Appendix B

What Has Natural Information to Do with Intentional Representation?[1]

> "According to informational semantics, if it's necessary that a creature can't distinguish Xs from Ys, it follows that the creature can't have a concept that applies to Xs but not Ys."
>
> (Jerry Fodor, *The Elm and the Expert*, p. 32)

There is, indeed, a form of informational semantics that has this verificationist implication. The original definition of information given in Dretske's *Knowledge and the Flow of Information* (1981, hereafter KFI), when employed as a base for a theory of intentional representation or "content," has this implication. I will argue that, in fact, most of what an animal needs to know about its environment is not available as natural information of this kind. It is true, I believe, that there is one fundamental kind of perception that depends on this kind of natural information, but more sophisticated forms of inner representation do not. It is unclear, however, exactly what "natural information" is supposed to mean, certainly in Fodor, and even in Dretske's writing. In many places, Dretske seems to employ a softer notion than the one he originally defines. I will propose a softer view of natural information that is, I believe, at least hinted at by Dretske, and show that it does not have verificationist consequences. According to this soft informational semantics, a creature can perfectly well have a representation of Xs without being able to discriminate Xs from Ys.

I believe there is some ambivalence in Dretske's writing about natural information, especially noticeable when comparing KFI to *Explaining Behavior* (1991, hereafter EB), but if we ignore some of Dretske's

1 This chapter is drawn from a paper that was originally presented at the Conference of the Royal Institute of Philosophy on Naturalism, Evolution and Mind in July 1999.

examples, the explicit statement of the theory in KFI is univocal. This theory is also strongly suggested in Fodor's work on mental content (1990, 1994, 1998) and seems to be consonant with J. J. Gibson's use of "information" as well.

According to Dretske,

> A signal *r* carries the information that *s* is F = The conditional probability of *s*'s being F, given *r* (and *k*), is 1 (but, given *k* alone, less than 1). (KFI, p. 65)

Dretske's "*k*" stands for knowledge already had about *s*. Knowledge that p is belief that is caused by information that p. It follows that a signal carries the information that *s* is *F* when either it alone, or it taken together with some other signal that has also been transmitted to the receiver, returns a probability of 1 that *s* is *F*. Thus, I suggest, we can drop the parenthetical "and *k*" in the formulation and just say that a signal carries the information that *s* is *F* if it is an operative part of some more complete signal, where the conditional probability that *s* is *F*, given the complete signal, is 1 but would not be 1 without the part. Thus we eliminate reference to knowing.

What is meant by saying, in this context, that the occurrence of one thing, "the signal," yields a probability of 1 that another thing, "*s* being *F*," is the case? In a footnote, Dretske explains:

> In saying that the conditional probability (given *r*) of *s*'s being F is 1, I mean to be saying that there is a nomic (lawful) regularity between these event types, a regularity which *nomically precludes r*'s occurrence when *s* is not F. There are interpretations of probability (the frequency interpretation) in which an event can fail to occur when it has a probability of 1 . . . but this is *not* the way I mean to be using probability in this definition. A conditional probability of 1 between *r* and *s* is a way of describing a lawful (exceptionless) dependence between events of this sort. . . . (KFI, p. 245)

and in the text he tells us:

> 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. (pp. 73–4)

The probability that *s* is *F* given *r* must follow, it appears here, *given merely logic and natural law*. That is, the necessity must be strict natural necessity.[2]

The next question concerns the reference classes intended when referring to "the probability that *s* is F, given *r*." *r* was said to be a signal and *s* being *F* would seem to be a state of affairs, but if there are causal laws necessitating the one given the other, these laws must be general. There must be certain general aspects under which we are considering *r*, and the fact that *s* is *F*, by which they are connected in a lawful way. They cannot be connected in a lawful way merely as an individual occurrence and an individual fact. It must be a certain type of signal that determines, with a probability of 1, a certain type of fact. And this will yield two reference classes for the probability, the class of "signals" of a certain type and the class of facts of a certain type, such that the probability that a signal of that type is connected with a fact of that type is 1. What reference classes are intended, then, when it is said that a certain *r* carries the information that a certain *s* is *F?* When Dretske says that the pointer on my gas gauge being at the ½ mark carries the information that my gas tank is half full, in which two reference classes are these two conditions being considered, so as to make that so?

Clearly the reference classes cannot be (1) all pointers on gas gauges that point to the one half mark and (2) all gas tanks that are in the same cars as those gauges. For some gas gauges are broken or disconnected or badly calibrated, and even if none were, it would not be a matter of natural law that they couldn't be broken or disconnected or badly calibrated. Rather, as Dretske emphasizes in KFI, a reference must be made here to the presence of certain "channel conditions." In this case, channel conditions consist in a fairly intricate collection of surrounding conditions including various connecting parts the presence of which is needed before natural laws will guarantee that the gas gauge will read half full only if the gas tank is half full. One kind of thing carries information about another in accordance with strict natural necessity only given specified channel conditions. The two reference classes concerned contain only members connected by these channel conditions.

We can contrast this to the notion of a *ceteris paribus* law. According to the classical view, a *ceteris paribus* law is one that is true in accordance

2 The necessity may go in either temporal direction. For example, an effect might carry information about another effect of the same cause.

with natural necessity given certain surrounding conditions, where exactly what these conditions are is not specified, indeed, may or may not be known. Usually the idea is, however, that whatever these conditions, they are for the most part realized in the contexts in which the law is used. The *ceteris paribus law*, then, makes reference to both kinds of probability that Dretske mentioned above. First, given the surrounding conditions to which it implicitly refers, it holds true with a probability of 1 in accordance with strict natural necessity. Second, the surrounding conditions to which it implicitly refers are themselves assumed to hold true with high statistical frequency.

But on the above reading of Dretske's definition of information, the second sort of probability is not involved. The frequency with which the channel conditions hold, relative to which a certain kind of signal bears information about a certain kind of fact, is not part of the definition of information. Suppose, for example, that many gas gauges are badly calibrated (indeed, they are) so that the gas tanks connected to them are half full when the pointer is on the one quarter mark, others when the pointer is on the three quarters mark, and so forth. In each case, when the gas tank is half full, no matter what it reads, the pointer carries the information that it is half full, relative to its own particular channel conditions. How often each of these various kinds of channel conditions holds is quite irrelevant. To be sure, Dretske often talks as if the relevant reference class in which *this* reading on *this* gas gauge should be put is restricted to those times when this very same gas gauge does or, counterfactually, would have given this same reading. Still, the assumption has to be that we are talking only about times when this very same gas gauge is found surrounded by the very same relevant channel conditions. Or suppose the reference class consists only of this particular reading on this particular occasion, the idea being just that if the tank had not been half full the pointer would not have pointed to this number. This way of thinking of the matter is in every way equivalent. The point is that the counterfactuals have to be run on the assumption that the relevant channel conditions still hold, and nothing has been said about how often conditions of this sort do hold in the world.

This is the only way I can see to interpret Dretske's definition and remarks on information quoted above. On the other hand, this way of interpreting Dretske's definition of information does seem to be inconsistent with certain things he says about "natural meaning," "natural signs," and "indication" in EB, despite the fact that he explicitly associates all three of these with signals that bear "information" in the sense of KFI (EB, p. 58). Dretske tells us, for example, that although otherwise such tracks would indicate quale, "[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 that it was a quale that made them" (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. And Dretske tells us that "[t]he red spots all over Tommy's face mean [natural meaning] 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" (p. 56). Contrast Fodor, who seems to use the term "information" more in the way we interpreted it above following the explicit definition in KFI. He says, "If the tokens of a symbol have two kinds of etiologies, it follows that there are two kinds of information that tokens of that symbol carry. (If some 'cow' tokens are caused by cows and some 'cow' tokens aren't, then it follows that some 'cow' tokens carry information about cows and some 'cow' tokens don't)" (1990, p. 90). Fodor also often speaks of "covariation" between represented and representation, which is plausible only if one imagines a reference to some one definite though unspecified channel of influence, making the signal depend nomically on whether *s* is F and vice versa. Fodor's usage fits not only Dretske's original definition but also a cautious physician's offering: "Those spots may mean Tommy has the measles, but they could also mean scarlet fever. I think we had better take a culture." Dretske's modified claim, that if some people with spots like that don't have the measles then those spots don't mean measles, apparently refers instead to statistical frequencies at the source.

Alternatively, perhaps it refers to the frequency of certain channel conditions. It might well be, for example, that given certain channel conditions, only measles virus would cause spots like that, but that given other channel conditions, only strep bacteria would. Just as, given certain channel conditions, only a half full tank of gas would cause that reading, but given other channel conditions, only a quarter full tank would. Then by Dretske's original definition, Tommy's spots might mean measles even if on another child they would mean scarlet fever. But if Dretske's modification here involves assigning certain channel conditions themselves a probability of one, such a probability would also seem to be merely a statistical frequency.

Indeed, both Dretske's KFI and his EB waver at points between the two kinds of probability in discussing information. Dretske tells us, both in KFI and in EB, that if his doorbell rings, that carries the information that someone is at the door. But in EB we are told:

> "It is partly the fact, presumably not itself a physical law, that animals do not regularly depress doorbells . . . that makes a ringing doorbell *mean* that some *person* is at the door. . . . As things *now* stand, we can say that the bell would not be ringing if someone were not at the door. It therefore indicates or means that someone is at the door. But this subjunctively expressed dependency, though not a coincidence, is not grounded in natural law either. . . . Normally, though, these things don't happen. . . . And this is no lucky coincidence, no freaky piece of good fortune. . . . There must actually be some condition, lawful or otherwise, that explains the persistence of the correlation . . . [for the doorbell to indicate a person].

But, of course, if the condition that explains the correlation is not lawful but "otherwise," then it is grounded in mere facts about the state conditions characterizing the world at certain times and places – either conditions at the source or existent channel conditions. It has the status merely of a local statistical frequency – based lawfully, perhaps, hence explainably, upon prior local statistical frequencies, but that does not change its essential nature as merely a statistical frequency.

The vacillation here seems to be twofold. First, it concerns whether or not mere statistical frequencies at the source, rather than strict natural law, should be allowed to determine signals as bearing "natural information." Second, it concerns whether we should count a signal that is not univocal except as harnessed to a particular information channel. But, of course, most of the interesting examples of signals carrying "information," defined Dretske's original way, are of a sort that either do not always carry the same kind of information (because channel conditions vary) or if they do, that is a matter of convenient empirical fact, not natural necessity. The fact that a signal carries "information," defined Dretske's original way, has no bearing whatever upon whether, by the mere fact that the signal has arrived, one can tell anything about *what* information, if any, it carries.[3]

I propose to stay for a while with Dretske's original definition of natural information. To my knowledge, no other well-defined notion of natural information is currently available. Allowing merely statistical

3 Dretske worries about something close to this in KFI, pp. 111–23, but he does so in the confusing context of worrying about what "knowledge" is, and thus he never confronts the basic problem. – Or so I would argue, but my main project here is not Dretske exegesis.

considerations on board poses an intractable problem concerning the reference classes within which the frequency of 1 should be required to hold. Do spots like that mean measles if small pox, though now extinct, used to, and may in the future, cause spots like that? If the Skinner-trained pigeons in your neighborhood start pressing doorbells, how close may my neighborhood be to yours for my ringing doorbell still to carry the information that a person is at my door? More important, mixing frequencies with natural necessities muddies the issues involved in trying to understand phenomena connected with intentional representation. These issues can be seen much more clearly if we separate issues of natural law from issues that concern mere frequencies. For clarity, I will call natural information purified of all mere frequencies, natural information as originally defined by Dretske, "informationL" (for "law").

InformationL is an entirely objective commodity and it is ubiquitous. Often its channels are complex, and such as seldom if ever to be duplicated. Channels that are often duplicated tend to be fairly simple channels, such as reflections in calm water. Channels carrying reflections in choppy water, though not much more complex, are seldom repeated. The more numerous and irregular the intervening media between source and signal are, the less likely repetition becomes.

InformationL is everywhere, but the problem, of course, is to interpret it. For no signal that makes up only part of the world can carry the informationL that its own channel conditions hold. And that means that it cannot carry the information that it carries informationL, nor what code this information is in. This opens the question why an organism could possibly care whether or not it ever encounters any of this ubiquitous but uncommunicative informationL. What good will it do an animal to have informationL?

The problem is twofold. First, a signal carrying informationL is, as it were, in code. It is of no use to an organism unless the organism can "read" the code. Second, the informationL that reaches an organism is not all in the same code.

Consider first the easy problem, that of reading the code. Suppose that the information all arrives in the same code. Then for a signal to be of use to a creature – to be "read" by it – it would only be necessary that the creature should be guided by the signal in a way that diverts it from activities less likely to benefit it to ones more likely to benefit it, this likelihood being contingent on the fact conveyed by the signal. For example, if the fact conveyed is the relative location of water, given that the creature is thirsty, all that is needed is that the signal should cause the creature to turn toward the location indicated. The beneficial activity need not, of course, be overt. It might be an inner state change. The basic idea here is well known, I believe, and has been given numerous expressions, for example, by Dretske and myself.

The "same code" problem is the harder one, and is itself two-sided. First, we have not yet offered a reason to suppose that informationL about the same thing always or ever reaches an organism in the same code or packaging. Second, we have offered no reason to suppose that the same packaging always or ever carries the same informationL, indeed, any informationL. Why suppose, for any signal that the organism receives, that all signals of that kind reaching the organism, carry the same informationL. But for the organism to be able to use the informationL it receives, the same kind of informational content needs to affect the organism in the same kind of way, and different kinds of informational content need to affect it in different ways. Information about the same must, as it were, look the same to the organism, and information about different things must look different to the organism. (This may put us in mind of Fodor's "formality constraint" [1980].)

A central tenant of contemporary ecological psychology of the sort introduced by J. J. Gibson is that there is far more consistency in the natural information received by an organism than was formerly supposed. The claim is, first, that if you look for the right aspect of the signals that arrive by way of the ambient energy surrounding an organism, you find that a surprising number of superficially or apparently different channels of informationL can be described as really being the same channel once you have located the right high order invariances in the signals. And it is these invariances, these univocal codes,[4] that the evolving animal has become sensitive to, so as to "pick up" the relevant information and use it.

Second, the Gibsonian claim is that the very same relevant channel conditions are present under environmental conditions that the animal frequently or nearly always finds itself in, or that it knows how to maneuver itself into. In the animal's normal environment, the relevant channel conditions are always the same, or always possible for the ani-

4 Gibsonians protest that the natural information used by organisms is not in the form of a "code." Their point, however, is merely that it is constituted by changing energy structures that do not require translation into some other medium in order to be used by the organism.

mal actively to intercept, so that relevant features of the source lawfully produce informationL about themselves in the same code. There are "ecological laws" such that the signals covary with the relevant environmental features.

Third, the Gibsonian claim is that informationL of this sort that is relevant to the animal's needs is much more complete than had previously been supposed. Information about exactly those environmental conditions to which the animal needs to adjust is frequently presented in an unequivocal way. "The stimulus is not impoverished."

These three claims are not generally separated in the tradition of Gibsonian psychology, but they are independent. Gibsonian "information" is not only informationL, but also lawfully carries complete information needed for guidance with respect to important aspects of the environment, and is frequently present in the environment, coming in always through the very same information channels, that is, exemplifying the very same *ceteris paribus* laws, arriving in a single code. All the animal has to do is to tap into these rich sources of information (for example, by developing eyes with lenses) and funnel them directly into guidance of appropriate behavior.

Mechanisms by which various perceptual constancies are achieved, such as recognition of same color, same shape, same size, same voice, and so forth, through a wide spectrum of mediating conditions, insofar as these constancies are sometimes detected over wide ranges of input in accordance with univocal principles, illustrate the use of Gibsonian information. Then it is a very complex signal indeed, one in which the significant invariances are (from the physicist's point of view) highly derived, that yields informationL through a complicated but still univocal channel in a single code. The job of tapping into such informationL channels and using the information to guide useful action is, as a biological engineering feat, extremely challenging. Yet natural selection has managed to solve many of these problems.

Surely there does exist in our world at least a certain amount of Gibsonian information, or at least very close to that, which serves as the bedrock foundation making sentient life possible at all. This foundation guides the most basic immediate responses to the environment of all animals, and also supports all information-gathering activities and faculties that make use of less tractable, less user-friendly, forms of information that are also naturally found in the environment. But there is also that useful information in the environment that is not fully Gibsonian.

InformationL becomes less Gibsonian, for example, as it becomes less ubiquitous. Darkness and dense fog, for example, impede transmission of normal visual information. InformationL becomes less Gibsonian as it arrives in more alternative packagings, in alternative codes. For example, we use a number of alternative visual cues for depth. More interesting are cases in which the same signal form varies in the information it carries. Consider light and sound when reflected off smooth surfaces. Like a gas gauge that carries informationL but reads "¼" when it is half full, reflections carry perfectly good informationL but informationL that needs to be read differently than usual. A puddle in the woods is not a hole in the ground with upside down trees hanging inside. Animals, after brief exposure, generally treat reflections simply as irrelevant noise in the data, holes in the normal flow of information. But a kitten's first experience with a mirror can be very amusing to watch, a dog will bark at its own echo, sometimes for hours, and a Canada goose once spent a whole afternoon doing a mating dance to his reflection in the basement window of our building on the Connecticut campus. We humans, on the other hand, are able to tap many such sources of informationL and to read them correctly. We can comb our hair in the mirror, we understand that Clinton is not inside the TV set nor our friends inside the telephone. We build gadgets to collect thousands of different kinds of informationL – various indicators, meters, gauges, scopes, audios, videos, and so forth – and we learn to read them correctly.

When a variety of channels of informationL about the same are intermittently available to an organism, the animal must understand when each is open, distinguishing informationL both from mere noise and from informationL arriving in similar vehicles but differently coded. Nor should we take for granted that an animal can integrate the sources of informationL that it uses. There is a story circulating (though probably apocryphal[5]) that certain venomous snakes strike mice by sight, trace the path of the dying mouse by smell, and find its head (so as to swallow it first) by feel, and that none of these jobs can be done using any other sensory modality. The lesson is, anyway, logically sound. InformationL about the same that comes in a variety of codes requires "translation" if it is to be used in a versatile way.

Suppose then that informationL about the same things arriving through a variety of media is translated by mechanisms in the organism

5 The original source seems to be the zoologist Sverre Solander, who gives no references and, despite requests, has offered no data yet to my knowledge.

into a common code.[6] Insofar as this result is achieved, whatever appears in that code is correlated always in the same way with the same source or kind of source in the environment, even when the channels that control this effect are variable. In this way, a great deal of informationL that is not fully Gibsonian as it originally reaches the organism may be translated into the practical equivalent of Gibsonian informationL inside the organism. As I will now argue, however, relatively few things that an animal needs to know can be communicated in this direct way.

InformationL depends on a channel between the information source and the signal producing a correspondence between the two in accordance with natural necessity. But unfortunately, relatively few things that an animal needs to know about can figure as sources for this kind of information. The mouse, for example, needs to know when there is a hawk overhead, but there are no natural laws that apply to hawks overhead and hawks only. The existence of hawks is not a matter of law, nor, for any given channel, is the nonexistence of things other than hawks that might cause the same effects as a hawk on the output of that channel a matter of natural necessity. Similarly, if there are channel conditions under which cows cause mental "cow" tokens as a matter of natural law, surely there can be none under which mental "cow" tokens are caused by cows. They might instead be caused by something that looked like a cow, or sounded like a cow, or smelled like a cow, or all three, but that wasn't a cow. It is the *properties* of objects like hawks and cows that enter into natural laws, not the hawks and cows themselves, and it is never a matter of natural law that only hawks or cows have these properties.

There is, of course, an old-fashioned way out of this difficulty. You can argue that it is a matter of nominal definition that cows and only cows have certain properties, and then argue that information concerning the copresence in one and the same object of all these defining properties could indeed be transmitted through an information channel. Then there might be natural informationL about the presence of a cow. As a preliminary, however, first notice that you can't take this route for information concerning individuals. Even quite primitive animals are often able to recognize and keep track of various of their conspecifics individually, to learn things about them, and so forth. But there are no

6 By "translated into a common code," I mean only that sameness or overlapping in content is marked.

laws that concern any individuals as such. No signal can carry the informationL that it is *Tommy* who has the measles. Second, although a classical position that some still occupy gives natural kinds such as gold and water definitions in terms of necessary and sufficient characteristics, it is no longer plausible that biological kinds, such as cow, can be defined that way. A large proportion of the kinds that we name in everyday speech are "historical kinds," kinds that are not defined by their possession of certain properties at all, but instead through "historical" connections – connections in the spatial/temporal/causal order – that their members have to one another (Millikan 1999, Chapter 2 above). Exactly as with individuals, these kinds cannot be subjects of informationL. They fall under no laws, not even ceteris paribus laws, and they support no counterfactuals.

Thus we are returned to the problem addressed earlier when Dretske observed that it is not a matter of natural necessity that your ringing doorbell "indicates" there is some person at the door. In what sense of "natural information" then, exactly, does the doorbell carry natural information? Is there a way to define a softer notion of "natural information" to do the work required here?

To answer this we must have firmly in mind what work it is that is required. What do we need a theory of natural information for? In this context, we require it to support a theory of "intentional" representation, in the sense introduced by Brentano. This is the kind of representation that displays Brentano's mark of the mental. Intentional representations can represent nonexistent things, for example, nonexistent facts. They can be misrepresentations. All agree, of course, that natural information is not itself intentional, that it cannot misrepresent or be false. "Informational semantics," as Fodor calls it, is an attempt to show how, despite this difference, intentional representation still rests at base on natural information.

How to move from a theory of natural information to a theory of intentional representation is, however, a problem. That is what Fodor's theory of "asymmetrical dependency" is designed to do (1990, Chapter 4). And that is what Dretske's addition of teleology is designed to do – his claim that it is only a function, not always realized, of intentional representations to carry natural information (1981, 1991). Fodor's asymmetrical dependency theory seems, quite explicitly, to rest on informationL, but I won't argue that case here. Rather, I will try to show how teleology can be combined with a theory of soft natural information to produce the variety in forms of intentional representation that animals require. But there has been some confusion about the relation of teleological accounts of intentionality to informational semantics. So let me first remark on that relation.

Naturalized teleological theories of the content of representations are attempts to explain Brentano's mark of intentionality: How can representations be false or represent nonexistent things? But teleological theories are only overlays, minor additions, veneers superimposed, on prior underlying theories of representation, and there can be considerable variety among these underlying theories. When looking at any teleological theory, the first thing to ask is on what kind of more basic theory of representation it rests.

Suppose, for example, that you think of mental representations as items defined in a classical functionalist way, in accordance with patterns of causal/inferential dispositions. And suppose that you have a theory that tells what dispositional relations one of these representations must have to others, and the collection as a whole to the world, for it to be a representation, say, of its raining. Then the teleological theorist, call her Tilly, will come along and point out that surely some of the causal roles of actual representations in actual people's heads correspond to bad inferences. What you must say, says Tilly, is that what the representation represents is determined by what its causal role *would* be if the head *were* operating correctly, that is, in the way it was designed, by evolution or learning, to operate. Similarly, suppose that you think of mental representations as items that "stand in for" the things they represent, running isomorphic to them, with differences in the representations producing differences in the behaviors guided by them, thus making the behaviors appropriate to the presence of the things represented. Then Tilly will come along and point out that some representations are false, that is, not isomorphic to things in the world as required to guide behavior appropriately. What you must say, says Tilly, is that the representations represent what would be in the world, running isomorphic to them, if the cognitive systems were operating correctly. That is, what a teleological theory of content does is to take some more basic theory of content, point out that the application of that theory to actual creatures requires idealizing them in certain ways, and then offer the teleological principle to explain which idealization is the right one to use in interpreting intentional contents, namely, the one that fits how the cognitive systems were designed or selected for operating. You give your naturalistic analysis of what a true or correct representation is like, and Tilly merely adds that systems designed to produce true representations don't always work as designed, claiming that correctness in perception and cognition is defined by reference to design rather than actual disposition.

Accordingly, the teleologist who is an information semanticist begins with the idea that representations are signals carrying "natural information" and then adds teleology to account for error. My claim is that adding teleology to informationL will not yield the rich variety of intentional representation that either we or the animals employ, but that there is a softer kind of natural information that does underlie all intentional representation. This softer kind, however, offers no help whatever to the verificationist.

Let us return, for a few moments, to the animal whose perceptual/cognitive systems are capable of translating informationL about the same things arriving through a variety of media into a common code. Whatever appears in that code is correlated always in the same way with the same source or kind of source of informationL in the environment. But, Tilly reminds us, it is not plausible that errors will never occur. If this arrangement has been built by natural selection, however, it will at least be a *function* of these mechanisms, which tap into and converge these channels of informationL, to produce signals that carry informationL in a univocal code. Their function is to transmit signals that are controlled by certain external sources of information so that these sources then control the behavior of the organism in ways that are adaptive. Surely this is the sort of thing that Dretske had in mind in saying that the function of a representation is to indicate (1986, 1991). Or, being very careful, what has really been described here is not the function of the representations themselves, but the function of certain mechanisms that produce representations. The first job of such a mechanism is to complete a specific type of channel of information flow, or to bring to focus in a single code a number of such channels, so as to produce an informationL-bearing signal in a specific code. This is the way to add teleology to the idea that intentional representation is, at root, natural informationL. False intentional representations result when such a mechanism fails to perform this job properly.

I say that I think this is what Dretske has in mind. Dretske has sometimes wavered, however, on whether it can be a function of information gathering systems to gather information about affairs that are distal to the organism. I will explain.

The job of bringing information arriving through different channels, perhaps through complex media, in different codes, to a focus is obviously difficult and very risky. Tilly is surely right that systems responsible for accomplishing this feat inevitably will sometimes fail. Recall the Canada goose in love with itself, and the dog trying to communicate with its echo. When this sort of thing happens, however, it is not usually because there is anything wrong with the organism. Without doubt, perhaps definitionally, almost none of the mistakes in informationL gathering that are made by healthy animals are due to malfunction of the animals' informationL-focusing systems. Mistakes are due to an uncooperative environment, which fails to supply those informationL channels that the animal has been designed or tuned to recognize and employ. Gibson to one side, concerning some informationL that an animal needs to gather, the environment may be rife with decoy channels, nor is there anything the animal can do about that, perhaps, without evolving completely different perceptual systems. Both Dretske (1986) and Neander (1995) have concluded from this, however, that the information-gathering systems of animals may not actually have the function of gathering information about *distal* affairs at all. The argument is that when representations of distal affairs are apparently mistaken, since typically this is not because the animal's information systems are failing to function properly, it must be that these systems do not have as their function to gather this kind of information. Neander then seriously claims that all representation must be only of proximal stimuli. The effect, of course, will be a very strong form of verificationism indeed. The organism can only represent what it can verify conclusively, granted it's not sick or damaged.

But the idea that nothing can have a purpose or function that it requires help from anything else to achieve is mistaken. Consider the can opener on the wall in my kitchen. It is not now opening cans. It is not now performing its function. It would need my help in order to do that. Certainly it doesn't follow that it is malfunctioning, or that opening cans is not its function.[7] In the case of information-gathering systems, exactly as with can openers or, say, with the famous walking

7 If, however, you do insist, as Neander does, that in the *ordinary* sense of "function" things really can't have distal functions, then I refer you to the definition of "proper function" stipulated in Millikan (1984), in accordance with which most of the many proper functions that most biological items have are distal, and I suggest that the notion of function we need to use to gain insight here is "proper function" as there defined.

mechanisms in cockroaches, a cooperative environment plays a lead role in helping them serve their functions. (Nor, of course, does it follow that it is the environment's function to help cockroaches walk or to help us focus information.)

Let us now look more closely at the result of adding teleology to natural informationL to produce intentional representation. The first job of a system that uses informationL to produce representations is to complete a specific type of natural-informationL channel so as to project that informationL into some standard code. But systems of this kind also have jobs beyond. The codes into which they translate informationL must be ones that the behavioral systems of the animal are able to use. The problem, posed first during evolutionary development, then to the developing individual animal, is to coordinate these two kinds of systems. Suppose, however, that the representation-producing systems and the behavioral systems fail to cooperate on some task. Suppose that a signal carrying informationL about one state of affairs is used by the behavioral systems in a way appropriate instead to some contrary state of affairs. For example, the informationL that the height to be stepped up is, say, eight inches, is coded in a representation that guides the legs to step up only seven inches. Which has erred, the perceptual side of the system or the motor side of the system? Is the representation wrong, or is its use wrong? Has the message been written wrong, or has it been read wrong? What does the *intentional* representation say, eight inches or seven inches?

Notice that the signal, as carrying informationL, definitely says eight inches. Compare the informationL carried by a miscalibrated gas gauge. The miscalibrated gauge carries informationL telling the *actual* level of the gas in the tank. If we interpret it wrongly, that does not make it carry the informationL we wrongly take it to carry. What it itself naturally means just is whatever it *actually* carries informationL about, even though in a difficult or uninterpretable code. In the same way, the coded informationL about the height of the step cannot be wrong. The attributes *right* and *wrong*, *true* and *false*, don't apply to the code considered as a natural sign.

Recall that a signal carries informationL, not as considered within the reference class of all items in the world having the same physical form, but only as a member of the class of signals linked to sources through the same kind of information channel, that is, in accordance with the same natural necessities implemented through the same mediating conditions. As an *intentional* representation, however, the representation of the height of the step is a member of a different reference class altogether. It is a member of the class of all representations like it in form,[8] produced by the same representation-producing systems, for use by the same representation-using systems. In this class there may also be representations identical to it but that carry natural informationL in a different code, and representations that carry no natural informationL at all. In which code, then, is its *intentional* content expressed?

Exactly here is the place to apply teleology, as I see it, to the analysis. We suppose that the system that codes and uses the information about the step is a system where the coding and using parts of the system have coevolved, either phylogenetically and/or ontogenetically. During evolution of the species and/or during learning or tuning, they have been selected or adjusted for their capacities to cooperate with one another. The operative features of both halves of the system have been selected for and/or tuned as they have because these features and settings have sometimes succeeded in guiding behavior appropriate to the informationL encoded. If this is so, inevitably it is true that these coordinations were achieved by settling on some single and quite definite code. Only if there was constancy or stability in the code employed by the representation maker and user could coordinations have been achieved systematically. It is this code then that the representation producer was designed to write in, and it is this code that the representation user was designed to read. And it is this code that determines the *intentional* content of the message about the height of the step. In any particular case of error, whether it is the representation producers or the representation users that have erred depends on whether or not the natural informationL appears in this code.

My proposal is now that we should *generalize* this result. Intentional representations and their producers are defined, are made to be such, by the fact that it is their job to supply messages that correspond to the world *by a given code*. That is the essence. But notice that *that* formulation makes no reference to informationL. If that is the essence of the matter, then the *mechanisms* by which the producers manage to produce messages that correspond by the given code drops out as irrelevant to their nature as intentional representation producers. If there exist systems with the function of supplying messages that correspond to the world by a given code but that manage to achieve this result, when successful,

8 More accurately, the class of all representations that the systems designed to use it are designed to identify as having the same content.

without tapping into any channels of natural informationL, they too will be producers of intentional representations. They will be producers of intentional representations that are not defined with reference to natural informationL. I will now argue that such systems do exist, indeed, that the bulk of our mental representations necessarily are of this type. Rather than informationL, they tap into channels of softer natural information. How should we define this "softer" form of natural information?

Dretske wishes to eliminate *de facto* perfect correlations that are "lucky coincidences" or "freaky piece[s] of good fortune" as possible supports for any notion of natural information. But does anything stand in the middle between, on the one hand, statistical frequencies resulting from lucky coincidence and, on the other, the necessity of natural law? The answer Dretske gave to this question, though inadequate, I believe, is still a very interesting one. He said, "[t]here must actually be some condition, lawful or otherwise, that explains the *persistence* of the correlation" [emphasis mine]. About this I remarked earlier that the fact that a local statistic is based lawfully upon prior local statistics, hence that a correlation is explainable, does not alter its nature as a mere statistical frequency. If the frequency of black balls in the urn today is 1, and if nothing disturbs the urn, then by natural necessity it follows that the frequency of balls in the urn tomorrow is 1. That does not change the probability of being black if a ball in the urn into a probability of some kind other than mere statistical frequency. It does not help being-a-ball-in-the-urn to carry the informationL being-black.

But it does do something else. It explains how, by sampling the urn today and adjusting my expectations of color accordingly, this adjustment in expectation can turn out to be adequate to my experience tomorrow, *not by accident but for good reason*. Many statistical frequencies *persist* over time in accordance with natural necessity, and many produce correlate statistical frequencies among causally related things, in accordance with natural necessity. If measles are producing spots like that in this community today, then measles will probably be producing spots like that in this community tomorrow. Measles, after all, are contagious. And if a nose like that is correlated with the presence of Johnny today it will probably be correlated with the presence of Johnny tomorrow. Johnny's nose, after all, tends to sustain both its shape and its attachment to Johnny. There are no laws that concern individuals as such, but there are many kinds of local correlations that do. Notice, however, that whether the persistence of a correlation may be explained in this sort of way does not depend on its being a perfect correlation. Conditional probabilities of 1 have nothing to do with the matter.

This yields a way that an organism may come to possess systems that produce representations that correspond to the world by a given code *often enough* to have been selected for doing that job, but that do this job without tapping into any natural informationL. Systems of this sort run on bare statistical frequencies of association – on correlations – but on correlations that persist not by accident but for good reason. Probably these correlations typically obtain between properties of the not-too-distant environment that *do* supply informationL to the organism, and more distal properties, kinds, situations, individuals, and so forth, of interest to the organism but that *don't* supply it with informationL. The intentional contents of representations of this sort are determined not by any natural informationL that it is their function to carry, but merely by the codes in which their producers were selected to write, so as to cooperate with the systems designed to read them.

It follows that a representation producer, basing its activities on past local statistical frequencies, may indeed be representing Xs, and yet be unable perfectly to distinguish Xs from Ys. It may have no disposition under any conditions infallibly to distinguish Xs from Ys. *To perform properly*, its representations of Xs – its code tokens of a certain type – must correspond to Xs, but this does not entail that there exist any information channels at all, actual or possible, through which it could infallibly discriminate Xs from Ys. That having grown up with gray squirrels around, I am thinking of gray squirrels has nothing to do with whether I can discriminate gray squirrels from Australian tree possums, even if someone introduces tree possums into my neighborhood. Similarly, the determinacy of content of my representation of cows is not threatened by the possibility of a new species arising that I couldn't distinguish from cows, or by the possibility of Martians arriving with herds of facsimile cows. The alternative that I should sometimes actually be at the other end of an informationL channel from cows is not even coherent.

Consider, in this light, Pietroski's tale about the kimus and the snorfs (1992). The snorfs are attracted by the red morning glow over their local mountain so that they climb up it each day. Thus they conveniently avoid their chief predators, the snorfs, who don't take to mountain terrain. Pietroski claims that since no current kimu would recognize a snorf if it ran into it head on, it is implausible that the perception of red means snorf-free terrain to the kimus. A mere correlation between the direction of the red glow and the direction of the snorfs is not enough to support intentional representation. Now first, we should note that the injection of phenomenology here is perversely distracting. The question is not whether a red qualia, should there exist such things, could mean *no snorfs this direction* rather than *red*. Bats perceive shapes by ear and, goodness knows, maybe squares sound to them like diminished seventh chords do to us. Pietroski's question should be whether *any* inner representation that merely directs the kimu toward the sunlight could represent for it the snorf-free direction. Nor should the idea be that the kimu reads or interprets the inner representation as meaning "the snorf-free direction" the way you or I would interpret a sign of snorfs. To interpret a sign of snorfs, you or I must have a prior way of thinking of snorfs, and that, by hypothesis, the kimus do not have. The question, put fairly, is whether something *caused* by red light could *constitute* an inner representation of the snorf-free direction for the kimus. Also, we should be clear that the kimus' sensitivity to and attraction by the red light is not supposed to be accidental, but is a result of natural selection operating in the usual way. Kimu ancestors that were not attracted to red light were eaten by the snorfs.

Put this way, the situation is parallel to that of certain tortoises, who are attracted to green things, because green correlates with edible vegetation. They will move on the desert toward any green seen on their horizon. Nor do the nutritious properties of the vegetation produce the green light. These properties are merely correlated with green light. Can the green mean "chow over there" to the tortoise? Obviously not in so many words. But your percept of an apple doesn't mean "there's an apple over there" in so many words either. If the green doesn't mean chow over there to the tortoise, then what on earth *could* mean chow over there to anyone? Is it really plausible that there could be a genuine informationL channel open to any of us, for you or for me, that would communicate the informationL that there was chow on the table? Does human chow, as such, figure in any causal laws? If not, then in what sense are we "able to discriminate" when it is chow time? Unless, that is, we rely on mere statistical correlations.

Besides natural informationL, then, we should recognize another equally important kind of support for intentional representation, resting on what may also be called "natural signs" carrying – to keep the terminology parallel – "informationC" (for "correlation"). Natural signs bearing informationC are, as such, instances of types that are correlated with what they sign, there being a reason, grounded in natural necessity, why this correlation extends through a period of time or from one part of a locale to another. One thing carries information about another if it is possible to learn from the one something about the other not as a matter of accident but for a good reason. But no vehicle of information is transparent, of course. How to read the information through its vehicle has to be discovered, and it has to be possible to learn this in an explainable way, a way that works for a reason. The vehicle carries genuine information only if there is an ontological *ground* supporting induction that leads from prior experience to a grasp of the information carried in new instances. There must be a connection between the various instances exhibiting the correlation, a reason for the continuation of the correlation. Correlations that yield true belief only by accident do not carry genuine information.

Natural signs carrying informationC are correlated with what they represent because each sign instance is connected with what it represents in a way that recurs for a reason. Typically, however, the correlations are not perfect, and informationC, like informationL, cannot be false by definition. A token indistinguishable from a natural sign but that is not connected in the usual way with its usual represented is not a natural sign. The correlations that support informationC may be weak or strong. For example, a particular instance of a small shadow moving across the ground is a natural sign carrying informationC that a flying predator is overhead if it is actually caused by a flying predator, but the correlation that supports this natural signing, though it persists for good reason, may not be particularly strong.

If we allow ourselves to use the term "natural information" to cover informationC as well as informationL, then, we must keep firmly in mind that this sort of natural information has nothing to do with probabilities of one. Nor does the presence of this kind of information directly require the truth of any counterfactuals. If a shadow is a natural sign of a predator it does not follow that if a predator weren't there a shadow wouldn't be there, hence that such shadows can be used to discriminate predators from nonpredators. Nor does it follow that if a shadow weren't there a predator wouldn't be there – not on a cloudy day. Thus it is that a creature can perfectly well have a representation of Xs without being able to discriminate Xs from Ys.

References

Atran, S. 1989. Basic conceptual domains. *Mind & Language* 4, nos. 1 and 2, 7–16.

Barsalou, L. W. 1987. The instability of graded structure: Implications for the nature of concepts. In Neisser 1987, 101–40.

Bartlett, F. C. 1932. *Remembering*. London: Cambridge University Press.

Billman, D. 1992. Modeling category learning and use. In Burns 1992, 413–48.

Block, N. 1986. Advertisement for a semantics for psychology. *Midwest Studies in Philosophy*, 10, 615–78.

Blum, P. 1998. Different structures for concepts of individuals, stuffs and real kinds: One Mama, more milk and many mice. *Behavioral and Brain Sciences* 21.1, 66–7.

Boyd, R. 1991. Realism, anti-foundationalism and the enthusiasm for natural kinds. *Philosophical Studies* 61, 127–48.

Boyer, P. 1998. If 'tracking' is category-specific a 'common structure' may be redundant. *Behavioral and Brain Sciences* 21.1, 67–8.

Bridgeman, B. 1982. Multiplexing in single cells of the alert monkey's visual cortex during brightness discrimination. *Neuropsychologia* 20, 33–42.

Burge, T. 1979. Individualism and the mental. In P. French, T. Uehling, and H. Wettstein (eds.), *Studies in Metaphysics: Midwest Studies in Philosophy* IV, 73–121. Minneapolis: University of Minnesota.

Burge, T. 1982. Other bodies. In A. Woodfield (ed.), *Thought and Object*, 97–120. Oxford: Clarendon.

Burge, T. 1986. Individualism and psychology. *Philosophical Review* 95, 3–45.

Burns, B. (ed.) 1992. *Percepts, Concepts, and Categories: The Representation and Processing of Information*. Amsterdam and New York: North-Holland.

Byrnes, J. P. and Gelman, S. A. 1991. Perspectives on thought and language: Traditional and contemporary views. In Gelman and Byrnes 1991, 3–27.

Calvert, G. A., Bullmore, E. T., Brammer, M. J., Campbell, R., Williams, S. C., McGuire, P. K., Woodruff, P. W., Iverson, S. D., and Davis, A. S. 1997. Activation of auditory cortex during silent lipreading. *Science* 276, 593–96.

Campbell, J. 1987/88. Is sense transparent? *Aristotelian Society Proceedings*, New Series. Vol. LXXXVIII, 273–92.