

PART II
NATURAL SIGNS AND INTENTIONAL SIGNS
CHAPTER THREE
LOCAL NATURAL SIGNS AND INFORMATION

Intentional signs are signs that can be false or unsatisfied. Natural signs cannot. Black clouds do not mean rain unless it actually rains, and an elevated temperature does not mean illness unless one is actually ill. True intentional signs are sometimes thought to be just natural signs that have been purposefully produced. The information they carry is thought to be just natural information that they happen to have been designed to carry. It is their purpose or function to carry this information. Certainly that is Fred Dretske's claim (Dretske 1986, 1988, 1995). I will agree that true intentional signs often do carry natural information. But to support this claim it will be necessary to understand the terms "natural sign" and "natural information" quite differently from the way Dretske defines them. Also, not all true intentional signs do carry natural information. Some of them are true by accident. Further, I will argue that it is not the purpose of an intentional sign to carry natural information. Carrying natural information is merely the usual means by which an intentional sign gets to be true.

In Knowledge and the Flow of Information (1981) Dretske defined the notion "natural information" quite strictly. I will argue that his definition is not adequate even to the simplest tasks that he wished natural information to perform, and propose a description of "local natural information" that I believe is more adequate. The most useful kind of local natural information is carried by "locally recurrent" natural signs. In this and the next chapter, I will introduce locally recurrent signs and local natural information and describe some of their many virtues. In Chapters Five and Six, I will explore their relations to intentional signs. The perceptual and cognitive systems of every animal are deeply dependent on the local natural information found both in the environment and within the organism itself. Without such information there could not be any intentional signs or intentional information.

Dretske's use of the term "natural information" vacillates in Knowledge and the Flow of Information (1981) and also in Explaining Behavior (1991).¹ Sometimes his usage conforms to his original strict definition in (1981) but other times it gestures more broadly. According to Dretske's original definition and the discussion immediately following it, a signal carries "information" about a source only if the signal's occurrence yields a probability of one, determined in accordance with strict natural and logical necessity, that the source is a certain way. On the other hand, many of Dretske's examples of natural information seem to rest not on strict natural necessity, but merely on statistical frequencies at the source. For example, he tells us that although in some woods a certain kind of tracks made by quail might carry the information that quail are present there, on the other hand, "[i]f pheasants, also in the woods, leave the very same kind of tracks, then the tracks, though made by a quail, do not indicate [carry the information] that it was a quail that made them" (1991, p. 56). Here, not natural law but statistical frequencies at the source end of the information channel appear to be determining whether the tracks carry natural information. Yet Dretske explicitly claims that statistical frequencies are not enough to determine the presence of natural information:

Even if the properties F and G are perfectly correlated...this does not mean that there is information in s's being F about s's being G...For the correlation...may be the sheerest

coincidence, a correlation whose persistence is not assured by any law of nature or principle of logic....All Fs can be Gs without the probability of s's being G, given that it is F, being 1. (1981, p. 73-4)

But surely, whether or not there are pheasants as well as quail in these particular woods is a matter of statistical frequency not natural law. If we look only at natural law, given that pheasants as well as quail could lawfully produce such tracks, there could never be a probability of one that such tracks are produced by a quail, no matter how far these tracks happened to be from any actual pheasants. Logic and natural law do not change over space and time with variations in the distribution of pheasants and quail.

This is an insurmountable problem, as I see it, for Dretske's explicitly stated theory of intentional representation. Nearly all of the kinds of information needed by us, and by all other organisms as well, for securing what we need in an inclement world, is information that cannot possibly be acquired without leaning on certain merely statistical frequencies. Consider a rabbit that needs to know when a predator is near. However she detects a predator, no natural law can require it to BE a predator that causes her predator detectors to fire. Whatever information channel she uses, it is always nomically possible that non-predators should exist who would activate it. Suppose for the sake of the argument (though very implausibly) that there are unbreakable natural laws that concern the effects of foxes on rabbit sense organs. Still, there surely are no laws that nothing else could possibly produce these same effects on rabbit sense organs. For example, no natural laws prevent the introduction of new species.

Moreover, Dretske is often quite explicit that it is only relative to certain channel conditions that the sign carrying natural information must correspond with certainty to what it signifies. Suppose that we contrast, for example, (1) the probability that it is the presence of helium under excitation that produces such and such a spectrum with (2) the probability that it is the gas tank's being half full that produces such and such a reading on the gas gauge. Perhaps the spectrum corresponds to the presence of helium with a probability of one, period. But the gas gauge reading corresponds to half full with a probability of one only on the assumption that there are connections of a rather exact sort between the tank and the gauge. Similarly, when the sounds of the recorded music get first louder and then softer, this may be a sign that the orchestra executed a crescendo and then a decrescendo. But this depends on stable placement of the microphone during recording, and on whether the recording engineer, or someone at the loud speaker end, twiddled the knobs. Channel conditions of this sort are not considered by Dretske to be part of the signs they transmit. The signal carrying natural information indicates neither what channel it comes through nor, supposing the channel known, what semantic rule to apply in order to read information coming through that channel. Both would have to be independently known.

The question arises, then, what use signals carrying bits of natural information could be to an organism. The mere fact that a signal carries certain natural information seems not to bear on whether a creature could learn anything from encountering that signal or know anything by virtue of harboring that signal in its brain. To learn anything from such a signal would seem to require that the organism have a means of detecting not merely the information bearing signal, but detecting also the presence of a particular kind of information channel, as well as a means of knowing what semantic mapping function, what translation rule, to apply. Returning to our rabbit, the difficulty is that no matter what means she uses for detecting the presence of whatever

particular kind of external information channel she knows how to use, it is always nomically possible that the presence of some other external circumstances should cause the very same signals to arrive as through this channel, but with different meaning. Just as no natural laws prevent the introduction of new species that would affect the rabbit's senses in the same way given these channel conditions, no natural laws prevent the introduction of new external circumstances not detectable by the rabbit that substitute for these channel conditions. If the capacity of an organism to represent something mentally were to depend on its ability to discriminate that thing from all others in accordance merely with natural law and logical necessity, it is clear that no organism could possibly represent anything distal.²

Nor, of course, does the rabbit care that there are no such laws. As long as in her actual environment, given its actual statistics, she is good at detecting foxes, all is well. Similarly for the things we humans detect, for all the various things we acquire empirical information about. Our abilities to represent these things couldn't possibly depend on natural information that fits Dretske's strict characterization, nor would the insertion of this kind of information make much difference to our practical lives.

I conclude that a theory of natural information that will help to explain how real animals manage to obtain useful information will need to introduce statistical considerations about the environment in some controlled way. The relevant statistics would have to bear both on conditions at the source of information and on channel conditions. Moreover, a reference to non-accidental stability of these statistics over trackable regions would need to be built in, for organisms must somehow keep track of the regions in which the relevant statistics continue to hold. I will try to describe such a new notion of natural information, information of a more user friendly kind than Dretske described, calling it "local information." Information defined by Dretske's strict rule above I will call "context-free" natural information.³

At the same time I will try to speak to another deep problem intrinsic to any informational semantics that defines information by reference to natural law only, ignoring local statistics. The difficulty is that there are no natural laws about individuals. No natural laws are laws just about George Bush or just about the White House or just for the moon. On a theory that recognizes only context-free information, no natural sign can carry the information that it is George Bush who is speaking from the White House or that the moon is out again. Context-free informational semantics is debarred in principle from explaining how there can be representations of individuals. Consonant with this, informational semanticists have invariably confined themselves to a discussion of predicative representations. For example, when Dretske explains what it is for a signal to carry the information that s is F, he concentrates entirely on the representation of F, telling us nothing about what it would be, say, for t rather than s to be represented as being F. The subject of the intentional representation that s if F is not represented. Compare: if Billy and Johnny are identical twins, that a photograph is a photograph of Johnny rather than Billy is not represented IN the photograph. Dretske describes a signal carrying the information that s is F exactly as if it were, in this way, like a photograph of Johnny. The description of local natural information that I will give explains how there can exist natural information concerning individuals, and how the subject term of a natural sign that s is F can be naturally represented.

We can begin by reflecting on what count as natural signs and their meanings in everyday life.⁴ What makes black clouds a sign of rain and birds flying south mean that winter is

approaching? Consider Dretske's bird tracks in the woods. If both quail and pheasant leave tracks of exactly the same kind in the woods, are these tracks signs of quail, or signs of pheasant, or are they not signs of either (as Dretske claimed)? If there's no way to tell the difference between quail tracks and pheasant tracks, it seems to me that a natural remark would be "those tracks might mean quail or they might mean pheasant, we'll just have to wait and see." That is, we'll have to wait and see which they really DO mean, for they do mean whichever was their actual cause. No probability of one is in view here, but only a real causal connection. Similarly, Dretske told us that "[t]he red spots all over Tommy's face mean that he has the measles, not simply because he has the measles, but because people without the measles don't have spots of that kind" (1991, p. 56). But suppose that scarlet fever can cause spots like that too. What a cautious physician will say, I believe, is something like this: "Those spots probably mean measles, but they could also mean scarlet fever. I think we had better take a culture." Indeed, a doctor might say this even if measles is ten times more frequent in the schools this fall than scarlet fever. Not even a high probability is always required for attribution of natural meaning.

And there are examples that are more extreme. Suppose that I take my daughter's mitten lying in the path to our front door to be a sign that she is home from school. For me to be right, the connection must not be accidental. If she dropped her mitten on the way to school rather than the way home, then even if she is home, her mitten is not a sign of that. As Dretske said, merely correlating (in this case, merely coinciding) is not enough to make one thing mean another. But what is the general probability of a mitten in the path meaning that a daughter is home from school? How would one calculate that? Also, what would count here as channel conditions? Again, suppose you are conspiring with Joe in an assassination attempt. He hands you a briefcase with a homemade bomb in it saying, "Here is the briefcase for you to leave in his office. Be careful though. When you hear it beep twice it will explode in two minutes. Leave the room quickly." The beeps will be a sign that the briefcase is about to explode and they will be causally connected with that. But surely beeping in briefcases doesn't generally precede their exploding.

What seems to be happening in these cases is that there is a real causal connection between two things such that in the circumstances one does depend on the other, and that given what one already knows, one is able to track that connection, hence come to know or suspect one thing on the basis of knowing the other. In the mitten case, the inference is to the best explanation one thinks of, and this explanation, as it happens, is correct. In the bomb case, the inference is based on knowledge of causal principles plus quite exact knowledge of the channel conditions. The central thing common to all these examples of natural signs, I suggest, is that in each case it is possible for a true belief to be reached about one thing from knowledge of the other, the transition from one belief to the other being based on prior knowledge or experience, where the truth of the belief reached will not be accidental because the connection in thought correctly and nonaccidentally tracks a dependency in nature. A natural sign of a thing is something else from which you can learn of that thing by tracking in thought a connection that exists in nature. The notion of a natural sign is at root an epistemic notion.

The use of isolated natural signs, such as the mitten in the path and the beeping briefcase, clearly depends on having a good bit of prior knowledge with which to combine one's observation of the sign. The use of isolated signs could not then serve as the original foundation on which perception and knowledge is built. Such a foundation would require signs that can be

used alone or one by one. Further, for a system to acquire the capacity to use a natural sign either by means of natural selection or by means of learning, the sign would need to be one that recurs, and recurs with the same natural signification or meaning. For these basic natural signs, it must be true with some generality that the same kinds of signs are connected to the same kinds of signifieds. There must be a correlation between similar signs and similar signifieds, and a correlation, as Dretske has said, that is not accidental.

Correlations are defined relative to reference classes. Within what reference class must A's be correlated with B's for a particular A to be an instance of a recurrent natural sign of B's? Dretske says that if in his woods only quail leave a certain kind of tracks —call them "ε-tracks" for the quail's three-toed foot— then ε-tracks in his woods are signs of quail. He implies that if in the woods next door only pheasants were to leave ε-tracks, then the ε-tracks next door would be signs of pheasant. But why take Dretske's wood and the wood next door as separate reference classes? Why not take both together as our reference class?

Suppose that looking world wide, we find ten more species that leave ε-tracks and there is no strong correlation between ε-tracks and any one of these species. Why not pool all these together and assert that ε-tracks are never recurrent signs of any particular species of bird? Or suppose that in my American woods only poisonous mushrooms have a certain reddish color underneath whereas in the French woods only edible mushrooms do. Is that certain reddish color underneath a mushroom ever a recurrent natural sign of anything? What reference class should we use?

We should be clear that this problem cannot be solved by relativising natural signing to an arbitrary reference class. We cannot just say that this particular ε-track in Dretske's woods exemplifies a recurrent natural sign of quail with reference to the class of ε-tracks in Dretske's wood though not with reference to the class of ε-tracks worldwide. We cannot just say that the reddish color under my American mushrooms is a recurrent natural sign of poisonousness with reference to the class of mushrooms in America but not with reference to mushrooms worldwide. For what is to prevent arbitrary divisions of the reference classes? Suppose I combine a highly discontinuous set of sectors from Jill's wood in Massachusetts with a highly discontinuous set from Jack's wood in Minnesota, such that although each of these small sectors has pheasants living nearby, mainly quail live within my gerrymandered class. Then I name the whole of this scattered terrain "Q-wood," claiming that each quail track in Q-wood exemplifies a recurrent natural sign of quail with reference to Q-wood. I also admit, of course that the infrequent pheasant tracks found in Q-wood nicely exemplify recurrent natural signs of pheasant relative to some other reference classes, just not with reference to Q-wood. Relativised in this manner the notion of a recurrent natural sign obviously spins its wheels. It does no work.

Well, what kind of work should it do? What we would like, I suggest, is for the notion of a recurrent natural sign to explain why a person might be able to use the recurrent sign as an indicator of its signified with some success. To be genuine, this explanation must have its footing in nature. Consider, for comparison, the fact that John, who happens to be 40 years old, 5 feet ten inches tall, and to like sports, has a mortgage. This fact would not be explained by citing the fact that the average man who is 40 years old, 5 feet ten inches tall and likes sports has a mortgage. Similarly, that the inference from "This is an ε-track in Q-wood" to "This is the track of a quail" is likely to yield a true conclusion is not explained by citing the statistics on quail in Q-wood. It

is a clear example of the fact, now well known, that a logical derivation from true premisses is not always an explanation. What is needed is some way to delineate relevant natural classes. We need to require that A and B are correlated within a relevant natural reference class for one to be a recurrent natural sign of the other.

Another problem needs to be solved here as well. Dretske told us that even a perfect correlation between A and B is not enough to ensure that A carries natural information about B because a perfect correlation could be perfectly accidental. Of course, a particular A does not mean B unless it actually coincides with a B, but Dretske demands more. He wants A's that mean B's to correlate with B's not by accident but for a reason. But what is it, exactly, for a correlation not to be accidental? The standard thought here is that a correlation between A's and B's is not accidental just when there is a causal connection between A's and B's, the A's causing the B's or vice versa, or perhaps something else causes both the A's and the B's. But suppose that this causing itself happens accidentally. The mechanisms by which the A's cause the B's are each time entirely different. Each A is in some circumstance such that it causes a B, but these circumstances are all different. From chaos a correlation accidentally emerges. Are we to say that A's are recurrent natural signs of B's?

Here is my suggestion. What we are trying to do is to construct a notion of recurrent natural sign that will help us to understand how it is possible for an animal nonaccidentally to recognize the recurrence of such a sign so as to use it effectively. The project is not conceptual analysis but theory construction. If A's in some reference class are to be recurrent natural signs of B's, we want it to be possible for an animal to come to learn of B's from encounters with those A's, where what counts as learning is acquiring true beliefs non-accidentally. From the standpoint of natural epistemology, the cause of one's moving from encounters with A's to representations of B's or to accommodations to B's should be connected with the reason these moves are sometimes correct. Now one's moves of this sort will be based, in central cases, on one's experience --or someone else's experience, or the experience of the species-- of a correlation of A's with B's within some sample. The inference then predicts that the correlation will continue to various new samples encountered. The inference will succeed "for a reason" rather than "by accident" only if there is a reason why the correlation persists from the old sample into the new. A natural reference class for a sign --the natural domain within which certain A's are "locally recurrent signs" of certain B's-- is a domain within which the correlation of A's with B's extends from one part of the domain to other parts for a reason, and a domain it is possible for an organism to track.⁵

Here are some simple examples. Suppose that every ball in the urn is black. Then every ball I pull out today will be black. That a ball comes from this urn is perfectly correlated with its being black. Now consider tomorrow. Granted nothing disturbs the urn overnight it will be no accident that the correlation I discovered today holds tomorrow. Balls left undisturbed do not change their colors overnight, nor do urns left undisturbed change their contents. The correlation continues for a reason. Being a ball in this urn is a local sign of being black for as long as the urn remains undisturbed.

What about the balls in the next urn? Depending on the histories of these two urns, there may or may not be a reason why the contents of one urn would reflect the contents of the other. If there is a reason for the contents to match, say, they came from the same factory in response to

the same middleman's order, then the relevant recurrent sign domain includes the contents of the second urn as well. Otherwise not --not even if all the balls in the second urn do happen to be black. What this example brings out, however, is that in order to make use of a local sign, one needs a method of tracking, or recognizing, or managing to remain within, its local domain. In this case, one would need a way of recognizing other urns that came from the same factory in response to the same middleman's order. But then, in the simpler case above involving only one urn, a way of tracking the relevant domain was also necessary. One would need to keep track of which urn it is that contained only black balls yesterday.

Suppose that this morning my gas gauge reads half full when it is three quarters full and, later today when the tank is half full, the gauge reads one quarter full. What about tomorrow? The channel conditions that obtain today between the height of gas in my tank and the gauge readings concern bits of hardware and wiring that, should nothing outside disturb them, strongly tend to remain in place in accordance with natural conservation laws. It will be no accident if my gas gauge continues to underestimate the amount of gas in my tank by about the same amount tomorrow. The readings on my gauge are locally recurrent signs of the amount of gas in my tank.⁶

Suppose that most ϵ -tracks now being made in this wood are by quail. Quail in this wood this week causally result in quail being in this wood next week and next year, and also in quail soon being in the wood next door. For good reason ϵ -tracks are repeating themselves, over and over, now and then, here and there, in these woods and in contiguous woods, each time connected to quail. The correlation persists and spreads over a local time and space for a reason. There is a causal stream in which the correlation holds.

In general, it is because many conditions persist and/or spread or replicate themselves over time and space, that there are locally recurrent natural signs belonging to recurrent sign domains. Identical signs that come from other domains, for example, from different regions of space and time, may be local natural signs of something quite different. Clouds of a certain sort on this side of the mountain may mean rain but on the other side not. That is because certain contours on each side of the mountain persist, and because the causes of air moving in and the places the air comes from on the two sides tend to persist. The local statistics for each place repeat themselves for a reason.

Suppose that all the spots that appear on the children in our school this week are from measles. And suppose that the children pass measles on to their friends, brothers and sisters. Thus it happens that spots due to measles appear on lots of other children in our school next week. In this school, spots will be a recurrent natural sign of measles until the epidemic is over. If an epidemic of scarlet fever should now intrude as well, then there will be two separate grounded persistent correlations in our schools, one between this kind of spots and measles, the other between this kind of spots and scarlet fever. Within the natural reference class that is our schools, some of these spots are recurrent natural signs of measles, others of scarlet fever. ("These spots might mean measles. But they might also mean scarlet fever. I think we had better take a culture.") On the other hand, each of these natural signs may be viewed as having its own separate gerrymandered geographic domain, determined by the contingent contacts of individual children carrying one disease or the other with children to whom they have passed these germs. Normally one would not attempt to track these two gerrymandered domains separately. But if one

of the diseases is serious enough, exactly this sort of tracking will be attempted by physicians and health authorities, in order to recognize natural signs of the disease more quickly and accurately.

Consider now the look, the appearance, of Johnny's face. Your encounters with that look are likely to be highly, indeed perfectly, correlated with your encounters with Johnny. This is because natural conservation laws plus principles of homeostasis built into Johnny tend strongly to preserve Johnny's appearance from one day to the next, indeed, over many years. The look of Johnny's face is a locally recurrent natural sign of Johnny. It carries local natural information as to the presence of Johnny. There are no laws that concern individuals as such, but there may be very numerous local well grounded correlations that do.

Successfully to interpret a locally recurrent sign you must keep within its natural domain. You must stay within the boundaries. On the other hand, it may not be necessary to discriminate the boundaries of the sign domain in order to stay within them. The rabbit is likely to be born and to die without ever leaving the domain containing various of the locally recurrent signs of fox that it recognizes, and you are not likely to leave the domain in which that face (it's look) reliably signifies Johnny. In many fields of knowledge, becoming an expert involves learning to recognize diagnostic signs for a variety of kinds, that is, local signs not having their boundaries within the environment one traverses. Or it may involve learning how to track various causal streams of signs. The modern doctor diagnosing signs of infectious diseases will be helped by understanding as much as possible about what causes the spread and the boundaries of each disease.

Knowing how to track is most obviously important if one needs to be able to identify each of various similar individuals, individuals whose signs are hard to distinguish. To be able to recognize my glass from among other identical glasses at a party so as to drink out of it and no other, I have to keep it in my hand, or remember where I put it down, or who took it away offering to refill it, or, should I accidentally drop it, I must know the way glasses are likely to fall and roll, and so forth. This means I must understand quite a lot about the possible trajectories of individual objects of that general kind. Squirrels behave differently from drinking glasses. If I want to recognize the same squirrel again as it jumps among the branches, I will have to use quite a different method of tracking. Squirrel shape, color, movement style, behavior (dropping hickory nuts from trees, chattering, squeaking), and other characteristic squirrel effects (ways the branches quiver and so forth) are each recurrent signs of squirrel. As such they are also signs of Chipper specifically, but only as manifested within a certain very much more local domain that may often shift in location.

Failure to account for our capacity to represent individuals in language and thought has been, perhaps, the most serious failing common to contemporary naturalist theories of content. This failing is not often discussed, but it should certainly be considered embarrassing for a tradition that has strongly rejected the classical analysis of individual concepts as composed of definite descriptions. If a cognitive mechanism is to have producing representations of individuals as one of it's functions, there had better be some practical way of producing these representations. That there exist local natural signs of individuals will help us to explain this possibility in Chapter Four.

But local sign domains may sometimes be recognized for practical purposes without tracking, and without understanding much about why they flow where they do. In northern

climes, mushrooms that look like that are always edible whereas in southern climes they are not, except in certain parts of South America. You can know that without knowing why. You can recognize a sign of a local domain. If the domain of a local sign is signed by another sign or signs, however, this does not make the whole complex involving both signs together into a context-free sign, for the sign of the domain, as a sign of a portion of the actual world, will not itself be context-free. There is no way of adding to a local sign so as to completely free it from context.

The notion of a recurrent local sign is designed to explain how it is possible for an animal to use natural signs to collect information about its world. The kind of knowledge that earthly creatures have is knowledge applicable in the domains they inhabit, not knowledge for arbitrary nomically possible worlds, nor for other domains, regions or eras within the actual world. If the question arises how strong the non-accidental correlation between A's and B's must be within a domain for the A's to count as locally recurrent signs, the following is what is important. A strong enough correlation to count in determining a local sign to be such is one that is strong enough to have actually influenced sign use, either through genetic selection or through learning. For the point of introducing locally recurrent signs is as a tool for understanding how perception and cognition are possible. Our concern is only with natural signs that are actually used by organisms. Suppose that an animal has only crude and inaccurate ways of tracking a local sign's domain. Suppose that the animal tracks only a rather vague domain that includes other similar signs as well and is unable to distinguish accurately among these signs. (Compare: both measles spots and scarlet fever spots are found in our schools and no one knows how to take a culture). Still, I will soon argue, the animal may be able to use these signs effectively in the process of producing intentional signs. Intentional signs, I will argue, do not have verificationist meaning. No probability of one, nor even a particularly high probability, needs to be involved anywhere in the origin or use of an intentional sign. But here I have gotten ahead of my story.

One last comment about locally recurrent signs. As I have described them, it is not necessary that there be any causal connection between a locally recurrent sign and the affair that it signifies. This is because the relation between tokens of the sign and affairs signified may be reiterated over a domain merely because the signs and the affairs each persist, maintaining the same relation to one another. A nice example of this comes from Dretske (1986), though it was not his intention to illustrate this point. The magnetosomes of certain anaerobic bacteria are tiny magnets that serve as sense organs. In northern hemisphere bacteria they steer the bacteria toward magnetic north, which means steering them towards geomagnetic north which means steering them into deeper water which contains less oxygen. (Oxygen poisons these bacteria.) The direction in which the magnetosomes point are locally recurrent signs of the direction of lesser oxygen. But there is no causal connection between the direction of magnetic north and the direction of lesser oxygen. The information is locally recurrent information, the reference domain of which is extended through time and space because magnetic north continues to equal geomagnetic north so long as no local bar magnets (and so forth) intervene, geomagnetic north continues to remain in the same relation to the earth hence to deeper water, and lesser oxygen continues, for entirely independent reasons, to remain where the deeper water is.

FOOTNOTES

1. See (Millikan 2000, Appendix B).

2. Jerry Fodor likes to explain the supposed laws that produce covariance between representation and represented in terms of "ceteris paribus laws," "counterfactuals," and what is true in "near possible worlds," but nowhere has he explained with care what ceteris paribus laws are, exactly how the counterfactuals are to be stated, or what the principles are that determine relevant possible-world nearness and why. Following his lead, irresponsible incantations of "ceteris paribus" everywhere pollute the literature on this subject.

3. In (Millikan 2000 Appendix B) I use the terms "soft natural sign" and "soft natural information" rather than "local (recurrent) sign" and "local information." I think the term "local" captures the idea we need a bit better. The central idea is that there is an historically positioned domain to which the sign is bound. One might also call these signs "bounded signs," but the correlative term "bounded information" somehow doesn't work. Those familiar with the description of "historical kinds" and "individual substances" from my

(2000) will notice that the domains of local signs are bound together in much the same way that historical kinds and individuals are. In both cases, the underlying claim is that to understand the possibility of any kind of cognition we must recognize that contingently existing historical regions of continuity as well as universal natural laws are tapped by cognizers.

4. The passages immediately following may look like conceptual analysis of the notion "natural sign." I have often complained about the tradition of conceptual analysis in philosophy. Philosophy should be concerned primarily with theory construction. Ordinary usage should be described by lexicographers, who do a much better job of it anyway. But it is true that a careful look at the subtleties of usage can sometimes help us to "assemble reminders" about the complexities of the phenomena that we are trying to address. Ordinary usage can be merely idiomatic, superficial, indecisive, erratic, vague, gerrymandered, inconstant, revealing nothing of theoretical interest. Other times it can be surprisingly deep and subtle, following underlying natural contours of great interest.

5. This description of local natural signs is designed, in part, as a way to clarify what it is for Dretske's "channel conditions" to remain constant and how it is possible for an organism to keep track of the domains in which channel conditions remain constant.

6. Dretske discusses exactly this kind of case in his (1981). I am hoping to articulate somewhat better what I believe he was groping for there in his discussion of channel conditions. I believe that his attempt to show how the idea of natural information could explain what knowledge is, coupled with his assumption that knowledge requires certainty, got in his way. He was also concerned about what he called the "Xerox principle." Information that there exists information that p ought itself to be information that p. For this reason too he wanted probabilities of one for his information. Validity of the Xerox principle for locally recurrent natural signs is discussed in Chapter Four.