MACHIAVELLIAN INTELLIGENCE

Oxford University Press
1988

14

The intentional stance in theory and practice

DANIEL C. DENNETT

In the first part of this chapter, one of today's most exciting philosophers elegantly charts the increasing 'levels' (embedding) of intentionality which may in principle underly primate vocal behaviour, and suggests a simple method for picking out the real level. Then, in the second part, he calls his own bluff. He visits the very primatologists on whose data he theorized, and discusses the difficulties in executing his simple test in practice.*

Section 1: Intentional systems in cognitive ethology

I. The problem

The field of cognitive ethology provides a rich source of material for the philosophical analysis of meaning and mentality, and even holds out some tempting prospects for philosophers to contribute fairly directly to the development of the concepts and methods of another field. As a philosopher, an outsider with only a cursory introduction to the field of ethology, I find that the new ethologists, having cast off the straitjacket of behaviourism and kicked off its weighted overshoes, are looking about somewhat insecurely for something presentable to wear. They are seeking a theoretical vocabulary that is powerfully descriptive of the data they are uncovering and at the same time a theoretically fruitful method of framing hypotheses that will eventually lead to information-processing models of the nervous systems of the creatures they are studying (see Roitblat 1982). It is a long way from the observation of the behaviour of, say, primates in the wild to the validation of neurophysiological models of their brain activity, and finding a sound interim way of speaking is not a trivial task. Since the methodological and conceptual problems confronting the ethologists appear to me to bear

* Section 1 extracted from *Behavioral and Brain Sciences*, 3, 343–350 (1983).

striking resemblances to problems I and other philosophers have been grappling with recently, I am tempted to butt in and offer, first, a swift analysis of the problem; secondly, a proposal for dealing with it (which I call intentional system theory); and thirdly, an analysis of the continuity of intentional system theory with the theoretical strategy or attitude in evolutionary theory often called adaptationism.

The methodology of philosophy, such as it is, includes as one of its most popular (and often genuinely fruitful) strategies the description and examination of entirely imaginary situations, elaborate thought experiments that isolate for scrutiny the presumably critical features in some conceptual domain. In *Word and Object*, W. V. O. Quine (1960), gave us an extended examination of the evidential and theoretical tasks facing the 'radical translator', the imaginary anthropologist-linguist who walks into an entirely alien community—with no string of interpreters or bilingual guides—and who must figure out, using whatever scientific methods are available, the language of the natives. Out of this thought experiment came Quine's thesis of the 'indeterminacy of radical translation', the claim that it must always be possible in principle to produce nontrivially different translation manuals, equally well supported by all the evidence, for any language. One of the most controversial features of Quine's position over the years has been his uncompromisingly behaviourist scruples about how to characterize the task facing the radical translator. What happens to the task of radical translation when you give up the commitment to a behaviouristic outlook and terminology? What are the prospects for fixing on a unique translation of a language (or a unique interpretation of the 'mental states' of a being) if one permits oneself the vocabulary and methods of 'cognitivism'? The question could be explored via other thought experiments, and has been in some regards (Bennett 1976; Dennett 1971: Lewis 1974), but the real-world researches of Seyfarth *et al.* (1980*a*) with vervet monkeys in Africa will serve us better on this occasion. Vervet monkeys form societies of sorts and have a language of sorts, and of course, there are no bilingual interpreters to give a boost to the radical translators of Vervetese. This is what they find:

> Vervet monkeys give different alarm calls to different predators. Recordings of the alarms played back when predators were absent caused the monkeys to run into the trees for leopard alarms, look up for eagle alarms, and look down for snake alarms. Adults call primarily to leopards, martial eagles, and pythons, but infants give leopard alarms to various mammals, eagle alarms to many birds, and snake alarms to various snakelike objects. Predator classification improves with age and experience. (Seyfarth *et al.* 1980, p. 801.)

This abstract is couched, you will note, in almost pure Behaviourese—the language of *Science* even if it is no longer exclusively the language of

science. It is just informative enough to be tantalizing. How much of a language, one wants to know, do the vervets really have? Do they really communicate? Do they mean what they say? Just what interpretation can we put on these activities? What, if anything, do these data tell us about the cognitive capacities of vervet monkeys? In what ways are they—must they be—like human cognitive capacities, and in what ways and to what degree are vervets more intelligent than other species by virtue of these 'linguistic' talents? These loaded questions—the most natural ones to ask under the circumstances—do not fall squarely within the domain of any science, but whether or not they are the right questions for the scientist to ask, they are surely the questions that we all, as fascinated human beings learning of this apparent similarity of the vervets to us, want answered.

The cognitivist would like to succumb to the temptation to use ordinary mentalistic language more or less at face value, and to respond directly to such questions as: What do the monkeys know? What do they want, and understand, and mean? At the same time, the primary point of the cognitivists' research is not to satisfy the layman's curiosity about the relative IQ, as it were, of his simian cousins, but to chart the cognitive talents of these animals on the way to charting the cognitive *processes* that explain these talents. Could the everyday language of belief, desire, expectation, recognition, understanding, and the like also serve as the suitably rigorous abstract language in which to describe cognitive competences?

I will argue that the answer is yes. Yes, if we are careful about what we are doing and saying when we use ordinary words like 'believe' and 'want', and if we understand the assumptions and implications of the strategy we must adopt when we use these words.

The decision to conduct one's science in terms of beliefs, desires, and other 'mentalistic' notions, the decision to adopt 'the intentional stance', as I call it (Dennett 1971, 1976, 1978*a*, 1981*a*, *b*, *c*), is not an unusual sort of decision in science. The basic strategy of which this is a special case is familiar: changing levels of explanation and description in order to gain access to greater predictive power or generality—purchased, typically, at the cost of submerging detail and courting trivialization on the one hand and easy falsification on the other. When biologists studying some species choose to call something in that species' environment food and leave it at that, they ignore the tricky details of the chemistry and physics of nutrition, the biology of mastication, digestion, excretion, and the rest. Even supposing many of the details of this finer-grained biology are still ill-understood, the decision to leap ahead, in anticipation of fine-grained biology, and rely on the well-behavedness of the concept of food at the level of the theory appropriate to it, is likely to meet approval from the most conservative risk takers.

The decision to adopt the intentional stance is riskier. It banks on the soundness of some as yet imprecisely described concept of information—not the concept legitimized by Shannon–Weaver information theory (Shannon 1949), but rather the concept of what is often called semantic information. (A more or less standard way of introducing the still imperfectly understood distinction between these two concepts of information is to say that Shannon–Weaver theory measures the capacity of information-transmission and information-storage vehicles, but is mute about the contents of those channels and vehicles, which will be the topic of the still-to-be-formulated theory of semantic information. See Dretske 1981 [and multiple book review in *BBS* 6(1) 1983] for an attempt to bridge the gap between the two concepts.) Information, in the semantic view, is a perfectly real, but very abstract commodity, the storage, transmission, and transformation of which is informally—but quite sure-footedly—recounted in ordinary talk in terms of beliefs and desires and the other states and acts philosophers call intentional.

II. Intentional system theory

Intentionality, in philosophical jargon, is—in a word—aboutness. Some of the things, states, and events in the world have the interesting property of being about other things, states, and events; figuratively, they point to other things. This arrow of reference or aboutness has been subjected to intense philosophical scrutiny and has engendered much controversy. For our purposes, we can gingerly pluck two points from this boiling cauldron, oversimplifying them and ignoring important issues tangential to our concerns.

First, we can mark the presence of intentionality—aboutness—as the topic of our discussions by marking the presence of a peculiar logical feature of all such discussion. Sentences attributing intentional states or events to systems use idioms that exhibit *referential opacity*: they introduce clauses in which the normal, permissive, substitution rule does not hold: This rule is simply the logical codification of the maxim that a rose by any other name would smell as sweet. If you have a true sentence, so runs the rule, and you alter it by replacing a term in it by another, different term that still refers to exactly the same thing or things, the new sentence will also be true. Ditto for false sentences—merely changing the means of picking out the objects the sentence is about cannot turn a falsehood into a truth. For instance, suppose Bill is the oldest child in class; then if it is true that

(1) Mary is sitting next to Bill,
then, substituting 'the oldest child in class' for 'Bill', we get
(2) Mary is sitting next to the oldest child in class,
which must be true if the other sentence is.

A sentence with an *intentional idiom* in it, however, contains a clause in which such substitution can turn truth into falsehood and vice versa. (This phenomenon is called referential opacity because the terms in such clauses are shielded or insulated by a barrier to logical analysis, which normally 'sees through' the terms to the world the terms are about.) For example, Sir Walter Scott wrote *Waverly*, and Bertrand Russell (1905) assures us

(3) George IV wondered whether Scott was the author of *Waverly*,

but it seems unlikely indeed that

(4) George IV wondered whether Scott was Scott.

(As Russell remarks, 'An interest in the law of identity can hardly be attributed to the first gentleman of Europe; 1905, p. 485.) To give another example, suppose we decide it is true that

(5) Burgess fears that the creature rustling in the bush is a python,

and suppose that in fact the creature in the bush is Robert Seyfarth. We will not want to draw the conclusion that

(6) Burgess fears that Robert Seyfarth is a python.

Well, in one sense we do, you say, and in one sense we also want to insist that, oddly enough, King George was wondering whether Scott was Scott. But that's not how he put it to himself—and that's not how Burgess conceived of the creature in the bush, either—that is, as Seyfarth. It's the sense of conceiving as, seeing as, thinking of as that the intentional idioms focus on.

One more example. Suppose you think your next-door neighbour would make someone a good husband and suppose, unbeknownst to you, he's the Mad Strangler. Although in one very strained sense you could be said to believe that the Mad Strangler would make someone a good husband, in another more natural sense you don't, for there is another—very bizarre and unlikely—belief that you surely don't have which could better be called the belief that the Mad Strangler would make a good husband.

It is this resistance to substitution, the insistence that for some purposes how you call a rose a rose makes all the difference, that makes the intentional idioms ideally suited for talking about the ways in which information is represented in the heads of people—and other animals.

So the first point about intentionality is just that we can rely on a marked set of idioms to have this special feature of being sensitive to the *means of reference* used in the clauses they introduce. The most familiar such idioms are 'believes that', 'knows that', 'expects (that)', 'wants (it to be the case that)', 'recognizes (that)', 'understands (that)'.

In short, the 'mentalistic' vocabulary shunned by behaviourists and celebrated by cognitivists is quite well picked out by the logical test for referential opacity.

The second point to pluck from the cauldron is somewhat controversial, although it has many adherents who have arrived at roughly the same conclusion by various routes: the use of intentional idioms carries a presupposition or assumption of *rationality* in the creature or system to which the intentional states are attributed. What this amounts to will become clearer if we now turn to the intentional stance in relation to the vervet monkeys.

III. Vervet monkeys as intentional systems

To adopt the intentional stance toward these monkeys is to decide—tentatively, of course—to attempt to characterize, predict, and explain their behaviour by using intentional idioms, such as 'believes' and 'wants', a practice that assumes or presupposes the rationality of the vervets. A vervet monkey is, we will say, an *intentional system*, a thing whose behaviour is predictable by attributing beliefs and desires (and, of course, rationality) to it. Which beliefs and desires? Here there are many hypotheses available, and they are testable in virtue of the rationality requirement. First, let us note that there are different grades of intentional systems.

A *first-order* intentional system has beliefs and desires (etc.), but no beliefs and desires about beliefs and desires. Thus, all the attributions we make to a merely first-order intentional system have the logical form of:

(7) x *believes that* p;
(8) y *wants that* q;

where 'p' and 'q' are clauses that themselves contain no intentional idioms. A *second-order* intentional system is more sophisticated; it has beliefs and desires (and no doubt other intentional states) about beliefs and desires (and other intentional states)—both those of others and its own. For instance,

(9) x *wants* y to *believe* that x is hungry;
(10) x *believes* y *expects* x to jump left;
(11) x *fears* that y *will discover* that x has a food cache.

A *third-order* intentional system is one that is capable of such states as

(12) x *wants* y to *believe* that x *believes* he is all alone.

A *fourth-order* system might want you to think it understood you to be requesting that it leave. How high can we human beings go? 'In principle', forever, no doubt, but in fact I suspect that you wonder whether I realize

how hard it is for you to be sure that you understand whether I mean to be saying that you can recognize that I can believe you to want me to explain that most of us can keep track of only about five or six orders, under the best of circumstances. (See Cargile (1970) for an elegant and sober exploration of this phenomenon.)

How good are vervet monkeys? Are they really capable of third-order or higher-order intentionality? The question is interesting on several fronts. First, these orders ascend what is intuitively a scale of intelligence: higher-order attributions strike us as much more sophisticated, much more human, requiring much more intelligence. There are some plausible diagnoses of this intuition. Grice (1957, 1969) and other philosophers (see especially Bennett 1976) have developed an elaborate and painstakingly argued case for the view that genuine *communication*, speech acts in the strong, human sense of the word, depend on at least three orders of intentionality in both speaker and audience.

Not all interactions between organisms are communicative. When I swat a fly I am not communicating with it, nor am I if I open the window to let it fly away. Does a sheep dog, though, communicate with the sheep it herds? Does a beaver communicate by slapping its tail, and do bees communicate by doing their famous dances? Do human infants communicate with their parents? At what point can one be sure one is really communicating with an infant? The presence of specific linguistic tokens seems neither sufficient nor necessary. (I can use English commands to get my dog to do things, but that is at best a pale form of communication compared to the mere raised eyebrow by which I can let someone know he should change the topic of our conversation.) Grice's theory provides a better framework for answering these questions. It defines intuitively plausible and formally powerful criteria for communication that involve, at a minimum, the correct attribution to communicators of such third-order intentional states as

(13) Utterer *intends* Audience to *recognize* that Utterer *intends* Audience to produce response *r*.

So one reason for being interested in the intentional interpretation of the vervets is that it promises to answer—or at least help answer—the questions: Is this behaviour really linguistic? Are they really communicating? Another reason is that higher-orderedness is a conspicuous mark of the attributions speculated about in the sociobiological literature about such interactive traits as reciprocal altruism. It has even been speculated (by Trivers 1971), that the increasing complexity of mental representation required for the maintenance of systems of reciprocal altruism (and other complex social relations) led, in evolution, to a sort of brain-power arms race. Humphrey (Chapter 2) arrives at similar conclusions by a different

and in some regards less speculative route. There may then be a number of routes to the conclusion that higher-orderedness of intentional characterization is a deep mark—and not just a reliable symptom—of intelligence.

(I do not mean to suggest that these orders provide a uniform scale of any sort. As several critics have remarked to me, the first iteration—to a second-order intentional system—is the crucial step of the recursion; once one has the principle of embedding in one's repertoire, the complexity of what one can then in some sense entertain seems plausibly more a limitation of memory, or attention span, or 'cognitive workspace' than a fundamental measure of system sophistication. Thanks to 'chunking' and other, artificial, aids to memory, there seems to be no interesting difference between, say, a fourth-order and a fifth-order intentional system; but see Cargile 1970 for further reflections on the natural limits of iteration.)

However, back to the empirical question of how good the vervet monkeys are. For simplicity's sake, we can restrict our attention to a single apparently communicative act by a particular vervet, Tom, who, let us suppose, gives a leopard alarm call in the presence of another vervet, Sam. We can now compose a set of competing intentional interpretations of this behaviour, ordered from high to low, from romantic to killjoy. Here is a (relatively) romantic hypothesis (with some variations to test in the final clause):

4th-order: Tom *wants* Sam to *recognize* that Tom *wants* Sam to *believe* that
there is a leopard
there is a carnivore
there is a four-legged animal
there is a live animal bigger than a breadbox

A less exciting hypothesis to confirm would be this third-order version (there could be others):

3rd-order: Tom *wants* Sam to *believe* that Tom *wants* Sam to run into the trees.

Note that this particular third-order case differs from the fourth-order case in changing the speech act category: on this reading the leopard call is an imperative (a request or command) not a declarative (informing Sam of the leopard). The important difference between imperative and declarative interpretations (see Bennett 1976, pp. 41, 51) of utterances can be captured—and then telltale behavioural differences can be explored—at any level of description above the second order—at which, *ex hypothesi*, there is no intention to utter a speech act of either variety. Even at the second order, however, a related distinction in effect-desired-in-the-Audience is expressed, and is in principle behaviourally detectable, in the following variations:

2nd-order: Tom *wants* Sam to *believe* that
there is a leopard
he should run into the trees

This differs from the previous two in not supposing Tom's act involves ('in Tom's mind') any recognition by Sam of his (Tom's) own role in the situation. If Tom could accomplish his end equally well by growling like a leopard, or just somehow attracting Sam's attention to the leopard without Sam's recognizing Tom's intervention, this would be only a second-order case. (Cf. I *want* you to *believe* I am not in my office; so I sit very quietly and don't answer your knock. That is not communicating.)

1st-order: Tom *wants* to cause Sam to run into the trees (and he has this noise-making trick that produces that effect; he uses the trick to induce a certain response in Sam).

On this reading the leopard cry belongs in the same general category with coming up behind someone and saying 'Boo!'. Not only does its intended effect not depend on the victim's recognition of the perpetrator's intention; the perpetrator does not need to have any conception at all of the victim's mind: Making loud noises behind certain things just makes them jump.

0-order: Tom (like other vervet monkeys) is prone to three flavours of anxiety or arousal: leopard anxiety, eagle anxiety, and snake anxiety.[1] Each has its characteristic symptomatic vocalization. The effects on others of these vocalizations have a happy trend, but it is all just tropism, in both utterer and audience.

We have reached the killjoy bottom of the barrel: an account that attributes no mentality, no intelligence, no communication, no intentionality at all to the vervet. Other accounts at the various levels are possible, and some may be more plausible; I chose these candidates for simplicity and vividness. Lloyd Morgan's canon of parsimony enjoins us to settle on the most killjoy, least romantic hypothesis that will account systematically for the observed and observable behaviour, and for a long time the behaviourist creed that the curves could be made to fit the data well at the lowest level prevented the exploration of the case that can be made for higher-order, higher-level systematizations of the behaviour of such animals. The claim that, in principle, a lowest-order story can always be told of any animal behaviour (an entirely physiological story, or even an abstemiously behaviouristic story of unimaginable complexity) is no longer interesting. It is like claiming that in principle the concept of food can be ignored by biologists—or the concept of cell or gene for that matter—or like claiming that in principle a purely electronic-level story can be told of any computer behaviour. Today we are interested in asking what gains in perspicuity, in predictive power, in generalization, might accrue if we

[1] We can probe the boundaries of the stimulus-equivalence class for this response by substituting for the 'normal' leopard such different 'stimuli' as dogs, hyenas, lions, stuffed leopards, caged leopards, leopards dyed green, firecrackers, shovels, motorcyclists. Whether these independent tests are tests of *anxiety specificity* or of the *meaning* of one-word sentences of Vervetese depends on whether our tests for the other components of our *n*th-order attribution, the nested intentional operators, come out positive.

adopt a higher-level hypothesis that takes a risky step into intentional characterization.

The question is empirical. The tactic of adopting the intentional stance is not a matter of replacing empirical investigations with aprioristic ('armchair') investigations, but of using the stance to suggest which brute empirical questions to put to nature. We can test the competing hypothesis by exploiting the rationality assumption of the intentional stance. We can start at either end of the spectrum; either casting about for the depressing sorts of evidence that will demote a creature from a high-order interpretation, or hunting for the delighting sorts of evidence that promote creatures to higher-order interpretations (cf. Bennett 1976). We are delighted to learn, for instance, that lone male vervet monkeys, travelling between bands (and hence out of the hearing—so far as they know—of other vervets) will, on seeing a leopard, silently seek refuge in the trees. So much for the killjoy hypothesis about leopard-anxiety yelps. (No hypothesis succumbs quite so easily, of course. *Ad hoc* modifications can save any hypothesis, and it is an easy matter to dream up some simple 'context' switches for leopard-anxiety yelp mechanisms to save the zero-order hypothesis for another day.) At the other end of the spectrum, the mere fact that vervet monkeys apparently have so few different things they can say holds out little prospect for discovering any real theoretical utility for such a fancy hypothesis as our fourth-order candidate. It is only in contexts or societies in which one must rule out (or in) such possibilities as irony, metaphor, storytelling, and illustration ('second-intention' uses of words, as philosophers would say)[2] that we must avail ourselves of such high-powered interpretations. The evidence is not yet in, but one would have to be romantic indeed to have high expectations here. Still, there are encouraging anecdotes.

Seyfarth reports (personal communication) an incident in which one band of vervets was losing ground in a territorial skirmish with another band. One of the losing-side monkeys, temporarily out of the fray, seemed to get a bright idea: it suddenly issued a leopard alarm (in the absence of any leopards), leading all the vervets to take up the cry and head for the trees—creating a truce and regaining the ground his side had been losing (see Chapter 16). The intuitive sense we all have that this is possibly (barring killjoy reinterpretation) an incident of great cleverness is amenable to a detailed diagnosis in terms of intentional systems. If this act is not just a lucky coincidence, then the act is truly devious, for it is not simply a case of the vervet uttering an imperative 'get into the trees' in the expectation that all the vervets will obey, since the vervet (being rational—our predictive lever) should not expect a rival band to honour *his*

[2] See Quine (1960) pp. 48–49, on second-intention cases as the 'bane of theoretical linguistics'.

imperative. So either the leopard call is considered by the vervets to be informative—a warning, not a command—and hence the utterer's credibility, but not authority, is enough to explain the effect, or our utterer is more devious still: he *wants* the rivals to *think* they are *overhearing* a command *intended* (of course) only for his own folk, and so on. Could a vervet possibly have that keen a sense of the situation? These dizzying heights of sophistication are strictly implied by the higher-order interpretation taken with its inevitable presupposition of rationality. Only a creature capable of appreciating these points could properly be said to have those beliefs and desires and intentions.

Another observation of the vervets brings out this role of the rationality assumption even more clearly. When I first learned that Seyfarth's methods involved hiding speakers in the brush and playing recorded alarm calls, I viewed the very success of the method as a seriously demoting datum, for if the monkeys really were Gricean in their sophistication, when playing their audience roles they should be perplexed, unmoved, somehow disrupted by disembodied calls issuing from no known utterer. If they were oblivious to this problem, they were no Griceans. Just as a genuine Communicator typically checks the Audience periodically for signs that it is getting the drift of the communication, a genuine Audience typically checks out the Communicator periodically for signs that the drift it is getting is the drift being delivered.

To my delight, however, I learned from Seyfarth that great care had been taken in the use of the speakers to prevent this sort of case from arising. Vervets can readily recognize the particular calls of their band—thus they recognize Sam's leopard call as Sam's, not Tom's. Wanting to give the recordings the best chance of 'working', the experimenters took great care to play, say, Sam's call only when Sam was neither clearly in view and close-mouthed or otherwise occupied, nor 'known' by the others to be far away. Only if Sam could be 'supposed' by the audience to be actually present and uttering the call (though hidden from their view), only if the audience could believe that the noisemaker in the bush was Sam, would the experimenters play Sam's call. While this remarkable patience and caution are to be applauded as scrupulous method, one wonders whether they were truly necessary. If a 'sloppier' scheduling of playbacks produced just as 'good' results, this would in itself be a very important *demoting* datum. Such a test should be attempted; if the monkeys are baffled and unmoved by recorded calls except under the scrupulously maintained circumstances, the necessity of those circumstances would strongly support the claim that Tom, say, does believe that the noisemaker in the bush is Sam, that vervet monkeys are not only capable of believing such things, but must believe such things for the observed reaction to occur.

The rationality assumption thus provides a way of taking the various hypotheses seriously—seriously enough to test. We expect at the outset that there are bound to be grounds for the verdict that vervet monkeys are believers only in some attenuated way (compared to us human believers). The rationality assumption helps us look for, and measure, the signs of attenuation. We frame conditionals such as

(14) if x believed that p, and *if x was rational*, then since 'p' implies 'q' x would (have to) believe that q.

This leads to the further attribution to x of belief that q,[3] which, coupled with some plausible attribution of desire, leads to a prediction of behaviour, which can be tested by observation or experiment.[4]

Once one gets the knack of using the rationality assumption for leverage, it is easy to generate further telling behaviours to look for in the wild or to provoke in experiments. For instance, if anything as sophisticated as a third- or fourth-order analysis is correct, then it ought to be possible, by devious (and morally dubious!) use of the hidden speakers to create a 'boy who cried wolf'.[5] If a single vervet is picked out and 'framed' as the utterer of false alarms, the others, being rational, should begin to lower their trust in him, which *ought* to manifest itself in a variety of ways. Can a 'credibility gap' be created for a vervet monkey? Would the potentially nasty results (remember what happened in the fable) be justified by the interest such a positive result would have? (Editors' Note: this has now been done, see Chapter 6.)

IV. How to use anecdotal evidence: the Sherlock Holmes method

One of the recognized Catch 22s of cognitive ethology is the vexing problem of anecdotal evidence. On the one hand, as a good scientist, the ethologist knows how misleading and, officially, unusable anecdotes are, and yet on the other hand they are often so telling! The trouble with the canons of scientific evidence here is that they virtually rule out the description of anything but the oft-repeated, oft-observed, stereotypic behaviour of a species, and this is just the sort of behaviour that reveals no particular intelligence at all—all this behaviour can be more or less plausibly explained as the effects of some humdrum combination of

[3] 'I shall always treasure the visual memory of a very angry philosopher, trying to convince an audience that "if you believe that A and you believe that if A then B then you *must* believe that B". I don't really know whether he had the moral power to coerce anyone to believe that B, but a failure to comply does make it quite difficult to use the word "belief", and that is worth shouting about' (Kahneman 1982).

[4] The unseen normality of the rationality assumption in any attribution or belief is revealed by noting that (14), which explicitly assumes rationality, is virtually synonymous with (plays the same role as) the conditional beginning: if x *really* believed that p, then since 'p' implies 'q'. . .

[5] I owe this suggestion to Susan Carey, in conversation.

'instinct' or tropism and conditioned response. It is the novel bits of behaviour, the acts that couldn't plausibly be accounted for in terms of prior conditioning or training or habit, that speak eloquently of intelligence; but if their very novelty and unrepeatability make them anecdotal and, hence, inadmissible evidence, how can one proceed to develop the cognitive case for the intelligence of one's target species?

Just such a problem has bedevilled Premack and Woodruff (1978), for instance, in their attempts to demonstrate that chimps 'have a theory of mind'; their scrupulous efforts to force their chimps into nonanecdotal, repeatable behaviour that manifests the intelligence they believe them to have engenders the frustrating side effect of providing prolonged training histories for the behaviourists to point to in developing their rival, conditioning hypotheses as putative explanations of the observed behaviour. [See the commentaries and replies in: 'Cognition and Consciousness in Nonhuman Species' (*BBS* 1(4) 1978; see also Premack: 'The Codes of Man and Beasts' *BBS* 6(1) 1983; and Chapter 13].

We can see the way out of this quandary if we pause to ask ourselves how we establish our own higher-order intentionality to the satisfaction of all but the most doctrinaire behaviourists. We can concede to the behaviourists that any single short stretch of human behaviour can be given a relatively plausible and not obviously *ad hoc* demoting explanation, but as we pile anecdote upon anecdote, apparent novelty upon apparent novelty, we build up for each acquaintance such a biography of apparent cleverness that the claim that it is all just lucky coincidence—or the result of hitherto undetected 'training'—becomes the more extravagant hypothesis. This accretion of unrepeatable detail can be abetted by using the intentional stance to provoke one-shot circumstances that will be particularly telling. The intentional stance is in effect an engine for generating or designing anecdotal circumstances—ruses, traps, and other intentionalistic litmus tests—and predicting their outcomes.

This tricky tactic has long been celebrated in literature. The idea is as old as Odysseus testing his swineherd's loyalty by concealing his identity from him and offering him temptations. Sherlock Holmes was a master of more intricate intentional experiments, so I shall call this the *Sherlock Holmes method*. Cherniak (1981) draws our attention to a nice case:

> In 'A Scandal in Bohemia', Sherlock Holmes' opponent has hidden a very important photograph in a room, and Holmes wants to find out where it is. Holmes has Watson throw a smoke bomb into the room and yell 'fire' when Holmes' opponent is in the next room, while Holmes watches. Then, as one would expect, the opponent runs into the room and takes the photograph from where it was hidden. Not everyone would have devised such an ingenious plan for manipulating an opponent's behaviour; but once the conditions are described, it seems very easy to predict the opponent's actions. (p. 161.)

In this instance Holmes simultaneously learns the location of the photograph and confirms a rather elaborate intentional profile of his opponent, Irene Adler, who is revealed to *want* the photograph; to *believe* it to be located where she goes to get it; to *believe* that the person who yelled 'fire' *believed* there was a fire (note that if she believed the yeller wanted to deceive her, she would take entirely different action); to *want* to retrieve the photograph without letting anyone *know* she was doing this, and so on.

A variation on this theme is an intentional tactic beloved of mystery writers: provoking the telltale move. All the suspects are gathered in the drawing room, and the detective knows (and he alone knows) that the guilty party (and only the guilty party) *believes* that an incriminating cuff link is under the gateleg table. Of course the culprit *wants* no one else to *believe* this, or to *discover* the cuff link, and *believes* that in due course it will be discovered unless he takes covert action. The detective arranges for a 'power failure'; after a few seconds of darkness the lights are switched on and the guilty party is, of course, the chap on his hands and knees under the gateleg table. What else on earth could conceivably explain this novel and bizarre behaviour in such a distinguished gentleman.?[6]

Similar stratagems can be designed to test the various hypotheses about the beliefs and desires of vervet monkeys and other creatures. These stratagems have the virtue of provoking novel but interpretable behaviour, of *generating anecdotes* under controlled (and hence scientifically admissible) conditions. Thus, the Sherlock Holmes method offers a significant increase in investigative power over behaviourist methods. This comes out dramatically if we compare the actual and contemplated research on vervet monkey communication with the efforts of Quine's imagined behaviouristic field linguist. According to Quine, a necessary preliminary to any real

[6] It is a particular gift of the playwright to devise circumstances in which behaviour—verbal and otherwise—speaks loudly and clearly about the intentional profiles ('motivation', beliefs, misunderstandings, and so forth) of the characters, but sometimes these circumstances grow too convoluted for ready comprehension; a very slight shift in circumstance can make all the difference between utterly inscrutable behaviour, and lucid self-relevation. The notorious 'get thee to a nunnery' speech of Hamlet to Ophelia is a classic case in point. Hamlet's lines are utterly bewildering until we hit upon the fact (obscured in Shakespeare's minimal stage directions) that while Hamlet is speaking to Ophelia, he *believes* not only that Claudius and Polonius are listening behind the arras, but that they *believe* he doesn't *suspect* that they are. What makes this scene particularly apt for our purposes is the fact that it portrays an intentional experiment: Claudius and Polonius, using Ophelia as decoy and prop, are attempting to provoke a particularly telling behaviour from Hamlet in order thereby to discover just what his beliefs and intentions are; they are foiled by their failure to design the experiment well enough to exclude from Hamlet's intentional profile the belief that he is being observed, and the desire to create false beliefs in his observers. See, for example, Dover Wilson (1951). A similar difficulty can bedevil ethologists: 'Brief observations of avocet and stilt behavior can be misleading. Underestimating the bird's sharp eyesight, early naturalists believed their presence was undetected and misinterpreted distraction behavior as courtship' (Sordahl 1981, p. 45).

progress by the linguist is the tentative isolation and identification of native words (or speech acts) for 'Yes' and 'No', so that the linguist can enter into a tedious round of 'query-and-assent'—putting native sentences to co-operative natives under varying conditions and checking for patterns in their yes and no responses (Quine 1960, Chap. 2). Nothing just like Quine's game of query-and-assent can be played by ethologists studying animals, but a vestige of this minimalist research strategy is evident in the patient explorations of 'stimulus substitution' for animal vocalizations—to the exclusion, typically, of more manipulative (if less intrusive) experiments (see note 1). So long as one is resolutely behaviouristic, however, one must miss the evidential value of such behaviour as the lone vervet quietly taking to the trees when a 'leopard stimulus' is presented. However, without a goodly amount of such telling behaviour, no mountain of data on what Quine calls the 'stimulus meaning' of utterances will reveal that they are communicative acts, rather than merely audible manifestations of peculiar sensitivities. Quine of course realizes this, and tacitly presupposes that his radical translator has already informally satisfied himself (no doubt by using the powerful, but everyday, Sherlock Holmes method) of the richly communicative nature of the natives' behaviour.

Of course the power of the Sherlock Holmes method cuts both ways; failure to perform up to expectations is often a strongly demoting datum.[7] Woodruff and Premack (1979) have tried to show that chimpanzees in their lab can be full-fledged *deceivers*. Consider Sadie, one of four chimps used in this experiment. In Sadie's sight, food is placed in one of two closed boxes she cannot reach. Then either a 'co-operative' or a 'competitive' trainer enters, and Sadie has learned she must point to one of the boxes in hopes of getting the food. The competitive trainer, if he discovers the food, will take it all himself and leave. The co-operative trainer shares the food with Sadie. Just giving Sadie enough experience with the circumstances to assure her appreciation of these contingencies involves training sessions that give the behaviourist plenty of grist for the 'mere reinforcement' mill. (In order to render the identities of the trainers sufficiently distinct, there was strict adherence to special costumes and rituals; the competitive trainer always wore sunglasses and a bandit's mask, for instance. Does the mask then become established as a simple 'eliciting stimulus' for the tricky behaviour?)

[7] I do not wish to be interpreted as *concluding* in this paper that vervet monkeys, or laboratory chimpanzees, or any non-human animals have *already been shown* to be higher-order intentional systems. Once the Sherlock Holmes method is applied with imagination and rigour, it may very well yield results that will disappoint the romantics. I am arguing in favour of a method of raising empirical questions, and explaining the method by showing what the answers *might be* (and why); I am not giving those answers in advance of the research.

Still, setting behaviourists' redescriptions aside, will Sadie rise to the occasion and do the 'right' thing? Will she try to deceive the competitive trainer (and only the competitive trainer) by *pointing to the wrong box*? Yes, but suspicions abound about the interpretation.[8] How could we strengthen it? Well if Sadie really intends to deceive the trainer, she must (being rational) start with the belief that the trainer does not already know where the food is. Suppose, then, we introduce all the chimps in an entirely different context to transparent plastic boxes; they should come to know that since they—and anyone else—can see through them, anyone can see, and hence come to know, what is in them. Then on a one-trial, novel behavioural test, we can introduce a plastic box and an opaque box one day, and place the food in the plastic box. The competitive trainer then enters, and lets Sadie see him looking right at the plastic box. If Sadie still points to the opaque box, she reveals,sadly, that she really doesn't have a grasp of the sophisticated ideas involved in deception. Of course, this experiment is still imperfectly designed. For one thing, Sadie might point to the opaque box out of despair, seeing no better option. To improve the experiment, an option should be introduced that would appear better to her only if the first option was hopeless, as in this case. Moreover, shouldn't Sadie be puzzled by the competitive trainer's curious behaviour? Shouldn't it bother her that the competitive trainer, on finding no food where she points, just sits in the corner and 'sulks' instead of checking out the other box? Shouldn't she be puzzled to discover that her trick keeps working? She should wonder. Can the competitive trainer be that stupid? Furthermore, better-designed experiments with Sadie—and other creatures—are called for.[9]

Not wanting to feed the galling stereotype of the philosopher as an armchair answerer of empirical questions, I will nevertheless succumb to the temptation to make a few predictions. It will turn out on further exploration that vervet monkeys (and chimps and dolphins, and all other

[8] It is all too easy to stop too soon in our intentional interpretation of a presumably 'lower' creature. There was once a village idiot who, whenever he was offered a choice between a dime and a nickel, unhesitatingly took the nickel—to the laughter and derision of the onlookers. One day someone asked him if he could be so stupid as to continue choosing the nickel after hearing all that laughter. Replied the idiot, 'Do you think that if I ever took the dime they'd ever offer me another choice?'

The curiously unmotivated rituals that attended the training of the chimps as reported in Woodruff and Premack (1979) might well have baffled the chimps for similar reasons. Can a chimp wonder why these human beings don't just eat the food that is in their control? If so, such a wonder could overwhelm the chimps' opportunities to understand the circumstance in the sense the researchers were hoping. If not, then this very limit in their understanding of such agents and predicaments undercuts somewhat the attribution of such a sophisticated higher-order state as the desire to deceive.

[9] This commentary on Premack's chimpanzees grew out of discussion at the Dahlem conference on animal intellience with Sue Savage-Rumbaugh, whose chimps, Austin and Sherman, themselves exhibit apparently communicative behaviour (Savage-Rumbaugh 1978) that cries out for analysis and experimentation via the Sherlock Holmes methods.

higher nonhuman animals) exhibit mixed and confusing symptoms of higher-order intentionality. They will pass some higher-order tests and fail others; they will in some regards reveal themselves to be alert to third-order sophistications, while disappointing us with their failure to grasp some apparently even simpler second-order points. No crisp, 'rigorous' set of intentional hypotheses of any order will be clearly confirmed. The reason I am willing to make this prediction is not that I think I have special insight into vervet monkeys or other species, but just that I have noted, as anyone can, that much the same is true of us human beings. We are not ourselves unproblematic exemplars of third-, fourth-, or fifth-order intentional systems. And we have the tremendous advantage of being voluble language users, beings that can be plunked down at a desk and given lengthy questionnaires to answer, and the like. Our very capacity to engage in linguistic interactions of this sort seriously distorts our profile as intentional systems, by producing illusions of much more definition in our operative systems of mental representation than we actually have (Dennett 1978*a*, Chaps. 3, 16; Dennett 1981*b*). I expect the results of the effort at intentional interpretation of monkeys, like the results of intentional interpretation of small children, to be riddled with the sorts of gaps and foggy places that are inevitable in the interpretation of systems that are, after all, only imperfectly rational (see Dennett 1981*a*, *c*; 1982).

Still, the results, for all their gaps and vagueness, will be valuable. How and why? The intentional stance profile or characterization of an animal—or for that matter, an inanimate system—can be viewed as what engineers would call a set of specs— specifications for a device with a certain overall information-processing *competence*. An intentional system profile says, roughly, *what information* must be receivable, usable, rememberable, transmittable by the system. It alludes to the ways in which things in the surrounding world must be represented—but only in terms of distinctions drawn or drawable, discriminations makable—and not at all in terms of the actual machinery for doing this work (cf. Johnston 1981 on 'task descriptions'.) These intentional specs, then, set a design task for the next sort of theorist, the representation-system designer.[10] This division of labour is already familiar in certain circles within artificial intelligence (AI): what I have called the intentional stance is what Newell (1982) calls 'the knowledge level'. Oddly enough, the very defects and gaps and surd places in the intentional profile of a less than ideally rational animal, far from creating problems for the system designer, point to the shortcuts and stopgaps Mother Nature has relied upon to design the biological system; they hence make the system designer's job easier.

[10] In the terms I develop in 'Three Kinds of Intentional Psychology (Dennett 1981*b*), intentional system theory specifies a semantic engine which must then be realized – mimicked, in approximation, by a syntactic engine designed by the subpersonal cognitive psychologist.

Suppose, for example, that we adopt the intentional stance towards bees, and note with wonder that they seem to *know* that dead bees are a hygiene problem in a hive; when a bee dies its sisters *recognize* that it has died, and, *believing* that dead bees are a health hazard, and *wanting*, rationally enough, to avoid health hazards, they *decide* they must remove the dead bee immediately. Thereupon, they do just that. Now if that fancy an intentional story were confirmed, the bee system designer would be faced with an enormously difficult job. Happily for the designer (if sadly for bee romantics), it turns out that a much lower order explanation suffices: dead bees secrete oleic acid; the smell of oleic acid turns on the 'remove-it' subroutine in the other bees; put a dab of oleic acid on a live, healthy bee, and it will be dragged, kicking and screaming, out of the hive (Gould and Gould 1982; Wilson *et al.* 1958).

Someone in artificial intelligence, learning that, might well say: 'Ah how familiar! I know just how to design systems that behave like that. Shortcuts like that are my stock in trade'. In fact there is an eerie resemblance between many of the discoveries of cognitive ethologists working with lower animals and the sorts of prowess mixed with stupidity one encounters in the typical products of AI. For instance, Roger Schank (1976) tells of a 'bug' in TALESPIN, a story-writing program written by James Meehan in Schank's lab at Yale, which produced the following story: 'Henry Ant was thirsty. He walked over to the river bank where his good friend Bill Bird was sitting. Henry slipped and fell in the river. Gravity drowned'. Why did 'gravity drown'? (!) Because the program used a usually reliable shortcut of treating gravity as an unmentioned *agent* that is always around pulling things down, and since gravity (unlike Henry in the tale) had no friends (!), there was no one to pull it to safety when it was in the river pulling Henry down.

Several years ago, in 'Why Not the Whole Iguana?' (Dennett 1978*b*) I suggested that people in AI could make better progress by switching from the modelling of human microcompetences (playing chess, answering questions about baseball, writing nursery stories, etc.) to the whole competences of much simpler animals. At the time I suggested it might be wise for people in AI just to *invent* imaginary simple creatures and solve the whole-mind problem for them. I am now tempted to think that truth is apt to be both more fruitful, and, surprisingly, more tractable, than fiction. I suspect that if some of the bee and spider people were to join forces with some of the AI people, it would be a mutually enriching partnership.

Section 2: Out of the Armchair*

In June of 1983, I had a brief introduction to ethological field work, observing Seyfarth and Cheney observing the vervet monkeys in Kenya. Once I got in the field and saw, first hand, some of the obstacles to performing the sorts of experiments I had recommended, I found some good news and some bad news. The bad news was that the Sherlock Holmes method, in its classical guise, had very limited applicability to the vervet monkeys—and by extrapolation, to other 'lower' animals. The good news was that by adopting the intentional stance one can generate some plausible and indirectly testable hypotheses about why this should be so, and thereby learn something important about the nature of the selection pressures that probably have shaped the vervets' communication systems.

A vocalization that Seyfarth and Cheney were studying during my visit had been dubbed the Moving Into the Open (or MIO) grunt. Shortly before a monkey in a bush moves out into the open, it often gives a MIO grunt. Other monkeys in the bush will often repeat it—spectrographic analysis has not (yet) revealed a clear mark of difference between the initial grunt and this response. If no such echo is made, the original grunter will often stay in the bush for 5 or 10 minutes and then repeat the MIO. Often, when the MIO is echoed by one or more other monkeys, the original grunter will thereupon move cautiously into the open.

What does the MIO grunt mean? We listed the possible translations to see which we could eliminate or support on the basis of evidence already at hand. I started with what seemed to be the most straightforward and obvious possibility:

'I'm going.'

'I read you. You're going.'

What, however, would be the use of saying this? Vervets are in fact a taciturn lot, who keep silent most of the time, and are not given to anything that looks like passing the time of day by making obvious remarks. Then could it be a request for permission to leave?

'May I go, please?'

'Yes, you have my permission to go.'

This hypothesis could be knocked out if higher ranking vervets ever originated the MIO in the presence of their subordinates. In fact, higher ranking vervets do tend to move into the open first, so it doesn't seem that MIO is a request for permission. Could it be a command, then?

'Follow me!'

'Aye, Aye, Cap'n.'

Not very plausible, Cheney thought.

* Excerpted from Dennett (in press, *a*, *b*).

Vhy waste words with such an order when it would seem to *go without saying* in ervet society that low-ranking animals follow the lead of their superiors? For nstance, you would think that there would be a vocalization meaning 'May I?' to ɔe said by a monkey when approaching a dominant in hopes of grooming it. And ou'd expect there to be two responses: 'You may' and 'You may not', but there is ıo sign of any such vocalization. Apparently, such interchanges would not be useful enough to be worth the effort. There are gestures and facial expressions which may serve this purpose, but no audible signals.

Perhaps, Cheney mused, the MIO grunt served simply to acknowledge and share the fear:

'I'm really scared.'

'Yes. Me too.'

Another interesting possibility was that the grunt helped with co-ordination of the group's movements:

'Ready for me to go?'

'Ready whenever you are.'

A monkey that gives the echo is apt to be the next to leave. Or perhaps even better:

'Coast clear?'

'Coast is clear. We're covering you.'

The behaviour so far observed is compatible with this reading, which would give the MIO grunt a robust purpose, orienting the monkeys to a task of co-operative vigilance. The responding monkeys do watch the leave-taker and look in the right directions to be keeping an eye out. 'Suppose then, that this is our best candidate hypothesis,' I said. 'Can we think of anything to look for that would particularly shed light on it?' Among males, competition overshadows co-operation more than among females. Would a male bother giving the MIO if its only company in a bush was another male? Seyfarth had a better idea: suppose a male originated the MIO grunt; would a rival male be devious enough to give a dangerously misleading MIO response when he saw that the originator was about to step into trouble? The likelihood of ever getting any good evidence of this is minuscule, for you would have to observe a case in which the Originator didn't see and Responder did see a nearby predator *and* Responder saw that Originator didn't see the predator. (Otherwise Responder would just waste his credibility and incur the wrath and mistrust of Originator for no gain.) Such a coincidence of conditions must be extremely rare. This was an ideal opportunity, it seemed, for a Sherlock Holmes ploy.

Seyfarth suggested that perhaps we could spring a trap with something like a stuffed python that we could very slyly and surreptitiously reveal to just one of two males who seemed about to venture out of a bush. The technical problems would clearly be nasty, and at best it would be a long

shot, but with luck we might just manage to lure a liar into our trap. However, on further reflection, the technical problems looked virtually insurmountable. How would we establish that the 'liar' had actually seen (and been taken in by) the 'predator', and wasn't just innocently and sincerely reporting that the coast was clear? I found myself tempted (as often before in our discussions) to indulge in a fantasy: 'If only I were small enough to dress up in a vervet suit, or if only we could introduce a trained vervet, or a robot or puppet vervet who could . . . ' and slowly it dawned on me that this recurring escape from reality had a point: there is really no substitute, in the radical translation business, for going in and talking with the natives. You can test more hypotheses in half an hour of attempted chitchat than you can in a month of observation and unobtrusive manipulation. However, to take advantage of this you have to become obtrusive; you—or your puppet—have to enter into communicative encounters with the natives, if only in order to go around pointing to things and asking 'Gavagai?' in an attempt to figure out what 'Gavagai' means. Similarly, in your typical mystery story caper, some crucial part of the setting up of the 'Sherlock Holmes method' trap is—*must be*—accomplished by imparting some (mis)information verbally. Manoeuvring your subjects into the right frame of mind—and knowing you've succeeded—without the luxurious efficiency of words can prove to be arduous at best, and often next to impossible.

In particular, it is often next to impossible in the field to establish that particular monkeys have been shielded from a particular bit of information. Since many of the theoretically most interesting hypotheses depend on just such circumstances, it is often very tempting to think of moving the monkeys into a lab, where a monkey can by physically removed from the group and given opportunities to acquire information that the others don't have and that the test monkey knows they don't have. Just such experiments are being done, by Seyfarth and Cheney with a group of captive vervets in California, and by other researchers with chimpanzees. The early results are tantalizing, but equivocal (of course), and perhaps the lab environment, with its isolation booths, will be just the tool we need to open up the monkeys' minds, but my hunch is that being isolated in that way is such an unusual predicament for vervet monkeys that they will prove to be unprepared by evolution to take advantage of it.

The most important thing I think I learned from actually watching the vervets is that they live in a world in which secrets are virtually impossible. Unlike orangutans, who are solitary, and get together only to mate and when mothers are rearing offspring; unlike chimps, who have a fluid social organization in which individuals come and go, seeing each other fairly often, but also venturing out on their own a large proportion of the time, vervets live in the open in close proximity to the other members of their

groups, and have no solitary projects of any scope. Thus, it is a rare occasion indeed when one vervet is in a position to learn something that it alone knows and knows that it alone knows. (The knowledge of the others' ignorance, and of the possibility of maintaining it, is critical. Even when one monkey is the first to see a predator or a rival group, and knows it, it is almost never in a position to be sure the others won't very soon make the same discovery.) However, without such occasions in abundance, there is little to impart to others. Moreover, without frequent opportunities to recognize that one knows something that the others don't know, devious reasons for or against imparting information cannot even exist—let alone be recognized and acted upon. I can think of no way of describing this critical simplicity in the Umwelt of the vervets, this missing ingredient, that does not avail itself explicitly or implicitly of higher-order intentional idioms.

In sum, the vervets couldn't really make use of most of the features of a human language, for their world—or you might even say their lifestyle—is too simple. Their communicative needs are few, but intense, and their communicative opportunities are limited. Like honeymooners who have not been out of each other's sight for days, they find themselves with not much to say to each other (or to decide to withhold). If they couldn't make use of a fancy, human-like language, we can be quite sure that evolution hasn't provided them with one. Of course, if evolution provided them with an elaborate language in which to communicate, the language itself would radically change their world, and permit them to create and pass secrets as profusely as we do. And then they could go on to use their language, as we use ours, in hundreds of diverting and marginally 'useful' ways. Without the original information-gradients needed to prime the evolutionary pump, however, such a language couldn't get established.

So we can be quite sure that the MIO grunt, for instance, is not crisply and properly translated by any familiar human interchange. It can't be a (pure, perfect) command or request or question or exclamation because it isn't part of a system that is elaborate enough to make room for such sophisticated distinctions. When you say 'Wanna go for a walk?' to your dog and he jumps up with a lively bark and expectant wag of the tail, this is not really a question and answer. There are only a few ways of 'replying' that are available to the dog. It can't do anything tantamount to saying 'I'd rather wait till sundown', or 'Not if you're going to cross the highway', or even 'No thanks'. Your utterance is a question in English, but a sort of melted-together mixture of question, command, exclamation, and mere harbinger (you've made some of those going-out-noises again) to your dog (Bennett 1976, 1983). The vervets' MIO grunt is no doubt a similar mixture, but while that means we shouldn't get our hopes too high about learning Vervetese and finding out all about monkey life by having conversations with the vervets, it doesn't at all rule out the utility of these somewhat fanciful translation hypotheses as ways

of interpreting—and uncovering—the actual informational roles or functions of these vocalizations. When you think of the MIO as 'Coast clear?' your attention is directed to a variety of testable hypotheses about further relationships and dependencies that ought to be discoverable if that is what MIO means—or even just 'sort of' means.