

Crawford L. Elder
Department of Philosophy, U-54
University of Connecticut
Storrs, Connecticut 06269-2054
U. S. A.

What vs. How in Naturally Selected Representations

True judgments are about things that really exist in the world, and represent them as being ways they really are. Do false judgements stand in the same relation of *representation* that true ones do, but to non-obtaining states of affairs? Are empty judgements *about* in just the same way, but about non-real objects? The first task of a naturalistic understanding of thought and talk is to find ways to answer No to both questions. The most promising implementation of naturalism, most now would agree, sees false judgements, and empty ones, as simply judgements which fail to do what biologically they are supposed to do: any judgement is supposed, by virtue of the biological proper function of the device or program that generates it, to be about the real and to represent it as it is, but tokens of a given biological type can quite well fail to perform the function that is proper to them.

But if accuracy-about-the-real is truly a *biological* “supposed to”, there is no reason to assume that the devices which generate judgements operate infallibly—not even in the historical circumstances in which they evolved. Evolution, after all, is content to select devices which succeed merely more often, or with fewer detrimental side effects, than those turned out by the historical alleles. So a naturalist should never assume that how a judgement-producer cues judgements of this or that semantic type tells the last word on the semantic content of judgements of that type. On the contrary, a naturalist should be ready to suppose that how a representation-producing device judges that *p*—that a fly is here now, that a mouse is near, that there is nectar close by to the north—is a rough and ready method, and often a bad guide to the *p* thus judged. The content of any representation may be more fine-grained, or distal, than the signs and marks to which the mechanism gears its tokenings of that representation.

Most current attempts to see how well naturalism can work, and how much of human belief and desire it might be brought to cover, have begun with the ostensibly easier cases of representations which natural selection has installed in simple organisms, causally downstream of characteristic sensory cues, and causally upstream of routinized behaviors. They have begun, for example, with the signal which triggers the frog's gulping at passing flies. But in asking about the content of even these naturally selected representations,¹ philosophers have time and again held the verificationist view that how their tokening is cued tells us the last word on what they represent. For example, many have held that the gulp-triggering signal in frogs—which the frog's visual system produces quite indiscriminately, whether confronted with a fly or a falling BB, or with any other small, dark, moving speck—has as its content “fly or BB”, or even “small, dark, moving speck” (Neander 1995; Neander 1996; Goode and Griffiths 1995; Agar 1993; cf. Rowlands 1997).² Thereby they make it very hard for the frog's visual system to be fooled, to make mistakes (Neander 1995: 130-33). They thus make naturalism about human judgement seem poorly equipped to shed light on how human judgements manage to misrepresent.³ And they do so, moreover, in contravention of what is truly biological about a biological naturalism. What leads to this self-frustration of the naturalist program, I shall argue, are mistaken views about causal explanation. The aim of this paper is show how a different view can redeem the promise of biological naturalism, and avoid projecting verificationism onto the representations which natural selection installs.

I

How does causal explanation fit into the debate over the semantic content of naturally selected representations?

Most contributors to the debate would say: because in asking about the content of such a representation, we are asking what it is the *proper function* of such a representation to track or to signal (Neander 1995; Akins 1996; Goode and Griffiths; Fodor 1990: 70-79; Dretske 1988: 68-70; Rowlands).⁴ And proper function is, indeed, a concept which has causal explanation at its core. On the picture framed

by Millikan and others, the proper function of an existing biological device (token) is what ancestors of the device—earlier devices (tokens) from which it was reproduced—historically did, that causally explains why they were reproduced more widely than competitors (Millikan 1984; Millikan 1989a; Neander 1991). Thus the proper function of a device (type) is what it has been selected for having done—what it is “supposed” to do.

But could it be the proper function of routines in the frog’s visual system to signal the presence of nearby flies (or of food or of small dark specks)—or the proper function of routines in the hoverfly to signal the proximity of mates? I agree with Millikan (1989b; 1986; cf. Akins) that it could never be the proper function of a biological device merely to signal some feature in the ambient world. For signaling, by itself, confers no selective advantage. A signal helps just to the extent that it gets heeded by other routines (“consumer” devices, Millikan calls them), typically in the same organism. The neural signal triggered by a passing fly, for example, helps just because it switches on and steers the frog’s motor routines for striking and gulping—and does so in a way that helps them succeed at bringing in nutritious items. The proper function of a representation-producing device, then, is to orient representation-consuming devices so as to enable *them* to perform *their* proper function(s). Mapping an outside feature comes in mainly as *how* a representation-producer has historically performed the job that led to its replication, not as *what* that job has been. For in general, the only strategy for switching on and steering motor routines that is simple enough for evolution to have hit upon, and useful enough for evolution to have liked, is one that gears the switching and steering to certain features of the surrounding world.

Still, even if signaling passing flies (or dark specks or bits of food) is not *what* some representation-producing device was selected for, it is *how* that device has performed its proper function on the historical occasions that have led to its proliferation. So there is some outside feature which any such representation-producing device is *supposed to* map or track. It is just that outside feature which was causally responsible, in the specific historical cases that led to the device’s replication, for the successes which the consumer mechanism historically scored when acting as steered by the representations. This, then, is where causal explanation connects with the semantic content of a naturally

selected representation. What such a representation represents is a state of affairs which causally explains why its steering was effective, on the occasions when it was effective—why the consumer device, acting as steered, succeeded on those occasions.

II

So what *does* the frog's famous neural signal (call it "*S*") represent? (To keep the discussion focussed, I shall temporarily stick to this shopworn example; other popular examples will be discussed in the fourth section.) It represents the presence, at just *that* angle in front of the frog, of food, of something nutritious—or, what says nearly the same, of a Gibsonian affordance of eatability.⁵ For only the presence of something nutritious, at just that angle in front of the frog, can causally explain why, when the frog made a gulping strike at just that angle, what resulted was success for the gulping routines, i.e., the presence of something nutritious in the frog's digestive system. This answer assumes of course, that the proper function (at least, *a* proper function) of the gulping routines is introduction of nutritious items for the frog's digestion, and therein lurks a controversy. What if the proper functions of a consumer device can be described in either a high-level or a low-level way (Neander 1995: 126), and one way says the proper function of the gulping routines is introduction of small dark objects? Soon (section III) I will answer by denying that a history of having introduced the small or the dark into the digestive system helps causally to explain the fact that the gulping routines got replicated.

Can a more specific semantic content be assigned? Perhaps *S* is supposed to signal the presence of a fly (Sterelny 1990: 127) at just that angle (or of a-fly-or-a-midge). It could still be true that *S* is supposed to signal the presence of something nutritious, for it could be that *S* is supposed to signal the nutritious *by way of* signaling the presence of a fly (or a-fly-or-a-midge).

Still, the suggestion appears, at first, implausible. Surely it is immaterial, to the gulping's bringing in something nutritious for the frog, that it bring in something belonging specifically to the fly's lineage or taxon (Agar). Surely it would be just as good if the gulping brought in a gnat or a mosquito.

The device in the frog's visual system which tokens *S* may *usually* track the presence of flies, but so far as its proper function goes—i.e., guiding the gulping to ingest something nutritious—this may seem purely accidental.

But it in fact may not be accidental at all, as Goode and Griffiths (1995: 106-107) point out. Consider Dawkins' (1982: Ch. 4) concept of an "arms race". An arms race begins when selective advantage causes some trait to spread throughout some prey species that renders its members hard to detect by its usual predator species. Natural selection may then cause some refinement of the predators' perceptual capacities to compensate. This will happen, of course, only if evolution does not favor a strategy of simply switching over to other prey that resemble the altered prey's earlier version. Thus other prey species, look-alikes of the former version, must be fairly rare or fairly lacking in nutritive value, and members of the altered prey species must continue to be plentiful enough, or nutritious enough, for it to be worth the evolutionary "cost" of re-tooling.

It is a bit awkward, but not impossible, to envision this scenario in the case of frogs and flies. Just imagine that flies were once bright red, and gained advantage at avoiding predation by frogs by turning dark, thus blending into the crowd of midges, gnats, and mosquitoes—not perfect safety from predation, but an improvement. The neural mechanism in the frog's visual system then got adapted so as to track the darkened flies. That mechanism continued to get replicated because it had enabled, in Sterelny's (1990: 127) phrase, "diachronic tracking" of flies.

In this case it will after all be true that what that mechanism is supposed to track is the presence specifically of *flies*, of organisms belonging to just *that* lineage. And this is perfectly compatible with saying that what *S* is supposed to represent is the presence of food. For it was only *by* coinciding with the presence of flies that *S* managed to coincide with the presence of food—that is (we have supposed) *why* the "arms race" took place. To say it in full-dress form: it was the presence, at the right times and the right angles in front of frogs, of nutritious stuff—which perhaps took, almost invariably, the specific form of *flies*—that causally explained why gulps were rewarded when ancestor frogs gulped as they were steered to by signals generated by the device which tokens *S*. That device, then, would be supposed to

track the presence of *flies*—that would be how it is supposed to discharge its proper function. And because *S* is the product of that device, it is the presence of *flies* that *S* would be supposed to represent. (It doesn't really matter whether one says that *S* itself has been all along what that device tokened, but with a changing neurochemical "shape"—first the shape of a neural message triggered by red things, then the shape of a message triggered by dark things—or says that *S* is a new representation.⁶ The device which tokens it has existed all along, we've supposed—we've said that it altered, got adapted. What that device is supposed to do and how it is supposed to do it determine what *S* is supposed to do and how *S* is supposed to do it.⁷)

III

But mightn't a parallel argument show that what *S* is supposed to represent is the presence of *small, dark, moving objects*? To be sure, what has causally explained the historical successfulness of the steerings which *S* has done has been the presence, at the right times and angles, of *food*. But couldn't one argue that, as a matter of historical fact, there has been no "arms race"—that in fact, the specific version or form of food that was present all along at the right times and angles was *small, dark, moving objects*?

The quick answer is that small, dark, moving objects are not a version or form of *food* at all. They are not even a version or form of *nutritious objects*, or of *objects nutritious for frogs*. They are, rather, *objects alike in their appearance to objects nutritious for frogs*. And the mere appearance of nutritious objects cannot, as such, cause any creature to be nourished, nor causally explain that the creature has been nourished.

The fuller (but related) answer is that there is an important difference between merely *circumscribing* or *pointing to* some explanatory state of affairs, on the one hand, and actually *reporting* that state of affairs, on the other.⁸ There is a difference between *showing where to find* the explanation for some event that happened, and actually *giving* the explanation. To say that small, dark, moving objects were present at such-and-such times, in the trajectories of tongue flicks by historical frogs, is to draw a

line around what explains the successes which historical tokens of *S* had at steering gulping routines. But it is not to *give* the explanation. It does not show what *about* the presence of the dark objects in those trajectories was explanatory. For the successes consisted in nourishment's entering the frogs' digestive systems. But that something dark got gulped is not even the beginning of an explanation of that. It only tells us where to look for the explanation. It tells us to find some feature, in a dark thing which got gulped, which will explain why a thing in the digestive system was nutritious. Saying just that the thing gulped *was dark* does not say what explains this, even if we know that the thing gulped and the thing digested were the same. The color, by itself, has nothing to do with nutritional value. Likewise, the size of an item ingested has nothing to do with the item's nutritional value (though it may have much to do with its ingestibility). Likewise, whether an item is moving.

But what if my opponent were simply to deny any distinction between merely circumscribing and actually reporting an explanatory state of affairs—might this be enough to revive the claim that what causally explained the success of gulpings historically launched by *S* was the presence, at the right times and places, of small dark objects? Since quite obviously it was crucial that the objects in question were not just small and dark, but nutritious as well, my opponent would be relying on a very vulnerable view about causal explanation. Most philosophers think that the contexts in “the fact that ____ causally explains the fact that ...” are not even fully extensional. That some object has such-and-such electrical conductivity, for example, may explain different states of affairs from what the object's having such-and-such thermal conductivity explains. That the President did so-and-so may explain something not explained by the fact that the owner of Socks did so-and-so.⁹ But my opponent would need something stronger than even full extensionality. For she would need to claim that the causal-explanatory role of the fact that nutritious bits flew near ensures, entails, the same causal-explanatory role for the fact that small dark objects flew near. Yet “is a nutritious bit” (or “is a fly”) isn't even coextensive with “is a small dark object”, not even within the frog's normal environment.¹⁰

IV

I began by drawing a distinction between *what* a representation-producing device is supposed to do, and *how* it is supposed to do it—between what its proper function is, and how it performed it in the historical circumstances accounting for its evolutionary successes. I now have completed the argument for a second distinction of *what* from *how*, a distinction between *what* affairs in the outside world a representation-producer is supposed to track (this is part of *how* it is supposed to steer consumer devices), and *how* it is supposed to track them. *What* the frog’s *S*-tokening program is supposed to track, I have argued, is the presence of food items—and perhaps, depending on the evolutionary history, specifically of flies. *How* the program has tracked their presence has been by a very rough-and-ready method. In recent history, at least, it has indiscriminately cued its tokening of *S* to the passing of just any moving dark specks.

Similarly, *what* a certain signal in the visual system of male hoverflies is supposed to represent is the nearness of fertile female hoverflies; *what* the signal in the snake’s nervous system signifies, when triggered by the requisite combination of warmth and movement, is the presence of prey (perhaps specifically of a mouse); and *what* the magnetosome represents (if anything)¹¹ is the direction of safe travel. *How* the devices that produce these representations are supposed to produce them—the detection procedures designed in the them—are, in all cases, decidedly fallible. Hoverflies can be fooled by passing males, and even by airplanes; snakes, by warm puppets of mice; and magnetosomes, by bar magnets.

What I am insisting is that it is crucial not to absorb *how* a naturally selected representation-producer represents into one’s account of *what* that representation-producer represents, on pain of canceling the promise which biological naturalism holds—the promise of explaining how error is possible, ultimately how human judgements can manage to misrepresent. But proponents of the idea that the frog’s *S* means “small, dark, moving objects”—proponents of what Neander calls “Low Church semantics”, and Dretske calls “modest” (as opposed to “inflated”) content (see Neander 1996)—have

anticipated, not the exact arguments of this paper, but the general insistence which these arguments are meant to support. So I must address some things Low Church semanticists have said about the cases just mentioned.

(1) Under how wide a set of circumstances can it plausibly be said that the frog's visual system is malfunctioning? Questions about the frequency of malfunction lie at the heart of Neander's objections against "High Church semantics", or the advocacy of "inflated" content (Neander 1995; Neander 1996). Now, "High Church semantics" is a view about what it is the proper function of a representation-producing device to represent, and so, strictly speaking, Neander's objections are not objections to the view on content advocated here. For this paper has sided with Millikan in denying that it is the proper function of any representation-producer to represent anything. Still, I have said that *what* the *S*-tokener is supposed to track, as part of *how* it is supposed to steer motor routines, is *food* (or *flies*), rather than just dark bits in general. So the position here—call it "CS", or "consumer semantics"—does entail that the *S*-tokener is failing to perform as it is supposed to if it tokens *S* in response to a BB. And Neander finds this implausible. Surely, she objects, the frog's visual system is not malfunctioning when it tokens *S* in response to a passing BB; it need not in any sense be "broken", and may well be functioning exactly as it "should", given the rather casual way it was designed by natural selection.

But CS does not say that the frog's visual system is malfunctioning, or is broken. CS says only that the frog's visual system is failing to perform its proper function. Here is an example of the distinction¹². In Texas, there is a row of Cadillacs, stuck headfirst in the soil, in a kind of imitation of Stonehenge. Their engines are not malfunctioning; there is no need here for a mechanic. But neither are they performing their proper function. They are not doing the job for which they were built.

(2) When the snake senses motion and warmth simultaneously, circuits close which trigger striking routines in its motor nerves; but if one asks what the closing of this circuit is supposed to signal, Rowlands (1997) maintains, a distinction must be drawn. The *organismic* proper function of the circuitry which then sends the signal (=, roughly, what the circuitry is supposed to do for the organism) is to signal "mouse here". But the *algorithmic* proper function of the same circuitry (=, roughly, what immediate

response or output the circuitry is supposed to produce, in response to what immediate cues or inputs) is just to signal “movement-cum-warmth”.¹³ This line of argument couples interestingly with another suggestion of Neander’s. Even in humans, Neander asks, isn’t it the proper function of “early visual processing modules” severally to deliver “information about shapes, colors, texture gradients, lines, and so on—something much more like ‘small, dark and moving’ than like ‘fly’ or ‘frog-food’?” (1996: 270). That is, begin with the thought that since the snake’s circuitry is designed to fire whenever the snake is presented with motion-plus-warmth—even if the source is a warm mouse-puppet—motion-plus-warmth must, on *some* level, be what the circuitry is supposed to signal. Then turn, following Neander’s example, to the “earlier modules”: doesn’t it seem that each of the two signals, which together cause the circuit to close, must have a more “modest” content still—just “motion here now”, and “warmth here now”, respectively?¹⁴

But the warmth-sensing device may not have, on its own, a proper function at all, and the motion-sensing device may not either (nor, by parallel reasoning, the several early visual processing modules). Quite possibly, each of them has as its proper function to team up with the other in switching on and steering, via the doubly-cued neural signal, the snake’s striking routines. So the question of *how* each device is supposed to perform its proper function—what each is supposed to map, or help map—may not even arise except insofar as each operates in tandem with the other. And even when each does operate in tandem with the others, to claim that what the two together manage to represent (or cause to be represented) is *warmth-plus-motion* is to ask the causal question at the wrong place. What matters is not what causes tokening of the doubly-cued neural signal (call it “*M*”). What matters is what caused the lungings which *M* historically launched to be successful, when they were successful. The wrong question asks about what invariably precedes the tokening of *M*. The right question asks about why something which follows on *M* has, on the actual propitious occasions, been well received by the world. How the snake is disposed to token *M* may be far from a perfect guide to what historically made the results of the tokening, and hence the tokening itself, be well received. Evolution is not required to pick infallible designs.

(3) The image projected onto a male hoverfly's retina by a passing female hoverfly will usually launch the male into flight at an angle and speed just right to enable mating. Many other retinal images (of a fly, a male hoverfly, a 747) will also launch a flight. Now, it's because *female hoverflies* have sometimes lain in the trajectories of flights historically launched by the male's visual system that the visual routines launching flights have been replicated. But why stop with "female hoverfly"? It was just as crucial, Neander (1995: 127) points out, that the female hoverflies targeted by the historically decisive flights were fertile; that they would later lay eggs that did not get eaten before yielding new hoverflies; that they were going to have not only "children", but "grandchildren"; etc. In short, the objection concludes, CS would have to assign a content to some signal in the male's neural system which is far too fine-grained and specific to be believable.

But much of the content that seems excessively fine-grained is ruled out by the requirements on a legitimate causal explanation. Consider a particular historical hoverfly, Harry, whose amorous exploits had much to do with the spread throughout the species of the visual devices that print out today's typical flight-launching neural signal, *H*. The crucial question is, what causally explains the fact that flights launched in Harry by *H* were followed by additions of Harry's genes to the hoverfly gene pool? Clearly it was causally important that female hoverflies were often present in the trajectories of those flights, and also that the females in question often were fertile. But we must keep out of the account of what *caused* Harry's advances to have such influence over the gene pool anything that *includes* the fact that they had this influence. That Harry approached females who would actually produce offspring cannot count as a part of what causally explains the fact that Harry's advances led to offspring—unless the females produced offspring by fathers other than Harry, in which case they are irrelevant to the story of Harry's success.

(4) Pietroski (1992) imagines the fictitious species of *kimus*, in which evolution has selected for a visual responsiveness to red objects, but by a very unusual route. Present-day kimus are quite fond of red objects, and their ability visually to pick them out leads the kimus, every day, to move towards just those parts of their native terrain which are never visited by their dangerous predators, the *snorfs*. (Ancestor

kimus insensitive to red all fell prey to snorfs.) What has explained the success of the movements spurred by their visual response to red has been the absence, in the trajectories of these movements, of snorfs. So, Pietrosky holds, what the visual response to red *R* represents, according to CS, is “absence of snorfs over there”. But kimus do not and cannot recognize snorfs! Contemporary kimus have never seen a snorf, and would not recognize a snorf for a snorf, even if one were three feet in front of them. So the content which CS assigns to *R* simply cannot be right.

This objection is carefully crafted, but is nevertheless easy to deflect. Consider the bacteria equipped with the famous magnetosomes, and suppose that (as the story is usually told) the tuggings of the magnetosomes have representational content. Then evolution has selected, in these bacteria, for a device which represents the direction in which safe water lies. But evolution has not selected for representations (or a representation-producer) which show what *about* the water in that direction is safe—which show, that is, what dangers are absent. The bacteria have not a single thought about oxygen, and could not recognize it if it were right in front of them. So it is misleading to suggest that the content of a given tug is “oxygen-free water thither”; it would be better to say, “safe travel thataway”. But then, for parallel reasons, the content of the kimus’ *R* is “safe foraging thataway”.

Moreover, that is *all* the content *R* has. To philosophers who believe in the reality of colors, it may seem obvious that *R* has more—if not altogether different—content: it signals the presence, in this direction or that, of the color *red*. Two different ideas may motivate this reaction. One has to do with “what it is like” for the kimus to undergo *R*. The other has to do with what discriminations *R* enables the kimus to make.

The first idea holds that the kimus’ response *R* is phenomenally just like our “sensation of red”. But as Pietrosky actually frames the example, the kimus are not said to be visually responsive to any chromatic color other than red—not even to green. Yet as the opponent process theory of color vision shows, our capacity to have “the sensation of red” is physiologically inseparable from our capacity to have “the sensation of green” (Hardin 1988). Indeed the kimus may, for all Pietrosky says, have a color solid that is barely a solid at all—something more like a point or a line. So it is hard to have confidence

in any claims about “what it is like” for them to have *R* (see Clark forthcoming: Ch. 1). Pietrosky could, of course, switch to saying that kimus have an entire system of color vision, similar to that of normal human trichromats. But the switch would threaten the plausibility of the claim that the proper function of *R*, and of the mechanism that generates it—*all* that the kimus’ system of color vision has ever done for them, on the historical occasions that have led to its proliferation—has been activation of their snorf-evading locomotory routines.¹⁵

The second idea holds that *R* enables kimus to do a very good job of discriminating red objects from objects of other colors, but only a relatively poor job of discriminating safe foraging zones from dangerous ones: *R* could sometimes indicate safety in the direction of danger, and could more often fail to indicate safe foraging where safe foraging is in fact to be found. But consider the frog. Its signal *S* leads it almost unfailingly to single out (and swallow) nearby objects which are *small, dark, and moving*. But for all that, it is an open question—so the size of the literature suggests!—whether the content of *S* is *small, dark, moving thing*. What can make the question seem closed is the habit of thought, just now noted in connection with Rowlands, of supposing that the stimuli which lawfully (or close enough) trigger a given neural signal must be that in the world which the signal represents. I have argued that content is determined by the routines that are the “consumers” of that signal, and by the factors in the world which have made those consumer routines successful *often enough*—i.e., far less than lawfully—for natural selection to have reproduced them. If the consumer device for *R* has been locomotory routines in kimus, and what the locomotory routines were selected for was safe direction (*often enough*) of the kimus’ foraging, then “safe foraging” is the content of *R*.

V

It really is defensible, then, to distinguish what a naturally selected representation-producing device is supposed to represent, from how it is supposed to represent that something. The benefits of drawing the distinction are considerable. It then is possible to show how naturally selected

representations can misrepresent reality, and be empty, and not just in extraordinary cases but more or less routinely. And that is crucial. That is the only way one can keep bright the promise of biological naturalism to explain how human judgements can be empty or inaccurate concerning real subjects. For biological naturalism *begins* with naturally selected representations, and cannot help treating them as (in some measure) *paradigmatic*. And one can achieve this benefit painlessly, simply by cashing in on what is biological about biological naturalism.

¹By “naturally selected representations”, I mean, roughly, representations produced by mechanisms or devices which evolution has “hard-wired” to produce representations of just *that one type*—representations of a single “shape”, confined to a single content. Thus, for example, the mechanism in beavers which induces them to slap their tails when danger approaches produces only tail slaps, and is only for signaling the presence of danger (and not, for example, easily felled saplings). But the phrase needs help from two (at least) caveats. First, it may be that, as Millikan holds, any biologically generated representation has at least some articulation, some variability of content: each beaver tail slap, for example, means “danger *here now*”, and both the “here”s and the “now”s meant will vary (Millikan 1984: Ch. 2). And as long as we must inevitably allow some variability in content, it would be useful, secondly, to construe “naturally selected representations” as covering representations such as bee dances—representations which unmistakably have some articulation (as to the direction in which the nectar lies), but which, because of “hard-wiring”, can’t have much (bees can’t signal the direction of hungry bears).

²Rowlands—to anticipate section IV—holds that when the frog’s visual system tokens the gulp-launching signal, the content of *its* state is “small black dot”, though the content of the state of *the frog as a whole* is then “eatability!”. Other important discussions of this sort of view are found in Dretske (1990) and Fodor (1990: 70-79).

³Having conceded that her view offers little room for misrepresentation in the arena of “innate processing”, Neander points out that much room for misrepresentation may still remain in “the processing

that depends on learning”; “[n]ot all misrepresentation need be of the same type” (1996: 269). But such a position argues against, not for, a strategy of attending closely to naturally selected representations—provided, that it, the ultimate aim is to come up with a naturalist understanding of inaccuracy and emptiness in human thought.

⁴There is slight variation in the way these authors say that the question of content reduces to the question of the representation’s proper function. Some say the root question is: what is the proper function of the device that *produces* this representation to track or signal? Others are willing to speak of the *representation* as having the proper function of signaling this or that. But the difference is more a matter of expression than of doctrine. The representation could not have a proper function if the device which produces it did not, and if the device which produces it has a proper function, the representation will eo ipso have what Millikan calls an “adapted and derived proper function” (Millikan 1984: Ch. 2). (The latter is a point which Rowlands fails to grasp; see 282.)

⁵Thus Rowlands, 287-88. Strictly speaking, *S* is what Millikan (1996) calls a “pushmi-pullyu representation”: it is neither an indicative icon (“Fly in front now”) nor an imperative icon (“Gulp thataway now!”), but an undifferentiated precursor of both. Gibsonian affordances, as Millikan points out, are best understood as that which is represented by a special kind of pushmi-pullyu representation, namely one that can be found in either a potentiated or a depotentiated state. Strictly speaking, *S* is not an instance. For it has no depotentiated state: when *S* is tokened, the frog gulps, even if his belly is full. See Millikan forthcoming.

⁶This question should have been canvassed at 106 in Goode and Griffiths’ paper; it would have kept them, and Nicholas Agar, from talking past each other.

⁷*S* has—to repeat the point made in footnote 4—what Millikan calls a *derived* proper function.

⁸I take the distinction from Hillel Ruben (1990: 163-64). I add some elaborations in my 1996.

⁹In my 1996, I offer this example: suppose that a certain distinctive shade of blue is in fact found nowhere else in nature than in certain crystals containing bearing a rare combination of trace metals. That

these crystals also have distinctive capacities as superconductors might well be explained, causally, by their containing just those trace metals—but not by their being of just that shade of blue. For an excellent discussion of the larger issues, see Ch. 5 of Ruben.

¹⁰If “the fact that___causally explains the fact that...” is not fully extensional, can it really pick out an objective relation, which obtains independently of thought and talk? The question is a sharp one for Millikan, as Putnam has pointed out, since she needs this relation to *ground* the aboutness of thought and talk (Putnam 1992: 28-31, and fn. 9-10 on 206). But it can be answered in the affirmative, as I argue in my 1996.

¹¹The magnetosome does not, as Mark Rowlands points out, communicate its pointings to any distinct device within the host bacterium—there is no “consumer device”, in Millikan’s terminology, for it to steer; and so, Rowlands concludes, it is misleading to think of the magnetosome as representing anything (299-300).

¹²I have given this example before, in 1995: 520.

¹³Rowlands does not actually supply definitions of “algorithmic” and “organismic”; I trust that my rough attempts are accurate.

¹⁴In fact, there is room to question whether the neurons populating the earlier modules are invariably attuned to single types of stimuli; “our understanding of face-to-face communication,” for example, “is aided by the response of *auditory* neurons to *visual* stimuli” (MacLennan, forthcoming).

¹⁵For suggestions about what the proper function of our actual system of color vision is, see my forthcoming, and cf. Clark forthcoming: Ch. 6.

References

[Author’s own writings omitted so as to preserve anonymity.]

Akins, Kathleen. 1996. Of sensory systems and the “aboutness” of mental states. *Journal of Philosophy* 93: 337-72.

Agar, Nicholas. 1993. What do frogs really believe? *Australasian Journal of Philosophy* 71: 1-12.

-
- Clark, Austen. forthcoming. *A Theory of Sentience*.
- Dawkins, Richard. 1982. *The Extended Phenotype*. Oxford: W. H. Freeman and Co.
- Dretske, Fred. 1988. Misrepresentation. In his *Explaining Behavior*. Cambridge, Mass.: M. I. T. Press.
- . 1990. Reply to reviewers [of *Explaining Behavior*]. *Philosophy and Phenomenological Research* 50: 824-26.
- Fodor, Jerry. 1990. *A Theory of Content and Other Essays*. Cambridge, Mass.: M. I. T. Press.
- Hardin, C. L. 1988. *Color for Philosophers*. Indianapolis: Hackett.
- MacLennan, Bruce. forthcoming. Finding order in our world: the primacy of the concrete in neural representations and the role of invariance in substance reidentification.. *Behavioral and Brain Sciences*
- Millikan, Ruth. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, Mass.: M. I. T. Press.
- . 1986. Thoughts without laws. *Philosophical Review* 95: 47-80. Reprinted in Millikan 1993.
- . 1989a. In defense of proper functions. *Philosophy of Science* 56: 288-302. Reprinted in Millikan 1993.
- . 1989b. Biosemantics. *Journal of Philosophy* 86: 281-97. Reprinted in Millikan 1993.
- . 1993. *White Queen Psychology and Other Essays for Alice*. Cambridge, Mass.: M. I. T. Press.
- . 1996. Pushmi-pullyu representations. In Tomberlin 1996.
- . forthcoming. Some different ways to think.
- Neander, Karen. 1991. The teleological notion of function. *Australasian Journal of Philosophy* 69: 454-68.
- . 1995. Misrepresenting and malfunctioning. *Philosophical Studies* 79: 109-41.
- . 1996. Dretske's innate modesty. *Australasian Journal of Philosophy* 74: 258-74.
- Pietrosky, Paul. 1992. Intentionality and teleological error. *Pacific Philosophical Quarterly* 73: 267-82.
- Putnam, Hilary. 1992. *Renewing Philosophy*. Cambridge, Mass.: Harvard Univ. Press.
- Rowlands, Mark. 1997. Teleological semantics. *Mind* 106: 279-303.
- Sterelny, Kim. 1990. *The Representational Theory of Mind*. Oxford: Blackwells.
- Tomberlin, James. 1996. *Philosophical Perspectives, IX*. Atascadero: Ridgeview Publishing.
