

CHAPTER SIX INTENTIONALITY

The teleologist needs a base theory of the representing relation on which to build his description of intentional representation. Fred Dretske (1986, 1988, 1995) placed his theory of natural signs and the natural information they carry at the base of his teleological theory (Chapter Three above). Some items, Dretske claimed, have the "function" of carrying natural information, and when they do this they come to represent intentionally. They become intentional representations, representations that can be false (Dretske 1986, 1991). Dretske's "functions" are what I have been calling "purposes," at least close enough for my purposes.¹ If we take it that a just plain representation, one that is not intentional in Brentano's sense, is just a bearer of natural information, then Dretske's theory of intentional representation is a very neat example of a completed teleosemantic theory of the sort described in Chapter Five.

Now I have claimed that there exists practically no natural information of the kind Dretske describes. If intentional representations were to be made by organisms out of natural information, it would have to be out of a more user-friendly kind of natural information. Let me introduce my suggestions on intentional representations by asking whether Dretske's program can be carried out using local natural information as a base, and if not why not. In particular, let us ask whether locally recurrent information might be used in this way. Might intentional representations be just things designed to carry local natural information? Might they be things produced by systems whose purpose is to produce locally recurrent natural signs?

The proposal on the table here puts no definite requirements on the probability that the systems designed to produce natural signs actually succeed in doing so. They need only to succeed often enough to offset their own production costs (energy and resources used) plus any negative effects resulting directly from failures. The rabbit, for example, can afford to be mistaken that a fox is near many times, so long as it also takes a fox to be near whenever one actually is. Having an intentional natural sign producer that is extremely fallible may be much better than having none at all. It is only necessary that the costs of being wrong are less than the costs of totally unrelieved ignorance. Similarly, an organism's capacity intentionally to represent certain kinds of states of affairs is not dependent on its ability reliably to discriminate those states of affairs. Putting things bluntly, no vestiges of verificationism remain under the proposed analysis.

But the teleologist who uses natural information at the base of his theory must explain at least why producing natural signs might sometimes be useful to the organism that produces them. How might an organism benefit from the production of natural signs? Why might a system within it have been selected for that job? It will pay to be very careful here. We need to distinguish a useful effect that also happens to result in the production of natural signs as a side effect from the production of the signs being itself a useful effect. For example, the production by the body of calluses where the wear is has the useful effect of protecting the skin from further damage, and where the calluses are is a natural sign of where the wear has been. But the fact that the calluses are a natural sign of where the wear has been is not, of itself, of any use to the body. The disposition to produce calluses was not selected for its effect of indicating where the wear has been. If I dig ditches to channel the water across my property when it rains, then if my purpose is fulfilled, the ditches will be natural signs of where the water will flow when it rains.

But my purpose was not to produce natural signs of where the water will flow when it rains. Perhaps geese have been designed by natural selection to react to frosty nights by flying south, the useful effect being that they fly south just before winter arrives. If this design works right, a side effect will be that frosty nights are a natural sign that geese will soon be flying south, and geese flying south will be a natural sign that winter will soon arrive. Both of these natural signs are the result of the operation of natural selection on geese, but the production of neither sign was selected for. Neither the calluses nor the ditches nor the frosty nights nor the geese flying south will be intentional signs, then, under the proposed analysis.

On the other hand, consider the distinctive clucking sound that a mother hen makes when she finds food. The clucking is a locally recurring natural sign that the hen has found food. Further, the mother hen's chicks respond to her calls by running to her and thus finding the food. Indeed, her disposition to call has been selected for having this effect on her chicks. The call is not merely a natural sign of food that happens to be used by the chicks. It has been purposefully produced in order to serve the chicks as a sign of food. On the proposed theory, her call is an intentional sign.

The lesson seems to be that if an intentional sign is the same as a purposefully produced natural sign, it must be designed to function as a sign for some kind of interpreter. The teleologist who claims that intentional representations are produced by systems whose function is to produce natural signs must also claim that intentional representations are designed to have the effect of correctly representing something to some other organism or interpreting system (combining the first and second possibilities for a teleological theory mentioned in Chapter Five). This accords with the governing idea, suggested in Chapter Three, that a useful notion of natural sign should define a category useful to the natural epistemologist. Signs are things apt for use by sign users. Intentional signs are signs purposefully produced for use by sign users. It follows that besides a theory of what natural signs are, we also need a theory of what sign use is.

A refinement is now needed for this proposed description of intentional signs. Suppose that I purposefully shoo away flies with a flick of my hand. My hand flicks cause the flies to depart because they serve as a natural sign of danger to the flies, and I intend this. Are my flicks then intentional signs? We might solve this problem an easy way by saying that they may be intentional signs, but they are not cooperative intentional signs, and that when people talk of intentional signs they usually have cooperative intentional signs in mind. Cooperative intentional signs are produced by systems designed to make natural signs for use by cooperating interpreting systems. That is, the sign-maker system and the sign using system must have evolved or been designed to function symbiotically. Cooperative intentional sign makers must be designed to cooperate with interpreting systems which have been designed, in turn, to cooperate with them. A cooperative intentional sign will always stand midway between two systems that have been designed to cooperate with one another. Perhaps one system is a clucking hen, the other a listening chick, or perhaps one is certain systems in the brain that make signs for other systems in the brain to use. However it goes, that the sign should be a natural sign must be of concern to both systems. What helps proliferate the chicks also helps to proliferate the hen, and what helps to proliferate one half of the brain generally helps to proliferate the other. By "intentional sign" I will now always mean cooperative intentional sign unless I indicate otherwise.²

But is it true that the intentional representations in one's head are, in general, recurrent

natural signs of the affairs they are about? A locally recurrent natural sign must fall within some natural domain within which the signs coincide with their signifieds all for the same reason (Chapter Three). But the perceptual systems of an animal may rely on numerous alternative and quite independent natural cues in constructing even elementary perceptual representations. For example, ocular disparity (convergence), partial occlusion of one object by another, tautness of the muscles that focus the eyes (accommodation), atmospheric haze, are all used in depth perception, jointly or alternatively. It is quite possible, of course, that any one of these natural signs of depth, taken by itself, bears a strong enough correlation to depth to be of use by itself to an organism, hence that the set of inner representations of depth derived by each of these various methods is, in turn, a separate set of locally recurring natural signs. But suppose we generalize the principle involved. Consider how many quite independent ways you may have of recognizing one of your parents, or your spouse, or your children as manifested to you in perception. You know them by the look of their faces or of various of their other body parts from a hundred angles. You know them by their postures, their mannerisms, their clothes and other belongings, their voices, their handwriting, by their characteristic linguistic expressions, by traces of their habits (who leaves the lid off the peanut butter in the afternoon?) and so forth. To each of these various kinds of natural signs, you make the same cognitive response, forming another representation token representing some fact about this same person. Numerous different kinds of natural signs of the same thing have all been converted into a common intentional currency, but each by a different method.³

Further, it seems likely that advanced perceptual systems, like connectionist perceptrons, do not rely mainly on single recurrent high quality natural signs. Often they are moved by co-occurrence of a diversity of very weak symptoms of a thing, to form a representation of its presence. Further, prior knowledge of present state conditions and of natural laws or uniformities is very often involved when a human takes one thing as a sign of another. Recall the beeping briefcase from Chapter Four. Intentional sign production, when taken entirely generally, seems quite different from recurrent natural sign production as defined in chapter Three. In the case of humans at least, a multiplicity of quite different mechanisms, operating in accordance with a wide variety of different principles, is what leads to coincidence of intentional signs with what they intentionally signify.

But perhaps we don't have to be so straight laced. The second requirement for locally recurrent natural signs was that there must be a reason why the mechanism producing coincidence between sign and signified continues to exist or to repeat itself throughout the domain of the sign. That the correlation continues must not be an accident. Take the limits of the sign domains for perceptions and beliefs to fall inside the organisms that produce them. Surely it might be said that the correlations between intentional signs and what they intentionally signify extends throughout that domain "for a single univocal reason."⁴ The convergence is no accident. There are systems designed by selection processes, processes that take place on a variety of levels, that see to this correlation.⁵ Let us relax the criteria for locally recurrent natural signs then so as to let these inner representations in.

So far I have been able to defend a neo-Dretsian program that would explain intentional representations as purposefully produced local natural signs. But there is one problem for such a theory that I think cannot be solved.

If the intentional representation producers have as their purpose to produce natural signs it must be because they will be aided by systems that use these signs to guide them in some kind of activity that is productive for the sign-making organism. Intentional signs must have ways of earning their keep. How will these signs be used? What will constitute their being correctly "interpreted"? In the obvious case, they will be used to guide their consumers in activities that succeed only by taking account of or conforming to the affairs that are signified by the natural signs. That is, these activities, whatever they are, will succeed by normal mechanisms (Chapter Five) only because the effect of the signs is to adapt these activities to the existence of the signified affairs. Variations in the world must correspond to variations in the sign which produce adaptive variations in the activities of the sign's interpreters or consumers. Given the way the consumers' activities are designed to vary with the sign, then, there will be some determinate semantic mapping function⁶ by which the sign must correspond to the world if its consumers are to perform their functions normally.

The difficulty is, that that is all that the sign's consumers need in order to perform normally. The consumers can do their jobs perfectly so long as the signs they consume correspond to world affairs by the required mapping function. It doesn't matter to them how the signs they use were produced, so long as they map onto world affairs the right way. Further, if we focus clearly on the function of the sign producers, carefully distinguishing their function from the normal mechanisms by which they fulfill this function, we see that their function is only to produce for their consumers what the consumers need. Their function is only to produce representations that correspond to world affairs by a certain mapping function. Their purpose or function is not to achieve this in any particular way. Thus they might sometimes fulfill their function by accident, as the eye-blink might keep sand out of the eye by accident when triggered by a sudden shadow. But it is the way the signs were produced that determines whether or not they are natural signs. So it is not a purpose of the intentional sign producers to produce natural signs. When they perform their functions by their normal mechanisms they produce natural signs. When they perform their functions by accident they produce only true intentional signs. Putting things intuitively, these representations are true but don't constitute knowledge.⁷ But the purposes of both the producer and the consumer may be fulfilled anyway.

What then is the upshot for a teleological theory of intentional representation? The need for dialectic is over, I think. Let me just lay down my position.

Most theories of representation deal with descriptive representations only --with representations that purport to represent facts. But directive representations are certainly equally important, as well as a third kind that I label "pushmi-pullyus"⁸ because they both describe and direct. Figures 1, 2 and 3 diagram these three kinds of intentional representations. In every case, when the production and use of these representations proceeds by normal mechanisms, they are natural signs. (That the directive representations are natural signs is evident if you keep in mind that natural signs can be signs of future affairs as well as past affairs, and that local signs can be flanked by things indistinguishable that are not signs.) The theory might thus be described as a sort of informational theory or natural sign theory exemplifying the third possibility for a teleological theory of intentional representation described in Chapter Five. Or it might be taken as a sort of functionalist theory, combining possibilities one and two from Chapter Five. In the peculiar case of human belief and desire, part of the functional role concerns use of these

representations in inference, prior to their eventual effects on action. That is part of how they are "consumed" or "interpreted." The theory might also be taken as a sort of picture theory, because intentional representations, like recurrent natural signs, necessarily come in systems involving a domain of possible signs running isomorphic to a domain of possible representeds (Chapter Four). In the diagrams, causation runs from left to right, and the single lines indicate mapping relations or isomorphisms. Isomorphisms are logical, not causal relations, of course. Where only single lines are drawn there may be no significant causal relations. It will be seen then that the theory of descriptive representations is not a causal theory of intentional representation. This is because recurrent natural signs are not always causally related to the affairs they signify (as was explained at the end of Chapter Three, using the magnetosome as an example).

In each of the diagrams there is a producer and a consumer. These will have been designed to cooperate with one another. Perhaps each is a separate organism, usually conspecifics. Or perhaps they are two parts or aspects of one organism. What the consumer does helps the producer, and what the producer does helps the consumer, and this is no accident but the result of some kind of selection or learning that has operated on both together. The presence of each is part of the normal mechanism by which the other fulfills its functions. In each diagram the producer produces a sign that will be true or satisfied only if it maps onto some affair in accordance with a definite mapping function determined by a history of joint successes of producer and consumer (or their ancestors).

For descriptive signs, that the sign maps in this way is a condition that is required for the consumer to perform its tasks, whatever they are, by the mechanisms normal for it. The content of the descriptive sign is not determined by the tasks its consumer performs. It is determined by what the sign needs to correspond to if the consumer is to perform its tasks in its normal way. The producer's job is merely to make a sign that corresponds in the right way to a world affair. If it does this in its normal way, by its normal mechanisms, the intentional sign it makes will also be a local natural sign. For directive signs, that the sign maps in the right way will be a result of the consumer's activity. The consumer's job is to cause the sign to map in this way by producing or causing a corresponding affair. It's job is to obey the producer's orders. But equally, it is the job of the producer to give orders that will benefit both it and the consumer. Only in that way can such a cooperative pair be selected for. Similarly, in the case of a descriptive representation, the consumer's job must be such as to use the representation in a way that will benefit both itself and the producer. Whatever more concrete jobs that consumer has, it must have this effect (or have had this effect) often enough to ensure selection of the cooperative pair.

Thus the cooperation between producer and consumer in production of natural signs can be accomplished in either of two basic ways. First, it might be that the producer is the one primarily responsible for making the sign correspond to the world. Then the sign vicariously guides the consumer in relation to the signified as the consumer performs some task mutually beneficial to itself and its producer. These are descriptive intentional signs. They are designed to stand in for world affairs, typically affairs outside the organism, and to vary according to these world affairs, controlling the animal's internal or external behavior as needed to adjust to these world affairs. Second, the consumer may be the one primarily responsible for making the world correspond to the sign. Then the producer's job is to make the sign be such that when the consumer has produced the signified world affair, the result is mutually beneficial to itself and to

the consumer. These are directive intentional signs. Directive signs guide the consumer in the production of world affairs that vary according to how the signs themselves vary. They are blueprints for what is to be constructed or brought about. Inner imperative signs are represented purposes of the organisms that harbor them, as represented purposes were discussed in Chapter One. They represent what it is their purpose to bring about.

But the most interesting kind of representation is the pushmi-pullyu representation. It is also the most primitive. Consider again the hen's food call. It is at once descriptive and directive. It is the hen's job to make the call coincide with the time and place of some food, and it is the chick's job to make the call coincide with the time it approaches that place. The call is false if there is no food; it is unsatisfied (not complied with) if the chicks do not come. Almost all animal signals are of this kind. For example, bee dances tell at once both where the nectar is and where the watching bees are to go. The vervet monkey's leopard, snake and flying predator calls tell what kind of predator is near and direct the response appropriate to that predator. In Part Four I will discuss pushmi-pullyu representations in detail, and explore the question how and why natural creatures should have evolved the capacity to form intentional representations also of more differentiated kinds.

On the theory proposed, intentional representations always come with propositional attitudes attached. It is essential to them that they have some kind of function, that they are designed for a particular kind of use. Frege's notion of sense, which implied that first you can represent a proposition, then add an intentional attitude to it, has done a lot of damage, I believe. There are not and could not be intentional representations that lacked attitude. There are no intentional representations without purposes, and having a purpose guarantees attitude. Hypothetical thinking, for example, or just thinking of possibilities, are extremely sophisticated activities, and ones which are only possible for a creature that sometimes uses the results in the production of ordinary descriptive and directive representations. It is because thoughts of this kind have the function of sometimes turning into more basic kinds of representations that they can exit at all.

In Knowledge and the Flow of Information, Dretske noted as a remarkable fact that our percepts 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." The percept "...skips over (or 'sees through') the intermediate links in the causal chain in order to represent ... its more distant causal antecedents" (1981, p. 158). And he worried about how abstract representations are possible, ones that carry only the information that, say, an object is triangular and not also that it is isosceles or equilateral, large or small, red or green and so forth. He introduced a special process called "digitalization" to solve this latter problem. Both problems are solved at a stroke, however, if we carefully distinguish intentional representations from natural signs. Local natural signs of distal affairs carry local information about all of the more proximal affairs on the route from them to those distal affairs (Chapter Four). But a natural sign that is also an intentional sign will carry only some of that information intentionally. The information that it carries intentionally is only the information that it has been selected for carrying, that is, only the information that is used by its cooperative interpreters. This information may be very abstract, and it may be about very distal affairs. If its consumers are designed to use only the information that something is triangular, then that is all the information that it carries

intentionally. If they are designed to use only the information that a predator is near, then it need not carry information about any more proximal affairs intentionally, such as patterns on the retina or properties that are local signs of the predator (Chapter Four). Similarly, of course, not every stimulus that an organism discriminates on the way to producing intentional representations is itself something intentionally represented. Nor must an organism be able infallibly to discriminate the distal objects, properties or kinds that it intentionally represents from similar ones. It only needs a fallible capacity to use some natural signs or other of these things under some conditions. Possibly it gets things wrong a large part of the time.

Dretske's example of the magnetosome (Chapter Three) illustrates these principles nicely, though, once again, it was not introduced by him for this purpose. The magnetosomes of northern hemisphere bacteria discriminate magnetic north which, when the normal mechanisms associated with their functions are in place, corresponds to geomagnetic north, to the direction of deeper water and to the direction of lesser oxygen. Of what, Dretske asks, is the magnetosome's orientation an intentional sign? Does it signify magnetic north, geomagnetic north, deeper water, or lesser oxygen? It is, of course, a recurring natural sign of all four. But it is an intentional sign only of lesser oxygen. This is because it needs, and needs only, to coincide with lesser oxygen to serve its purpose. Aerate the deeper water with oxygen and the direction the magnetosome points continues to be a natural sign of magnetic north, of geomagnetic north, and of deeper water, but the bacterium dies. Place the bacterium in southern hemisphere waters and the direction it points continues to be a natural sign of both magnetic and geomagnetic north, but the bacterium dies. These other natural signs are not what interests the bacterium, or rather, not what interested natural selection in selecting magnetosomes to build into the bacterium. Corresponding to magnetic north, to geomagnetic north and to deeper water is merely the normal mechanism by which the magnetosome manages to point to lesser oxygen. Of course Dretske is right that the magnetosome that directs the bacterium in the wrong direction because someone holds a bar magnet overhead 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 its function when no one has put any coffee 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.

Taking another kind of example, although each of the signs that emerges during sensory perception carries natural information about affairs at many levels of distality, each lies on its own level of intentional representation. For example, the edge-detector cells in early vision represent edges, not light intensity gradients across the retina. Their function is appropriately to guide internal acts of identification of contours and shapes, given the presence of certain edges. That the edges are where the detectors say they are is sufficient for them to do their jobs properly. Whether or not it was indeed gradients across the retina that caused the edge detectors to be properly aligned with edges on a given occasion, hence whether or not they carry natural information about gradients, is not relevant to the intentional information they carry. Similarly, the representations of contours and shapes that are produced from the edge detectors are not intentional representations either of edge detectors (they are not intentional signs of signs) nor intentional representations of edges. And the representations of interesting objects that are produced from the representations of shapes are not representations of object shapes but of fully

identified objects.

Similarly, the English sentence "It is raining" is a recurrent natural sign that the speaker believes it is raining. But it is not an intentional sign that the speaker believes it is raining. Its memetic function, derived compositionally from the combined memetic functions of its significant components, is to produce beliefs that it is raining, not beliefs that speakers believe that it is raining (see Chapter Two; also Millikan 1984, 2001a, 2001b). But of course it may be read as a natural sign by a hearer with the appropriate cognitive skills and in this way produce the belief that the speaker believes it is raining. That is another matter entirely.

I have emphasized the sense in which the teleological theory I am proposing is akin to an informational theory of representation and the sense in which it is akin to a functionalist theory. Let me also emphasize its kinship with picturing theories of representation. Like locally recurring natural signs, intentional signs are always members of a domain exemplifying a system of possible signs, the entire system running isomorphic to the domain of its signified affairs. The notion of a sign, I suggested, is best developed as at root an epistemological notion (Chapter Three). The embedding of a sign in a system of signs affords an interpreter a capacity to learn new things from signs, or in the case of inner signs, to perceive or think new things with signs. The value of a system of representation lies in its productivity. This depends, in turn, on there being some kind of isomorphism, in the abstract mathematical sense, between the domain of the signs and the domain of their signifieds. There is no need to place any limit, however, on the complexity of the semantic mapping functions that might map intentional representations onto their representeds. Isomorphisms can be defined by functions that are as bizarre, as gruelike, as you please. A bizarrely coded secret message from a CIA agent may be as much an "icon" or "picture" that maps onto a certain world affair in accordance with a generalized semantic mapping function as any sentence or diagram. Signs must be things apt for use by sign users, but sign 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 representations 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 further 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.

Notice that in describing intentional representation as closely analogous to the representation of natural signs, I have made no reference to inference or rationality or to representations being "calculated over." Opposing a central current in contemporary American philosophy, I claim that rationality is not "the mother of intentionality" (Dennett). Rationality will play an important role in the discussion of inner representation before this essay ends, but rationality is a characteristic only of systems using representations in certain ways, and most intentional representations are not used in this way.⁹

The theory of intentional representations that I have presented in this chapter is explained in a rather different way, and with some differences, too, in terminology, in (Millikan 1984)

chapter 6. Many critical questions about the theory as presented there have come and gone, but there are two that seem to cling. Let me conclude this chapter by saying a word about them.

In (Millikan 1984) chapter 6, I said that the content of a descriptive or "indicative" representation was determined by what its consumer needs the representation to map onto if it is to perform all its functions in accordance with a "most proximate Normal explanation." "Normal explanations" are what I have been calling "normal mechanisms" here. Karen Neander (1995) has objected that among normal conditions that must be mentioned, say, for the male hoverfly's female-detecting systems to carry out all of their functions are that the female is fertile and that she won't be eaten before she reproduces. But if it were a function of the hoverfly's signal producers to signal when a fertile female who won't soon be eaten passes by, there would have to be a normal mechanism by which these producers had historically performed this function (Chapter five). There would have to be a systematic way that it managed to produce representations of fertile females who were not about to be eaten. For this sort of simple perceptual device, the explanation would have to be that the hoverfly's perceptual systems were sensitive to some kind of recurrent natural sign of the affair they were to represent. But hoverflies do not encounter recurrent natural signs that the things causing images on their retinas are fertile or that these things won't be eaten. The point is not that the hoverfly has no way of discriminating natural signs of females that are fertile and won't be eaten from other images crossing its retina. The point is that there simply are no such signs crossing its retina. The domain in which the hoverfly operates is one in which the chance that the shadow crossing its retina, assuming that it is of a female hoverfly, is also of a fertile female not about to be eaten is no higher than the chance of any arbitrary female hoverfly being fertile and not about to be eaten. By contrast, assuming that it is the shadow of a hoverfly, the chance of the shadow being that of a female is considerably higher than the chance of an arbitrary hoverfly being female. This is because only female hoverflies cruise, while male hoverflies hover. Similarly, given where the male hovers, the chance of the shadow crossing his retina being that of a hoverfly rather than of some other small particle of matter is also very much raised. He purposefully hovers in a recurrent sign domain where such a shadow is very likely to be that of a female hoverfly. (An unarticulated image crossing the retina can produce an intentional sign that the thing spotted is both a hoverfly and female, just as the bee dance can represent nectar, hive and sun without mentioning any of them explicitly.)

A second query concerns the possibility that there might be biological systems whose jobs are to produce false representations. For example, there is some evidence that people who are overconfident are more successful at performing certain tasks than people who evaluate their skills more accurately. First we should notice that any such system responsible for producing overconfidence would not have been selected for the fact that it produced false beliefs, but for the fact that it produced lots of confidence. For example, if a person were always 100% accurate at performing a certain task, there would be no profit in his believing falsely that he was less competent than that. Falseness itself could not be the point. More important, however, we should note that very many biological systems ride piggyback on systems developed earlier for other purposes. Systems whose jobs were to distort certain beliefs would have to ride on more general systems whose basic jobs were to produce true beliefs. Otherwise there would be no standard mapping rules according to which the distorted beliefs were designed to map world affairs hence

according to which they were false.

FOOTNOTES

1. At least they appear to be in (Dretske 1995).
2. Kim Sterelney (1995, 2001) overlooks the requirement of cooperation for producer and consumer laid down for intentional icons in (Millikan 1984, 1993)
3. Kim Sterelney (1995) wishes to require this kind of flexible tracking of information for intentional representations. One can use one's terms as one likes, so long as the ideas are clear and useful.
4. This is how Nick Shea put it to me. He nudged me into taking this more flexible position on locally recurrent signs, whereas I had been vacillating. The decision is, in a way, merely verbal. But some ways of using words help one to think more clearly than other ways do. There is one comment I feel obliged to add, however. In (Millikan 2000) I argued that there is no proof that inner representations that are understood as representations of the same by the cognitive systems need to be carried by similar vehicles. But the various tokens of a recurrent natural sign do have to be similar to one another. But even so, inner representations would still be local natural signs in the sophisticated sense that the mitten in the path in Chapter Three was a natural sign. The cognitive systems do seem to have systematic ways of determining what they are local signs of in this more sophisticated sense. (How one knows what one's inner signs are about is also discussed in (Millikan 2000) Chapter 13).
5. For conception, these systems are discussed in (Millikan 2000) Chapter 7.
6. I would say "semantic rule" if it were not that the notion of a "rule" tends to have prescriptive overtones that I wish, by all means, to avoid.
7. For defense of the position on knowledge suggested by this cryptic remark, see (Millikan 1993) Chapter 12.
8. After Hugh Lofting's charming double-facing creature by that name. For a general essay on pushmi-pullyu representations, see (Millikan 1996).
9. In this disclaimer lies my main disloyalty to my teacher Wilfrid Sellars. It should be clear that in many other ways I am very deeply indebted to him.