

Abstract

Reading Mother Nature's Mind

I wonder how Dennett and I, both firm naturalists and believers in the relevance and importance of natural selection for understanding the human mind, should have come to such different conclusions about holism in the theory of meaning.

I try to focus our differences by examining the relation between what Dennett has termed "the intentional stance" and "the design stance." Dennett takes the intentional stance to be more basic than the design stance. Ultimately it is through the eyes of the intentional stance that both human and natural design are interpreted, hence there is always a degree of interpretive freedom in reading the mind, the purposes, both of Nature and of her children. The reason, or at least a reason, is that intentional interpretation is holistic, hence indeterminate, for the kinds of reasons given by Davidson and Quine. In contrast, I take the design stance to be more basic than the intentional stance. Intentional attributions express our best guesses about the locations of effects of certain kinds of natural design. And although there is often indeterminacy, ambiguity, or vagueness concerning what it is that natural selection (or learning) has selected for, these indeterminacies and vaguenesses are local, not holistic. There is reason to suppose that the better portion of Nature's purposes and the intentional states of her children are determinate in content within quite closely defined limits. I propose to defend this position as well as I can, so as to call from Dennett his own views on precisely where our paths separate (if they really do).

Reading Mother Nature's Mind

Ruth Garrett Millikan

University of Connecticut

How does it happen that Dennett and I, both firm naturalists and believers in the relevance and importance of natural selection for understanding the human mind, should have come to such different conclusions about holism in the theory of meaning? I will explore the possibility that the difference stems from different understandings of the relation between the intentional stance and the design stance. Dennett takes the intentional stance to be more basic than the design stance. Ultimately it is through the eyes of the intentional stance that both human and natural design are interpreted. But the correctness of intentional-stance interpretation is not a completely determinate matter. There is always a degree of interpretive freedom in reading the mind, the purposes, both of Nature and of her children. The reason, or at least a reason, is that intentional interpretation is holistic, hence subject to Quinean/Davidsonian indeterminacy.¹ On the other hand, I take the design stance to be more basic than the intentional stance. Intentional attributions express our best guesses about the location of (read the next phrase transparently) effects of certain kinds of natural design. And although there is often indeterminacy, ambiguity, or vagueness concerning what it is that natural selection (or learning) has selected for, these indeterminacies and vaguenesses are local, not holistic. There is reason to suppose that the better portion of Nature's purposes and of the intentional states of her children are determinate in content within quite closely defined limits.

I propose to defend my position as well as I can, so as to call from Dennett more precisely where our paths separate, if indeed they do.

Let me begin by remarking on a possible equivocation in the notion of the design stance. On the one hand, prediction from the design stance seems to be just predicting that a thing will indeed do what it was designed to do. Suppose I believe that this object over my head is a smoke alarm, that is, I believe it has been designed to sound an alarm if it encounters smoke. Without having any idea what is inside it, or how it is supposed to accomplish this task, from the design stance I may confidently predict that it will sound an alarm if it encounters smoke. On the other

1

Dennett often refers to Davidson and, especially, to Quine's thesis on the indeterminacy of radical translation (Quine 1960 ch 2), most relevantly for this discussion, perhaps, in (1987, pp. 37-42; 1990, p.180). Quine's thesis is that translation is always radically indeterminate. And should Quine's reasons for saying this hold up, no weaker conclusion would suffice. Dennett, on the other hand, seems to vacillate on the ubiquity of indeterminacy. For example, in (1990) he moves from saying in his text on p. 180 that Quine claims that "there may be no deeper facts that settle the matter [of correct interpretation of people's intentional attitudes]" to saying in his footnote on this very sentence: "That there are no such deeper facts is also argued for at length in [Dennett 1987]" (italics mine). Part of what I hope for from this essay is clarification from Dennett which of these is really his own position and why. -- Not to preclude the possibility (recognized by my own position as well) that this very question may happen to have, or happen so far to have had, no determinate answer.

hand, Dennett often speaks of the design stance as though to use it one would need beliefs also about how a thing is designed accomplish its task(s). One would have to know, for example, not just that the tournament chess-playing computer is designed to win at chess, but something about the program it is designed to implement in order to win at chess. Shifting the example here makes the equivocation more difficult to spot, for it is evident in the case of the chess-playing computer, as it is not in the case of the smoke detector, that predicting that the device will actually accomplish its goal of winning from knowing only that it was designed to win would be a risky business. In this latter case, a reasonable prediction could proceed only by knowing something of how the machine was supposed to go about winning, for example, at minimum, that it was designed to win with legal moves. This makes it easy to assimilate what one must know in order to predict from the design stance to what one must know to give a Cummins-style functional analysis of how a system works (Cummins 1972).²

This threatened equivocation invites a sister one that lurks in the background of much current thinking about Cummins-style functional analysis. Although Cummins introduced his notion of functional analysis as explicitly not teleological, as not employing the word "function" in a way that connects with either purpose or etiology, he explained his idea by reference to what circuit diagrams, flow charts, and computer programs tell us about systems. But what these sorts of items, as found in the real world, generally tell us about actual systems is not how they do operate but how they were designed or intended to operate. The only reason for including a circuit diagram in the literature that comes with your clothes dryer is that there may come a time when your dryer fails to accord in its workings with this diagram, and then knowing how it was designed to work may help in repairing it. If you were to move from a look at the circuit diagram to a prediction about the effect of certain settings on the dryer's dials, that would be design stance prediction—prediction from a teleological stance—not a direct outcome of Cummins-style analysis. The circuit diagram is not a description of the actual dispositions of your particular dryer (a Cummins analysis), though hopefully it does accord with these dispositions. It is a description of the dispositions the dryer was intended to have.

What I will mean by "the design stance" in what follows is the teleological stance, not a Cummins-style stance. The design stance is the predictive stance that moves from what a thing was designed to do to a prediction that it will do that thing. On this account, design stance prediction is possible starting from beliefs about any aspect or aspects of a thing's design. It is possible starting with beliefs only about the most general "specs" for a thing's design, for example, the thing is designed, perhaps, to show the correct time, or to effect that the missile tracks its target. Compare here David Marr's "first level of analysis" or "task analysis" for his theory of vision (Marr 1982). Alternatively, it is possible starting with completely detailed knowledge of how a thing is designed to work, for example, of how the clock wheels

²

A Cummins-style analysis explains a complex capacity that a system has by showing how simpler capacities or dispositions possessed by it and/or by various of its parts add up to that capacity.

are supposed to engage, or of the program the computer in the missile-tracker is intended to run, and so forth. Compare here Marr's higher levels of analysis.

Compare the inferences that are involved when one makes a prediction from the intentional stance. Dennett does not define intentional systems with reference to origin. Rather, an intentional system is one that currently displays a certain pattern of behavior or, being more careful, is one that has a current disposition to display certain kinds of patterns of behavior. There is no need here to look too closely at what defines these patterns of behavior. Deferring to Dennett's use of "rational," I will call them "real rationality patterns," and note only something of what would seem to distinguish them from patterns of behavior that are "rational" merely in the sense of being reliably need-fulfilling, reliably gene-propagating or the like.

It is very rational for tortoises to grow shells and then to pull their heads inside when in danger. This is such a smart thing to do that tortoises have outlasted nearly every other largish kind of animal on earth. But tortoises themselves are not very smart. Nearly as old as the tortoise is the alligator/crocodile family, with members up to twelve feet long but with brains no bigger than a peanut. They too are built in smart ways but are not themselves very smart. Turtles and alligators are built such that they need not notice in much detail what situations they are in. They can handle most situations effectively with the same small stock of simple behavioral tricks. But there are other animals that care very much exactly what situations they are in, and are built to be able to change either themselves or the situations they are in accordingly. Only certain of the latter display "real rationality patterns."³ Real rationality patterns are, roughly, dispositions to respond to a wide variety of environmental situations with a wide variety of wide-context-sensitive responses that promote the animal's interests. Especially, the ability to have the very same proximate external situation help to produce quite different but reliably helpful responses depending on the wider or more distal context the animal is in is a move toward displaying more ideal rationality patterns. Inevitably, it is a move toward giving "individual belief-like states more to do, in effect, by providing more and more different occasions for them to serve as premises for further reasoning." (Dennett 1987 p. 30).

What kind of inference is involved then when we make a prediction from the intentional stance? Dennett is very clear about some kinds of inferences that are not involved. To ascribe intentionality to a system is not to make any bets about how the Cummins-style analysis of the system will go. Certainly no intentional system is made of jelly inside, but one might contain the analogue of a huge how-should-I-respond-to-this-stimulation look-up chart, another a hugely intricate lattice that determines potentiations to potentiate potentiations to action, while a third processes mental sentences. Likely, many contain collages encompassing many different kinds of principles all operating at once.

Predictions from the intentional stance make no reference to actual internal mechanisms. Nor, it seems, do they make any reference to design in the historical

3

In his (1996) Peter Godfrey-Smith has an extended discussion of why intelligence is not, just in general, a good thing, that is, of why it is not always smart to be smart.

sense. Dennett is clear that the intentional stance is separate from the design stance.⁴ Attributions of intentionality go no deeper than claims about current patterns of external behavior. Thus prediction from the intentional stance seems to go from the attribution of rationality—attribution of the disposition to display real rationality patterns—directly to predictions about instances of such patterns. It has the form "All of O's behaviors fit the real rationality pattern so O's next action will fit the real rationality pattern" or, allowing for idealization, "Most (many) of O's behaviors fit the rationality pattern so probably (perhaps) O's next behavior will fit the rationality pattern."

If this is the form of inference involved, it has two noteworthy peculiarities. First, it does not appear to be a form of inference that supports explanation. It seems to have the same general form as "All of the boys in this room were born on weekdays, so Johnny, being one of the boys in this room, was born on a weekday," which derives that, but does nothing to explain why, Johnny was born on a weekday. At most it explains why one should believe that he was, if one has a certain prior belief. Similarly, on this view, that O's behaviors fit the rationality pattern does not seem to explain why O behaves as O does, but merely why one should expect O to behave that way, given that one already believes O's actions fit the rationality pattern.

The second peculiarity of this pattern of inference is an unclarity about how one rationally acquires belief in the premises. "O's behaviors fit the rationality pattern" is short for "O has a disposition to produce (only, mainly, many) behaviors that fit the rationality pattern." But dispositions themselves are not, of course, directly observed. What is observed can only be a certain number of actual behaviors that are consonant with the possibility that they express a general disposition to fit the rationality pattern. We would like to treat the inference from some behaviors fitting the rationality pattern to all behaviors fitting the rationality pattern as a simple induction, of course. But fitting part of the rationality pattern is a very complex and rather disjunctive thing to do, not a simple thing like what emeralds do in support of the induction "all emeralds observed so far have been grue so the next emerald will be grue"! Compare: "The ink marks on the very small part of this paper that I now can see look quite a lot like the queer configuration of streets right around Peter's house in Stockholm, so the marks on the rest of the paper will look like the rest of Stockholm." What makes us think that fitting some very small (a nearly vanishingly small) part of the rationality pattern is projectible to the whole?

4

At (1987, p. 73) Dennett says, "One can view the intentional stance as a limiting case of the design stance: one predicts by taking on just one assumption about the design of the system in question: whatever the design is, it is optimal. This assumption can be seen at work whenever, in the midst of the design stance proper, a designer or design investigator inserts a frank homunculus (an intentional system as subsystem) in order to bridge a gap of ignorance." Being designed to do X in an optimal way is not, however, being designed to be rational. Recall the tortoises and the alligators.

Hint: what kind of history would one have to hypothesize for the part of the paper one sees in order rationally to make an inference to a map of the whole of Stockholm?

Why should partial exhibition of a rationality pattern, that is, apparent exhibition of a disposition to rationality, be a projectible predicate? Why, especially if there are so "many internally different ways of skinning the behavioral cat" (Dennett 1994, p. 520)? If there are so many different ways, there must surely be many more ways to produce temporary false appearances of rationality patterns but that would have failed to follow through in slightly different circumstances and will surely fail to follow through in the future. Is there nothing to go on, then, but the so-far-mysteriously-true meta-inductive premise that in the past, inferences from apparently-rational-at-t to apparently-rational-at-t+1 have often held up?

I think Dennett thinks this question is answered the same way I think it is. There is nothing that exhibits apparently rational patterns for any time or in any detail that was not designed to do so, either by natural selection, or by something that natural selection designed. Not only are there no swampmen, there are no apparent swampmen. There are no accidents that apparently exhibit coherent, rational behavior for a time but not due to any underlying general dispositions to do so. That they should express certain kinds of⁵ real rationality patterns is one of the "specs" for certain of nature's designs. That is, showing rationality patterns of one kind or another is often an excellent way to get yourself selected for, granted you don't have a strong enough shell or large and strong enough jaws to get by without.

True, Dennett takes spandrels (Gould and Lewontin 1979) and exaptations (Gould and Vrba 1982) seriously in this context. Sometimes what a thing was designed for is not a good guide to what it is used for, nor what it is used for a good guide to what it was designed for (see, especially, Dennett 1990). Elsewhere I have argued against the importance of the notions of spandrels and exaptations in the context of determining biological functions (Millikan 1993 ch. 2; 1999. See also Godfrey-Smith 1994 and Dennett 1995). The case against true rationality dispositions being mere exaptations is much stronger than the general case, however. Real dispositions to exhibit rationality patterns are very sophisticated, subtle, finely tuned dispositions indeed, ridiculously improbable dispositions if not assumed to be shaped by natural selection. A spandrel or exaptation that resulted in a frog's being accidentally disposed to flick out its tongue in exactly the right direction in response to exactly the right sort of angle of motion of a fly image on its retina is (barely) conceivable. Then we might, I suppose, take the intentional stance towards the frog, saying that it (accidentally) knew when and at what angle a fly is passing

5

I do not want to endorse the position that there is such a thing as the one perfect ideal of "rationality" in Dennett's sense of that term, to which different organisms approximate more or less closely. That the little thing plays a good game of chess does not give us reason to suppose it will be smart and keep itself away from harmful magnetic fields. Rather, there are many different kinds of smartness, many different ways to be smart and many different kinds of things to be smart about. Some of these ways are surely more versatile than others, but there is no single dimension or apex involved here. Perhaps Dennett agrees?

and (accidentally) desired to eat flies. But this is exactly the kind of simple inflexible case, I believe, where Dennett would agree there is no use in talk of rationality. We are less likely to be misled if we talk only of reflexes.

Indeed, if some actual organism had acquired fully rational dispositions to behavior totally by accident, and if the observed appearances of rationality resulting from these dispositions were known by us to have resulted from accident, we would be foolish indeed to project these appearances into the future. Not knowing that the absurdly improbable—full and real rationality patterns without design—had indeed occurred, we would take it as far more likely that these were freakish false appearances of rationality resulting from accidental interactions of unsystematic features irrelevant to real rationality. Ridiculously unlikely that these features should continue to produce apparently rational outcomes in future. Clearly, exaptations for rationality would not be projectible.

If this is so, then predictions made from the intentional stance are really grounded implicitly in something like inference to the best explanation and out again. From enough apparently rational behavior one can infer design for rationality, just as one can infer design for seeing from good sight. And from design for rationality, one can infer real dispositions to rationality patterns, as opposed to mere temporary illusions of such dispositions.

It thus appears that the intentional stance must be underwritten by the design stance, rather than vice versa. Then too, the fact that the organism is rational, indicating that selection pressures have slowly designed it to be rational, serves as a genuine explanation of its behavioral patterns, not merely as a redescription of them. (Dretske would say that the explanation was by way of the "structuring cause" of the behavior—Dretske 1988.)

... imagine posing scientists the following Swampman-style questions. Suppose that you discovered a thing that attracted iron but [the molecules inside were] ... not m-aligned (like standard magnets). Would you call it a magnet? Or: suppose you discovered a thing that was not M-aligned but did attract iron. Would you call it a magnet? The physicists would reply that if they were confronted with either of these imaginary objects, they would have much more important things to worry about than what to call them. Their whole scientific picture depends on [this]...and the "fact" that it is logically possible to break this deep regularity is of vanishing interest to them....

... If I ever encounter a plausible believer-candidate that violates [the idea that "a brain filled with jelly or sawdust could not sustain beliefs"], what to call it will be the least of my worries, since my whole theory of mind will be sunk. (Dennett 1994, p 519)

Similarly, I should think, for rationality patterns not designed to be such by natural selection. Intentional systems are as essentially designed as magnets are aligned. An intentional system IS a designed system. Rationality IS something that has been selected for as such. Rationality is not, then, merely a disposition to display rationality patterns. It is not something that resides merely in the present dispositions of a thing.

If we wish to know how determinate in content our various mental states are, then, the question we should ask is, first, how determinate is it what natural selection

has selected for? Second, given that what actual animals and their behaviors are like often strays quite far from what natural selection selects for, we should ask, how determinate in content are our various mental states when they stray from Nature's ideal?

It is crucial not to confuse the question whether what natural selection has selected for is determinate with the question how well what has been selected for is determined by the evidence we can collect or have collected. This would confuse epistemological determinacy with ontological determinacy. Indeterminate evidence for history is not indeterminate history. What kind of indeterminacy might there be then in selection histories themselves?

A number of kinds of indeterminacies, I will soon argue. But there is one kind of indeterminacy that it seems to me will not be found. What will not be found, I believe, is any parallel to holistic Quinean/Davidsonian indeterminacy.

The indeterminacies that Quine argued for resulted from a very particular theory of linguistic meaning arising, originally, out of 20th century empiricists' struggles (Carnap, Norman Campbell, Reichenbach, Braithwaite, Hempel) to understand the language of theoretical science. The results were then applied to everyday language by philosophers such as Sellars, Feyerabend and Quine and soon became dogma in empiricist circles: concepts are nodes in an inference net or a sentence-association net receiving input from sense, predicting the ongoing course of sensory stimulations and guiding action decisions. The semantic content of any concept thus depends on the contents of the concepts inferentially surrounding it, even when the concept plays out part of its role in observation judgments or in desires that directly produce basic actions. Thus the meaning of any concept depends on the meaning of many other concepts, and indeterminacies result from a variety of possible kinds of holistic re-mappings. Of course Dennett does not accept this theory of thought, certainly not with its original realist interpretation as a theory of thought mechanics.⁶ But it is worth noticing why a theory of this sort cannot possibly be applied realistically when we turn to the question of the determinacy of Nature's intentions.

One problem, of course, is the likelihood that this theory of mental semantics is in the end incoherent.⁷ But supposing it coherent for the case of humans, there still would be no way to apply it realistically to the "purposes" of Nature. First, classical theories of the indeterminacy of meaning are theories about the nature of

6

Dennett's theory of intentionality officially leaves it open, of course, that some intentional systems might actually work this way -- better, perhaps, than sawdust inside.

7

This theory of meaning is decidedly parochial looked at from a historical perspective, and there are many signs of its slow demise. Recently Fodor and Lepore (1992) have summarized many strong arguments against holism in the theory of meaning. In (Millikan 1984, 1993, 1998a, 1998b, 1998c) I have offered many details for construction of a decidedly non-holist theory of meaning for both language and thought.

representations, including of course, those that represent purposes and, on representational theories of mind, those that embody people's explicit beliefs and desires. But no one supposes that Nature anywhere represents her purposes (except, of course, in so far as her creatures are part of her and some of them may represent purposes). There is no sense in talk of a holistic re-mapping of Nature's beliefs and desires parallel to the classical theory of indeterminacy of translation, because if Nature has anything analogous to beliefs and desires, they are not represented. But second, and far more important, there is no sense in such talk because there is nothing in nature analogous to beliefs and nothing that so much as reminds one of inference. Nature turns out products that easily remind us of products purposefully designed by humans. But there is nothing in the process by which she does so to remind us beliefs, reasoning, inference nets, and so forth, hence of holistic re-mappings.

On Dennett's real-patternist theory of intentionality, the first disanalogy—Nature's failure to represent anything—is not relevant. But the second is decisive. The only way to read inference into Nature would be mechanically to interpret every product of natural design, such as an eye, with this sort of ritualistic formula: "Nature wanted her creature to see and she believed that employing such and such principles of optics and using such and such natural materials would enable it to see so (you see) that's just what she did!" Suppose we ignore Dennett's warnings against the barrenness of the intentional stance in cases where the supposed "individual belief-like states" are not given "more to do." Still, that formula won't yield any "beliefs" the contents of which are unclear because of their dependence on the contents of surrounding "beliefs" and "desires." Similarly for Nature's "desires." Quinean holistic indeterminacy depends on the assumption of a very particular and peculiar theory of how human meanings are determined, a theory that very evidently does not apply to Nature's purposes, determined, as they are, by the facts of natural selection.

It is important, I have said, not to confuse indeterminacy in the evidence for selection pressures with indeterminacy in the pressures themselves. This, of course, is compatible with there being, in fact, indeterminacy in the selection pressures as well. But conflict or lack of clear direction in selection pressures is a completely different sort of thing from Quinean/Davidsonian holistic indeterminacy. Unclear and conflict in selection pressures is local, at least in the sense required here. It has no tendency to contagion through any analogue of inference dispositions.

In Elliott Sober's terms, there is not just selection of features but selection for features. [Sober 1984a] And without this "discriminating" prowess of natural selection, we would not be able to sustain functional interpretations at all.

Certainly we can describe all processes of natural selection without appeal to such intentional language, but at enormous cost of cumbersomeness, lack of generality, and unwanted detail. We would miss the pattern that was there, the pattern that permits prediction and supports counterfactuals. The "why" questions we can ask about the engineering of our robot, which have answers that allude to the conscious, deliberate, explicit reasonings of the engineers (in most

cases) have their parallels when the topic is organisms and their "engineering". (Dennett 1990, p. 189)

Not quite, I believe. If we dropped all talk of function or purpose in the biological world we would indeed be unable to discern most of the important patterns that are there. But talk of beliefs and "conscious, deliberate, explicit reasonings" would be otiose in biology. There is no need to drag the whole intentional stance into biology in order to perceive nature's handiwork and the principles of natural design.

But these arguments do not settle the question we started with. That question concerned not whether holism was at work undermining determinacy in natural design generally, but whether it undermines determinacy in human beliefs and intentions. So far, I have argued only that if there is indeterminacy in human intention it is not rooted in the necessity of taking the intentional stance when interpreting nature as a whole. Nature does not have to be interpreted holistically, but perhaps people do.

Once again, it is important not to confuse epistemology with ontology. If it were true that the only way to guess the contents of another person's beliefs, desires and intentions was by forming a hypothesis about the whole or large parts of the whole of that person's intentional attitudes, that would not make the contents themselves indeterminate. Nor, of course, should we inadvertently slip into the familiar view that our thoughts are indeterminate for Quinean reasons. That strange picture of thought might conceivably be true but we certainly can't assume so, nor does Dennett usually do so. Rather, he has told us, there are many ways to skin the behavioral cat. Finally, that there are many ways to skin the behavioral cat does not mean, of course, that any individual cat does not have a determinate way to be skinned. The mechanism inside may be complex in the extreme, working in accordance with dozens of different principles each accounting for a different aspect or moment of the cat's rationality, but that just in itself would not cause indeterminacy any more than for bodily functions.

The theory of content I espouse for the whole person I espouse all the way in. The neurobiological theory of content is homuncular functionalism, to dress it in its most vivid metaphorical costume, and hence the very same principles of interpretation are used to endow subpersonal parts with contents as are used to endow whole persons.The way in which personal-level attributions of belief and other intentional properties get confirmed (in the crunch) by subpersonal attributions of (nonordinary) intentional properties is roughly parallel to the way in which one might confirm one's attributions of culpable motives to, say, the British Empire, or the CIA, or IBM, by discovering a pattern of beliefs, desires, intentions, among the agents whose joint activity compose the actions, beliefs, and intentions of the superpersonal agent... (Dennett 1994, p. 528)

If this sort of technique were necessary in epistemological practice, however, that would not argue for a correlate no-fact-of-the-matter ontology. The rationality of the CIA, should that body happen to be rational, was surely not designed by natural selection, but the rationality of a person was so designed. Why not expect there to be some definite principle of design, then, that would determine what a person's mental representations were each intended by nature to mean?

Can we suppose, perhaps, that natural selection is as blind as we are when trying to see the insides of an organism's head? It can see only whether the emerging behavior is rational, not how it was caused to be rational. So there is no way of selecting for determinate principles on which to rest rationality, or for determinate vehicles to implement it? It is true, of course, that in a sense, all natural selection ever sees about any organism is whether it gets to the next generation or not—never how. The reliable result for every plant and animal, however, is good design, usually of a vast number of intricately interlocking inner parts, each working in accordance with entirely determinate principles. Occasionally nature's designs are cumbersome or inelegant, and often they work only in rather specialized circumstances, but there are always good reasons why these designs work when they do. Moreover, evolutionary history characteristically displays progressive perfecting over the years of certain parts for certain roles. Clearly Nature has specific principles of operation for these mechanisms "in mind." Just as each animal is designed to make its living in a definite way, or in a definite set of alternative ways, each rational animal must be designed to be rational in some definite way, or perhaps in some large but definite set of alternative ways.

About Kripke's Pierre (Kripke 1979) Dennett says,
Which propositions, please, should be inscribed on Pierre's belief list?... Pierre is an imperfect believer, as we all are...Psychological-attitude talk is a huge idealized oversimplification of the messy realities of psychology. Whenever push comes to shove in borderline cases, its demands become unanswerable. This is my pretty pernicious instrumentalism showing...propositional attitude claims are so idealized that it is often impossible to say which approximation, if any, to use....Biologists shrug when asked whether herring gulls and lesser black-backed gulls are different species...How close to the (ideal) "specs" does something have to be to count as a genuine FM tuner?...a gradation of cases from truly embedded or encapsulated subdoxastic states to more and more versatile cognitive states... (1994, p.525-6)

It is tempting to fuss about details here. Why does the transition from truly embedded to more versatile cognitive states produce determinacy of content? Bee dances and rabbit thumps, each of which has only one thing to do, are not in the same pickle as Pierre's belief about London. And if the indeterminacies are rooted in holistic mappings, why should push come to shove in borderline cases rather than over the whole map? But forcing questions of this sort feels rather like squeezing soap. Instead, let me offer an alternative story about indeterminacy of content, and see with how much of it Dennett may agree.

Yes, there can be indeterminacy of biological function, even in a system's basic design. I don't have in mind here a biological trait that was originally designed for one purpose and now is used for another. Such a trait will generally be under current selection pressures to maintain its present form precisely for serving the new function, just as it originally invaded the gene pool under pressure to maintain this same form in service of the old function. Strictly speaking, of course, natural selection only selects, never designs, and having been selected once before for another function does not cancel its being in process of selection now for a new

one—selection, that is, over less perfect variations and other accidental junk thrown up by mutation.⁸ If there is indeterminacy at the transition from selection for one function to selection for the next, this is temporary and uninteresting indeterminacy, not really worthy of attention. More typically, the transition period probably finds it being selected for the service of both functions at once.

Interesting indeterminacy arises, however, when there is vacillation over alternating stretches of time or space in the selection pressures on a trait, say, from one short ice age through a warming period and into another, or from one kind of terrain that a species inhabits to an alternative terrain, so that its current form represents a compromise between two incompatible more ideal forms. Or it can arise when the trait is under selection pressures for two different functions at once, these pulling towards different forms for it. Thus the size and shape of the peacock's famous tail is, famously, assumed to be an awkward compromise between the functions of helping it fly and attracting a mate. Or imagine that a perceptual organ, or a neural response to a kind of sensory input, was torn by selection pressures between accurately carrying two somewhat different, not-quite-extensionally-equivalent, kinds of information, one needed for one purpose, the other for another purpose. It would seem then to carry equivocal content. Any such equivocation would be local, however, not holistic.

Another kind of content indeterminacy resting on basic design might better be called "vagueness." What does the frog's eye tell the frog's tongue? I have argued that it doesn't say "little ambient black speck here now" or "shadow now crossing the retina," because neither of these conditions is directly causally implicated in the process that moves from a directional flick of the tongue to arrival at the next generation of frogs (Millikan 1991a; see also Elder 1998). Neither is of relevance in the use to which the proffered information is put. But what exactly is the causal explanation of the efficacy of the flicking of the tongue when it does help the frog? Does the cause involve presence of a fly, or of nutrition for a frog, or of protein molecules abc? How abstractly should the content be described? What counts as the correct causal explanation is vague in a way that probably cannot be eliminated in any principled way. But again, the problem is local, not holistic.

Much more interesting is indeterminacy in content that arises when a creature's perceptual/cognitive mechanisms labor under conditions that are not normal for proper performance of their functions. In the last quotation above, Dennett seemed to portray nature's sculpting of rational animals as aiming toward a single determinate ideal, the same for all intelligent species but, woefully, always missing that mark by a mile, even in our case. I propose instead, that those animals that are designed to collect together information from which to make inferences that govern certain of their behaviors, are all perfectly designed rational animals. On the other hand, they are generally designed to use only specified kinds of information in this way. Moreover, they are designed to operate in this way only given certain quite definite supporting conditions, and these supporting conditions often are absent. It is not even coherent to suppose an animal designed to exhibit real rationality patterns

8

For more recent reservations on this point, see (Millikan forthcoming; Schwartz forthcoming).

in all possible worlds. Each is designed to handle only certain actual kinds of information, available only through certain actual media, that have actually been prevalent enough in its historical environment. Each is designed to use this information in the production of responses that lead to results reasonable for the animal, but lead there only given supporting conditions that have been prevalent enough in its historical environment. Let me illustrate.

Perceptual representations produced by the human eye may often be vague, but there is no reason to suppose they are, in the usual case, equivocal. They can become equivocal, however, if certain supporting conditions that historically have nearly always been present in the species's environment are artfully removed. This happens, for example, when you look through a stereoscope. Then you "see just one picture," but which one do you see, the one shown to your right eye or the one shown to your left? The object you "see" is equivocal between these two. The Optometrist also has another piece of equipment through which you can see one picture as two, and with suitable preliminaries for adapting the eyes, the two may even appear different colors. Thus you can see (you can "visage" -- Millikan 1991b) a contradiction without any awareness of this.

Kripke's Pierre thinks a contradiction without any awareness of this. How does this come about? Under the right conditions, our perceptual/cognitive systems are remarkably accurate in their ability to keep track of the identities of individuals, of natural kinds, of natural stuffs, and so forth. Indeed, they are designed so that they progressively learn to keep better and better track of identities of these sorts. We may learn to recognize each individual that we know in more and more ways with experience, under more conditions, by a greater diversity of symptoms or signs, and so forth. We may learn to recognize each of the natural kinds and stuffs with which we are acquainted in more and more ways over time (Millikan 1998a, 1998b, 1998c). How this happens is properly studied by developmental psychologists (and by the philosophy and history of science). As with all other biological mechanisms, however, the cognitive mechanisms that accomplish these tasks are not designed to work in all possible worlds, but only in the kind of world humans evolved it. Moreover, as with all other biological mechanisms, they were selected not because they always worked right, but because they worked better than competitors, for example, because they worked right under more conditions, or under conditions that were more prevalent. There will always be conditions under which they can be made to fail. For a simple example, although we are especially talented at recognizing people by remembering their faces, this ability assumes an environment in which different people have faces sufficiently different in just those ways by which we are designed to tell faces apart. Identical twins can confuse us, and if person-cloning became very common, we might face quite a serious problem.

Where failures to reidentify correctly have occurred we may have two thoughts of one thing (Pierre's thoughts of London) or one thought of two things (identical twins we have confused, mass and weight) or thoughts that are so equivocal as to have no definite object at all (phlogiston). And, to be sure, when our inner representational systems are corrupted in this manner, the "thoughts" that we

have can no longer be unequivocally described in intentional idiom. But these problems are local, not holistic.⁹

Similar remarks apply to Anscomb's example of the person who says "Now I press button A" while reaching out and pressing button B (Anscomb 1993, p. 57, cited by Dennett 1990), and to Dennett's signaling "out" while he simultaneously shouted "safe" (Dennett 1990, p.181). But all this has nothing to do with holism.

There are more extreme cases in which intentional description becomes inappropriate.

No one is perfectly rational, perfectly unforgetful, all observant, or invulnerable to fatigue, malfunction or design imperfection. This leads inevitably to circumstances beyond the power of the intentional strategy to describe, in much the same way that physical damage to an artifact, such as a telephone or an automobile, may render it indescribable by the normal design terminology for that artifact. How do you draw the schematic wiring diagram for an audio amplifier that has been partially melted, or how do you characterize the program state of a malfunctioning computer? In cases of even the most familiar cognitive pathology—where people seem to hold contradictory beliefs or to be deceiving themselves, for instance—the canons of interpretation of the intentional strategy fail to yield clear, stable verdicts about which beliefs and desires to attribute to a person.

This is certainly true and important. Design stance attributions are of no help in describing either sufficiently mutilated artifacts or sufficiently mutilated organisms. But Dennett continues thus:

Now a strong realist position on beliefs and desires would claim that in these cases the person in question really does have some particular beliefs and desires which the intentional strategy, as I have described it, is simply unable to divine. On the milder sort of realism I am advocating, there is no fact of the matter exactly which beliefs and desires a person has in these degenerate cases... (Dennett 1987, p. 28)

But the strong realist should not make such a claim. The strong realist who takes the contents of intentional states to be determined by phylogenetic and ontogenetic history is free to reject numerous behavior-influencing states as failing to have any content at all, and others as failing to have determinate content. What the strong realist must think is only that it is a determinate matter which of these states are contentless, which determinate and which indeterminate.

If my remarks have been right, then we should expect it to be pretty definite for the most part what jobs the various inner parts and aspects of the perceptual and cognitive systems are designed to be doing, when and what they are supposed to be representing. In so far as they are working in accordance with design, rather than smashed or laboring under conditions that fail to support them properly, they will represent what they were designed to represent given the circumstances. Or they will represent what they have learned or been tuned, in accordance with design, to

⁹

I have argued that the problems tend to be local even in the case of developing scientific theory, See (Millikan 1998c).

represent given the circumstances. It is no accident that we think, for the most part, pretty unequivocal (though often rather vague) thoughts.

One more of Dennett's wonderfully apt examples deserves comment. A chess playing computer, he famously says, may be correctly described as thinking that it ought to get its queen out early if that is its constant disposition, even if its program contains no such instruction (Dennett 1978, p. 107). It makes a difference, I believe, whether the disposition is a logical or merely a causal result of design. If a logical result, then it seems to be true that the computer was indeed designed to get its Queen out early, even if that part of its design was not independently "selected for." Similarly, my digestive system was designed to digest, among other things, french fries and chocolate muosse, even though these aspects of its design were not independently selected for.

But there is also this kind of case. William has a way of insulting everyone he wants to impress by belittling their accomplishments over against his own. James has a way of winning people over by warmly admiring their children and pets, for he adores all children and animals perfectly sincerely and quite indiscriminately. These ways of William and James are not intentional any more than pointing your eyelashes toward your toes is purposive when you blink. These dispositions are real, and can form the basis of well- evidenced predictions of the effects of William's and James's behaviors. But systematic predictable effects of behaviors often do not have legitimate intentional descriptions or explanations. We all systematically and predictably depress the carpets on which we walk, and kick doctors who apply small rubber hammers just below our kneecaps. These effects are not correctly explained in intentional terms. Nor, without doubt, are a great many of the effects that traditional psychoanalysis attempted to explain in intentional terms.—Has anyone yet proposed a Freudian explanation of the (clearly retaliatory) knee jerk response?¹⁰

¹⁰

Thanks to Gunnar Björnsson for help with this essay.

References Cited

- Anscomb, G.E.M. Intention (Oxford: Blackwells, second edition)
- Dennett, D.C. 1978 Brainstorms (Bradford Books).
- Dennett, D.C. 1987 The Intentional Stance. Cambridge MA: MIT Press.
- Dennett, D.C. 1990 "The Interpretation of Texts, People, and Other Artifacts" Philosophy and Phenomenological Research, vol. 50: 177-94.
- Dennett, D.C. 1991 Consciousness Explained (Boston: Little, Brown).
- Dennett, D.C. 1994, "Get Real," Reply to My Critics, in Philosophical Topics, Vol 22, no. 1&2, Spring and Fall.
- Dennett, D.C. 1995 Darwin's Dangerous Idea (New York: Simon&Schuster)
- Dretske, F. 1988 Explaining behavior. Cambridge MA: MIT Press.
- Elder, C., 1998 "What versus How in Naturally Selected Representations," Mind 107: 349-363.
- Fodor and Lepore 1992 Holism: a Shopper's Guide (Cambridge: MIT Press)
- Gould, S.J., and R.C. Lewontin. 1979. The Spandrels of San Marco and the Panglossian program. Proceedings of the Royal Society of London 205:281-288. Reprinted in Sober 1984b.
- Gould, S.J., and E.S. Vrba. 1982. Exaptation--a missing term in the science of form. Paleobiology 8(1):4-15.
- Godfrey-Smith, P. 1996 Complexity and the Function of Mind in Nature. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. 1994 "A Modern History Theory of Functions" Nous 28: 344-362.
- Kripke, S.A. 1979. A puzzle about belief. In Meaning and Use, ed. A. Margalit, 239-83. Dordrecht: Reidel.
- Marr, D. 1982 Vision. San Francisco: W.H.Freeman and Co.
- Millikan, R.G. 1984 Language, Thought, and Other Biological Categories (Cambridge: MIT press)
- Millikan, R.G. 1991a "Speaking Up For Darwin", in eds. G. Rey and B. Loewer, Meaning in Mind; Fodor and His Critics, Oxford: Blackwell, pp. 151-164.
- Millikan, R.G. 1991b "Perceptual Content and Fregean Myth", Mind 100.4: 439-59.
- Millikan, R.G. 1993 White Queen Psychology and Other Essays for Alice (Cambridge: MIT Press).
- Millikan 1998a "A Common Structure for Concepts of Individuals, Stuffs, and Basic Kinds: More Mama, More Milk and More Mouse", Behavioral and Brain Sciences 22.1 (February 1998), pp. 55-65. Reprinted in E. Margolis and S. Laurence eds., Concepts: Core Readings, MIT Press 1999.
- Millikan, R.G. 1998b "With Enemies Like these I Don't Need Friends", Behavioral and Brain Sciences Vol 22.1 (February 1998) pp. 89-100.
- Millikan 1998c "How We Make Our Ideas Clear," The Tenth Annual Patrick Romanell Lecture, Proceedings and Addresses of the American Philosophical Association. November 1998
- Millikan, R.G. 1999 "Wings, Spoons, Pills, and Quills; A Pluralist Theory of Function" Journal of Philosophy
- Millikan, R.G. (forthcoming) "Bio-functions: Two Paradigms" in R. Cummins, A. Ariew and M. Perlman, eds, Functions in Philosophy of Biology and Philosophy of Psychology, Oxford University Press.

- ????Schwartz, P. "Proper Function and Recent Selection" in R. Cummins, A. Ariew and M. Perlman, eds, Functions in Philosophy of Biology and Philosophy of Psychology, Oxford University Press.
- Sober, E. 1984a. The Nature of Selection. Cambridge MA: MIT Press.
- Sober, E. 1984b. Conceptual Issues in Evolutionary Biology. Cambridge MA: Bradford Books/MIT Press.