal mechanisms, but it still tells the truth. On the other hand, if a bar magnet attracts the bacterium to its destruction by pointing it towards oxygen, what the magnetosome says is false despite its being produced through a normal mechanism. Attributions of truth and falsity do not rest on whether descriptive icons are produced by normal mechanisms, but only on normal mechanisms associated with their use. (An icon's function cannot be to have been produced by something! Functions are effects, not causes.) Of course Dretske is right that the magnetosome that directs the bacterium in the wrong direction because a bar magnet is held over it is not broken or malfunctioning. In that sense it is functioning "perfectly properly" (Dretske 1991). But it doesn't follow that it is succeeding in performing all of its functions, anymore than a perfectly functional

coffee maker is performing all its functions when it is turned on but no coffee has been put in it. Very often things fail to perform their functions, not because they are damaged, but because the conditions they are in are not their normal operating conditions.

Directive intentional icons that correspond to world affairs by their normal correspondence rules are usually said to be "satisfied" rather than "true." The semantic rule associated with the icon determines a condition or state of affairs that its consumers are to produce. The consumer's job is to bring it about that a state of affairs corresponds by rule to the icon. Its job is to "obey" the producer's "orders." (Equally, of course, it is the job of the producer to give orders that will benefit both it and the consumer.) Thus the bee dance is, in part, a directive intentional icon because it will correspond by a certain rule to the direction of flight of the watching bees if their dance-interpreting apparatuses success in serving their functions normally. Suppose that you have a desire to eat yogurt. The systems designed to be moved by your desires --the "consumers" of your desires-- are the systems that make practical inferences eventually turning the conclusions into practical action. If your desire to eat yogurt affects these systems as designed, it will guide them to effect the fulfillment of your desire to eat yogurt. The chief function of a desire is to get itself fulfilled. This is not to say that, on the average, desires do get themselves fulfilled. On average they generally get eaten up by bigger opposing desires first, or perhaps no means are known to fulfill them. It is very common for a trait or capacity to have been selected because it sometimes performs a useful function, occasional performance being better than none. The point is that people would not have the capacity mentally to represent various states of affairs as desirable unless these desires were sometimes fulfilled. The capacity would otherwise be useless.

Thus the cooperation between producer and consumer in the production and use of intentional icons can be accomplished in any of three basic ways. It may be that the producer is the one primarily responsible for making the icon correspond to a certain state of affairs, it may be that the consumer is the one primarily responsible, and in the case of pushmi-pullyu icons, it is the responsibility of the producer to make the icon correspond to one kind of affair and the consumer's job to make it correspond to another. In each of these basic cases, granted the cooperation between producer and consumer comes about through the normal causal mechanisms, the intentional icon will also be a "local natural sign" carrying "local natural information" about the affair or affairs to which it corresponds (Millikan 2004, chapters 3-4). Local natural signs are, in part, abstract pictures of what they represent. Although there is not room to unpack the notion of local natural information here, I mention this because it follows that when intentional systems are functioning normally in accordance with normal explanations, intentional icons represent both by being pictures of what they

represent and by carrying natural information as to what they represent. For example, a bee dance is often a local natural sign both of where there is nectar and of where the watching bees will go. In the fundamental sense, actually representing involves both picturing and carrying natural information. As such it is not a matter determined by a history of selection. Representing intentionally, however, is a matter of having a certain kind of history. Also, attributions of truth and falsity and of satisfaction or unsatisfaction make sense only by reference to function, hence by reference to a history of selection. The terms "true" versus "false" and "satisfied" versus "unsatisfied" do not apply to natural signs, the most basic kind of representations.

A common question raised about informational theories of representation concerns how mental representations can carry "information about a distant causal antecedent ... without carrying information about the more proximal members of the causal chain ... through which this information ... is communicated," for it seems that such representation "...skips..over (or 'sees through') the intermediate links in the causal chain in order to represent ... its more distant causal antecedents" (Dretske 1981, p. 158). Similarly, there is a worry about how abstract representations are possible, ones that carry only the information, say, that an object is triangular and not also that it is isosceles or equilateral. These problems do not arise for the theory of intentional icons. Information carried by normally operating intentional icons is a form of natural information. It does not follow that all of the natural information carried by an intentional icon is carried intentionally. The information that a natural sign carries intentionally is only the information it is its function to carry, the information that its cooperative interpreters know how to use. This information may be very abstract, and it may be about very distal affairs. If its consumers are so designed that they can use only the information that something is triangular, then that is all the information that the icon carries intentionally. If they are so designed as to use only the information that a predator is near, then the intentional icon will not intentionally carry information about any more proximal affairs, such as patterns on the retina.

Similarly, not every stimulus that an organism discriminates on the way to producing an intentional icon is represented intentionally. Nor must an organism be capable of infallibly discriminating the distal objects, properties or kinds that it intentionally represents from all others that are similar. It needs only a fallible capacity to recognize some natural signs or other of these things under some local conditions. Possibly it even gets things wrong a large part of the time because a large part of the time supporting conditions on which its mechanisms of icon production rest are absent. Similarly, the rabbit's danger thump may be elicited more commonly when rabbit danger is absent than present. (What matters is the converse --that when danger is present it should usually be elicited.)

The basic theory of representing on which biosemantics rests is a picture theory and an informational theory, but equally it is a functionalist theory. The basic idea is that what makes something into a representation, for example, into a mental representation, is not, of course, what it is made of, but what functions it performs and/or how it performs these functions. Items that function in certain ways are representing, and if they have been designed to function in these ways, they are representing "intentionally." They are representing, that is, in accordance with natural purposes and such that they can be said to be true or false, satisfied or unsatisfied. (The intimate connection of functions that have been selected for with purposes is argued in (Millikan 1984; 2004 chapter 1)).

According to biosemantics, basic mental representations always represent complete states of affairs. Mental terms are not endowed with meaning first and then used to build mental sentences. "What makes the mental term 'horse' stand for horses?" is not the place to begin. Parts and aspects of complete intentional icons represent parts and aspects of complete states of affairs only as abstracted from completed icons. This general point is most apparent, perhaps, with the simplest and most common cases of pushmi-pullyu intentional icons, exemplified by nearly all animal's signals and also by the ubiquitous chemical signals running in the blood stream that direct responses from various organs and cells. These signals, taken along with their times of occurrence and sometimes with their places of occurrence, are intentional icons because variations in the times and places of occurrence correspond to variations in the times and places of the complete affairs represented. For example, the time and place of the mother hen's food call to her chicks descriptively shows the time and place of food and directively tells the time and place her chicks are to come. But it is evident that there is no sense in which a particular time stands for itself or a particular place for itself outside of the context of some such signal. Similarly, in telling what direction the bacterium is to go the direction in which the its magnetosome points stands for itself, but there is no prior sense in which a direction stands for itself.

The biosemantic account also implies that there are not and could not be intentional icons that lacked attitude. Rejected is the Fregean idea that first a proposition is represented, then an intentional attitude added. Intentional icons always have, as such, functions and their functions automatically create attitudes. Hypothetical thinking, for example, or just thinking of possibilities, is an extremely sophisticated activity, and one that is only possible for a creature that sometimes uses the results in the production of ordinary descriptive and directive representations. It is because thoughts "of possibilities" have the function of sometimes turning into more basic kinds of representations that they exit as representations at all. Similarly, desires are intentional icons only because they are designed to turn, under certain conditions, into full-blown intentions, whose functions are to effect their own fulfillment more directly. There would be no

benefit in the capacity to have desires if desires did not sometimes travel the whole route through intention into action.

Intentional icons represent complete states of affairs. This implies that they represent not only properties but also the things that have those properties. When produced normally, intentional icons also carry natural information corresponding to what they represent. This implies that there can be natural information as to what things have what properties --including what individuals have what properties. Contrast Dretske's description of the natural information carried by signals. As Dretske describes the matter, although a signal can carry the information that an individual x is F, there is no part or aspect of the signal that carries the information that it is x that is F. For example, the gas gauge on your car may carry the information that your gas tank is half full, but no aspect of the signal indicates which tank is half full (Dretske 1981, 1991). You have to know that independently. The gas gauge reading does not represent its subject, nor could it on Dretske's theory of natural information, because there are no natural laws that pertain to any individuals just as such. A necessary and central feature of the theory of local natural information (Millikan 2004) is that it explains how a natural sign can signify which individual it carries information about. The result is an explanation of how intentional representations of individuals are possible, something for which, to my knowledge, no other naturalized theory of intentionality accounts.

A common question raised about the program of biosemantics is how representations such as human beliefs and desires, which in numerous instances are entirely unique to the individual who has them, can have acquired functions through a process of selection. In outline the answer is straightforward. Compare the design of a camera. A camera is not designed to take any particular picture, but to vary the picture it produces depending on the scene in front. If a particular pattern is in front, it will function properly if it produces a likeness of that pattern. Similarly, of course, one's eyes are not designed to see any particular object but, roughly speaking, to see whatever object lies in front of them. If a particular person is in front of your open eyes, it is a function of your eyes to help produce an accurate perception of that person. Similarly, the function of an adding machine is not to give any particular answer, but to give the sum of the numbers put into it. Natural selection has designed cognitive systems not to turn out particular products, say, particular beliefs and desires, but to turn out quite different beliefs and desires depending on environmental circumstance. In the case of human beliefs and desires, however, the matter is considerably more complicated than with the camera. In order to turn out beliefs that will vary depending on states of affairs in the environment and in order to tune the systems that use these beliefs during practical and theoretical deliberation and in the production of useful action, humans must first develop adequate empirical

concepts. Indeed, to complete the biosemantic program, a rather long story needs to be told about conceptual development. We must explain how the producers and consumers of beliefs and desires can learn or be tuned to employ empirical concepts cooperatively without actually practicing together through the production of concrete actions. We must explain how their representation production and use dispositions can be tailored in advance to fit one another. That story is told in (Millikan 1984, chapters 15-19; 2000, especially chapter 7; 2004, part IV, especially chapter 19).

Many critical questions about the biosemantic theory first presented in 1984 and 1989 have come and gone, but there are three that have been especially tenacious. I will say a word about each.

What an intentional icon represents descriptively is an affair to which it must correspond if it's consumers are to perform their functions by normal mechanisms. They will perform their functions through normal mechanisms only if external conditions are such as to allow these mechanisms to operate properly. Taking for her example the female-hoverfly detector in a male hoverfly's visual system, Karen Neander (1995) has objected that among the external conditions needed for the detector's consumers to perform all their functions are that the female is fertile and that she won't be eaten before she reproduces, hence that on the biosemantic theory these facts about the female must be part of what is represented by the detector in the male's visual system. What this overlooks, however, is that an intentional icon must also have a producer and that it must be a function of the producer to make an icon that corresponds to the condition it represents. If the producer has a function there must be a normal mechanism by which it performs that function. This, however, would require the male hoverfly's visual systems to be sensitive to natural signs of fertility in female hoverflies and of liability not to be eaten. But on no theory of information, certainly not on the theory of local natural information, does the male hoverfly use or even encounter any such natural information (Millikan 2004, Chapter 6).

A second question concerns the possibility of biological systems whose jobs are to produce false representations. For example, people who are overconfident may be more successful at performing certain tasks than people who evaluate their skills correctly. Notice first, however, that it will not be the falseness per se, but only the high confidence that contributes to success. If one is completely and perfectly competent at a task, there certainly will be no gain in believing one is not! Notice second that there are many biological systems that ride piggyback on systems developed earlier for quite different purposes. If there actually were systems whose jobs were to distort certain beliefs they would have to ride on more general systems whose basic jobs were to produce true beliefs. Otherwise there would be no semantic rules in place according to which the distorted beliefs would be false. The various concepts out of which any belief are

formed are designed to serve purposes in arbitrary belief contexts. The systems responsible for concept development tune these concepts and the systems that normally use them for general purpose use, not for any one specific use such as increasing one's confidence. The semantics of mental representation is productive. That is what the "picturing" or "mapping" guarantees.

Third, consider Pietroski's tale about the kimus and the snorfs (1992). The kimus are attracted by the red sunrise glowing over their local mountain so that they climb up it each morning. Thus they conveniently avoid their chief predators, the snorfs, who pass by each morning below. Moreover, this is how the attraction to red light got selected for in kimus. Those not attracted by red light got eaten. On the biosemantic view, Pietroski claims, "...kimus climb the hill because they believe the hill is snorf-less," and when they approach red things that are not snorfs "they are acting on the belief that the area in question is snorf-free" [emphases Pietroski's]. Given that kimus "can't reliably discriminate snorfs from non-snorfs," it is implausible, he claims, that the kimus have any beliefs about snorfs. In summary, on the biosemantic account "A system can have the belief that P is instantiated without having any systematic ability to tell whether P is instantiated (in a given region at a given time). Indeed, instantiations of P can be completely irrelevant to the system's tokening of the belief that P is instantiated. The corresponding intentional explanations of such a system's behavior will...be very implausible."

What Pietroski describes in the kimus seems to be a simple tropism. They are attracted to red the way a moth is attracted to light. Apparently they have neural pushmi-pullyu intentional icons or signals that tell where the snorfs are fewer and hence where to go. Of course these icons are not at all like beliefs. Beliefs are formed only after the acquisition of concepts, which generally rest on multiple ways of recognizing. Further, the functions of beliefs involve participation in inference. (For a discussion of empirical concepts and of inference, see (Millikan 2000)). Pietroski seems to assume that an "intentional explanation" of an animal's behavior must not only be belief-desire explanation but must also be a straightforward causal explanation. Why intentional explanations are causal but not straightforward is explained in (Millikan 1993). Moreover, full intentional explanations do not begin with the presence of intentional representations, but explain also how the intentional representations get formed. The red light is definitely involved in an "intentional explanation" of how the kimus avoid snorfs.

Finally, there are no distal objects or stimuli that any organism has the capacity to discriminate under all conditions. All successful discrimination of distal affairs depends on merely local natural information. Local natural information rests on correlations that are not perfect but that are not accidental either, for they must persist throughout a spatial-temporal region for a reason (Millikan 2004 chapters 3-4). On this analysis the kimus do get local natural

information each morning about the direction of fewer snorfs. Similarly, although there is no causal connection, the correslation between magnetic north and lesser oxygen used by the anaerobic bacteria persists in the northern hemisphere for a reason. It carries local natural information about the location of lesser oxygen. The intentionality that characterizes pushmi-pullyu icons responsible for simple tropisms of this kind is the limiting case of intentionality. It is intentionality in the way zero is a number. If your theory doesn't count in these cases you will find that it fails to account for any of the obvious cases either.

Dretske F 1986 "Misrepresentation" In Radu Bogdan (ed.) Belief: Form, Content, and Function (New York: Oxford), 17-36.

Dretske F 1991 Explaining Behavior (Cambridge MA: MIT Press)

Millikan RG 1984 Language, Thought and Other Biiological Categories (Cambridge MA: MIT press).

Millikan RG 1989 "Biosemantics" The Journal of Philosophy 86.6:281-297.

Millikan RG 1993 "Explanation in Biopsychology", in J. Heil and A. Mele, eds.,Mental Causation, Oxford University Press 1993 pp. 211-232. Reprinted in RG Millikan, White Queen Psychology and Other Essays for Alice (Cambridge MA: MIT Press) 1993. Reprinted in C. Macdonald and G. Macdonald, Philosophy of Psychology: Debates on Psychological Explanation, Oxford University Press 1995 (Parts VII-X).

Millikan RG 2002 "Biofunctions: Two Paradigms" in R. Cummins, A. Ariew and M. Perlman, eds., Functions: New Readings in the Philosophy of Psychology and Biology (Oxford: Oxford University Press) 113-143.

Millikan RG 2000 On Clear and Confused Ideas (Cambridge: Cambridge University Press)

Millikan RG 2004 Varieties of Reference; The Jean Nicod Lectures 2002 (Cambridge MA: MIT press)

Neander K 1995 "Misrepresenting and Malfunctioning" Philosophical Studies 79: 109-41

Pietroski PM 1992 "Intentionality and Teleological Error. The Pacific Philosophical Quarterly 73, 267-82.