Behaviorism, Spring 1987, Vol. 15, Number 1

CONTRASTING APPROACHES TO THE LEGITIMATION OF INTENTIONAL LANGUAGE WITHIN COMPARATIVE PSYCHOLOGY

Cecilia M. Heyes
Department of Experimental Psychology
University of Cambridge

ABSTRACT: Dennett, a philosopher, and Griffin, an ethologist, have recently presented influential arguments promoting the extended use of intentional language by students of animal behavior. This essay seeks to elucidate and to contrast the claims made by each of these authors, and to evaluate their proposals primarily from the perspective of a practicing comparative psychologist or ethologist. While Griffin regards intentional terms as explanatory, Dennett assigns them a descriptive function; the issue of animal consciousness is central to Griffin's program and only tangentially related to Dennett's. The philosopher's arguments are founded upon a more coherent metaphysics, but Dennett neglects to substantiate his claim that animal competences can be most readily modelled by artificial intelligence (AI) specialists when they are described in intentional terms. Both authors assume that some examples of animal behavior should not be given an intentional characterization, but neither provides adequate guidelines for the identification of cases belonging to this negative set.

Throughout most of this century comparative psychologists have tried to avoid using intentional language to describe animal behavior; experiencing a certain guilt when references to an animal's beliefs, wishes or expectations have crept into their texts or dialogue. The historical origin of this tradition, commonly associated with Lloyd Morgan's Canon, is well known. In the last decade Griffin has launched an influential campaign challenging the tradition and seeking to legitimize the use of intentional language by students of animal behavior. Dennett has recently written in support of the same goal, but on very different grounds. At the core of Griffin's thesis is a relatively simple ontological claim concerning the existence of mental states in non-human species, while Dennett claims that intentional language is a useful descriptive tool for comparative psychologists, regardless of whether animals have mental states in any conventional sense. The contrast between these two programs is often overlooked by members of their target audience, cognitive ethologists, who tend to regard Dennett's arguments as the source of a philosophical imprimatur for Griffin's views. With or without this "support," several recent commentators, for example Silverman (1983) and Burghardt (1985), have concluded that the case for animal consciousness is now strong enough to persuade comparative psychologists to make cautious but unashamed use of intentional terms. In this paper I query that conclusion by exploring the contrast between Griffin's and Dennett's programs, and by arguing that while the latter is more conceptually coherent, neither provide the practicing psychologist or

I am grateful to Kathleen Akins, Donald Campbell, Daniel Dennett and David Oakley for very helpful discussion. This paper was presented at the Twelfth Annual Meeting of the Society for Philosophy and Psychology at Johns Hopkins University, Baltimore, June 1986, and several of its themes are discussed in Heyes (1987).

ethologist with adequate guidelines for distinguishing between examples of animal behavior which do, and do not, warrant intentional description by their lights.

CONSCIOUS INTENTION

I will consider Griffin's program first. His most recent book, *Animal Thinking*, begins:

> What is it like to be an animal? What do monkeys, dolphins, crows, sunfishes, bees and ants think about? Or do nonhuman animals experience any thoughts and subjective feelings at all? People have always been fascinated by the question of animal consciousness, because both pets and wild animals arouse our admiration and curiosity. They tempt us to put ourselves into their skins and imagine what their lives are like. But is this possible? Have students of animal behavior learned enough to constrain our speculations constructively, about the thoughts and feelings of other species? (Griffin, 1984, p. 1)

This fragment of text introduces most of the important elements of Griffin's program: First, it addresses two questions: Are animals conscious, and if so, What are they conscious of? In attempting to answer, Griffin does not confine his attention to a particular kind of conscious experience, he is concerned with both "thoughts and subjective feelings" (inferential and perceptual beliefs), and often neglects to distinguish between them. Second, Griffin regards species from almost any phylum as serious candidates for the attribution of consciousness, in a way that might shock any interested philosopher or scientist of the last century, with the possible exception of Romanes (1883). Griffin allies himself with the layman, the affectionate pet owner and curious huntsman. He calls for a re-examination of the issue of animal consciousness by biologists and psychologists on the grounds that their denial or disinterest is in conflict with the layman's intuitive conviction. Finally, of the many kinds of conscious experience about which Griffin is willing to speculate, he proposes that "the ability to think about objects and events, whether or not they are part of the immediate situation," should be the focus of empirical investigation.

Griffin is well known for his elegant studies of echolocation in bats, but his contemporary research largely consists of accumulating instances of behavior which can, in his opinion, be plausibly explained as the outcome of conscious thinking. For example (one chosen at random from a collection which is very impressive in the diversity of species and functional categories of behavior that it samples), Griffin regards the protective case building behavior of caddis fly larvae as "suggestive of intentional action." These larvae construct their hard outer covering from pieces of leaf which they cut into rectangles and bind together in a complex staggered arrangement using silk secreted from their own bodies. If you cut away the front of a completed case, giving it a continuous smooth edge, the larva will clip leaves into shapes which differ from those used in the original process of construction, and glue them into place restoring the staggered arrangement (Hansell, 1972). It is this flexibility which Griffin finds impressive.

It is only with difficulty that one can decipher the exact nature and novelty of Griffin's claims, and there seems to be a tendency among his commentators to attempt to "normalize" Griffin's views. For example, Oakley (1985) suggests that Griffins's definition of consciousness is similar to his own, "that consciousness can be defined, irrespective of its subjective accompaniments, as synonymous with the use of mapping or representational strategies of information processing" (p. 148). If this were the case then Griffin's claim that

certain vertebrates are conscious would be entirely compatible with current opinion in animal cognitive psychology, but I doubt the accuracy of Oakley's interpretation. Despite expressing a commitment to "emergent" or "reductive materialism," an appeal to layman's intuition is dominant among Griffin's attempts to express what he means by "consciousness." He cites dictionary definitions that are based on such an understanding, and endorses Armstrong's (1981) suggestion that we can appreciate the meaning of "consciousness" by contrasting our states of mind in relation to the actions involved in driving a car when we are learning, and later when the skill has been thoroughly acquired. This strongly suggests that Griffin is conventional among cognitive psychologists in referring to "the use of mapping and representational strategies of information processing" as "thought," and reserving the terms "conscious" and "conscious thinking" to indicate that the legendary internal light is *on*.

In a further attempt to normalize Griffin's views, it might be suggested that his frequent failure to discriminate among sub-categories of conscious experience makes his claims seem more controversial than they really are. If he were suggesting that species from many phyla experience pleasure and pain, but that only certain birds and mammals are rational or have conscious inferential beliefs, then his position would be similar to that of Lorenz (1963), or Paul Churchland (1979), who distinguishes between feeling and judging that one feels, and permits that almost any animal is capable of the former. It is clear that Griffin's claims are not compatibly modest when one considers his response to Humphrey's (1978) thesis that consciousness may have arisen in human evolution when social groups reached a size and degree of inter-dependence that made it important for each member to understand his companions' moods and thoughts in order to predict their behavior. Agreeing that this describes an important function of consciousness, Griffin contests Humphrey's conclusion that these capacities are confined to certain primates, and argues that, since they would also be useful for the social insects, we can also expect to find them there.

The informality of Griffin's prose should not be allowed to obscure the novelty of his position. He rejects two "rules" for the attribution of consciousness to other animals which have been respected by most students of animal behavior since the turn of the century. The first is the "similar brains" argument. Griffin challenges this specific version of the argument from analogy on the grounds that too little is known about the relationship between neurophysiological and mental events to eliminate the possibility that a wide variety of different neural mechanisms could be associated with the same conscious experience. He could be said to adhere to a "similar nervous tissue," but not a "similar brains" rule, and within this scheme finds himself able to grant that ants, but not machines, may be capable of forming inferential beliefs.

A second break with tradition is marked by Griffin's claim that instinctive behavior may have conscious accompaniments. He attributes the denial of such an association to selection of inappropriate introspections from which to reason by analogy. To be sure, we are not always aware of blinking but why, asks Griffin, should we treat blinking as representative of human instinctive acts? Some of the complex acts of which we are fully aware may be instinctive even though we have failed to identify them as such because of technical difficulties involved in teasing apart nature and nurture. Taking a different tack, Griffin tries to persuade his audience that even if complex human behaviors are not instinctive, those of other species might well be, and concludes that the denial of consciously sensed instinctive behaviors might be based on "unwisely anthropocentric" reasoning.

The goal of Griffin's program is not to *prove* that animals are conscious, but to convince

students of animal behavior that the burden of proof should fall upon those who deny it; that the existence of animal consciousness should be the null hypothesis. Of course, this is an attempt to reverse a bias which many believe to have been established by Lloyd Morgan's Canon. In fact, Lloyd Morgan, and even Thorndike in his early writings, expressed opinions on animal consciousness not at all dissimilar to Griffin's (Walker, 1983). The latter allies himself with Darwin and Romanes in their use of anecdotal evidence, consideration of subjective feelings as well as thoughts, and belief that instinctive acts could have conscious accompaniments. Indeed, in his more speculative moments Griffin is their contemporary representative, but when he gets down to identifying what he regards as the most interesting and persuasive evidence for animal consciousness he homes in on the same sort of behavior as Lloyd Morgan, that is, the non-habitual utilization of previous experience. The fate of the line of enquiry which Griffin seeks to resurrect after some seventy years was sealed, not by Lloyd Morgan, but by the strict empiricism of Watson and his followers. Is the climate now ripe for it to flourish?

Griffin thinks that it is, while acknowledging his share in the predicament of underdetermination which Nagel (1974) described so plausibly for animal consciousness. Griffin writes:

> I have suggested in earlier chapters that versatile coping with new challenges provides suggestive evidence of conscious thinking, but in every case a behaviorist can argue that a completely unconscious organism could behave in the same effectively adaptable fashion. (1984, p. 208)

Griffin's grounds for preferring the attribution of consciousness under these conditions appear to be threefold: (i) It is highly plausible or intuitively true; (ii) "extensive and mutually reinforcing evidence" is accumulating in its favor, and (iii) other hypotheses, like those concerning evolutionary lines of descent, are regarded as both useful and respectable even though they cannot be definitively confirmed or falsified by current methods.

Concerning the first of these, the problem is that Griffin's thesis is *too* plausible. In their classic work Heider and Simmel (1944) demonstrated that intentions and awareness are often attributed to filmed neutral objects, such as triangles and circles, exhibiting patterned movement. Similar attributions have been observed with respect to natural disasters, sporting events, gambling, and in a host of other circumstances where the objects of these attributions could not possibly be conscious within any accepted sense of the term (Felson & Gmelch, 1979; Lewis, 1963). These data may well lend weight to some variant on the thesis that the evolutionary origin of human consciousness lies in the necessity to predict environmental contingencies (e.g. Humphrey, 1982), but it should make us cautious about projecting our experience onto members of other species. Griffin's empirical program is utterly impractical because it does not provide the scientist with principles by which to exercise that caution. He gives an example of what he regards as "unwisely anthropocentric" reasoning, but no clue about the general criteria that could be applied in making other, similar judgements; to use Jaynes (1976) terminology, no indication of when the human case should be the metaphier and when the metaphrand.

Griffin's second reason for preferring the attribution of consciousness could be regarded as a claim that his theory is "robust" (Campbell & Fiske, 1959; Wimsatt, 1981); that the accumulating evidence has been derived by different methods, based on semi-independent sets of assumptions. In fact, all the data which Griffin counts as evidence does so by virtue of the single assumption that an animal is conscious at a given time to the extent that its behavior at that time resembles the conscious activity of a person.

Consideration of Griffin's third reason suggests the most serious flaw in his approach. In view of the acrimonious debate over adaptationism, the choice of hypotheses concerning lines of descent as examples of those which are untestable and yet respectable was a poor one. However, in contrast with Griffin's, these are supported within a sound conceptual framework, that of evolution by natural selection. Under the unbrella of "emergent materialism" (the nature of which is left unspecified), Griffin's commonsense approach attributes causal force to conscious experience; not to the associated brain states or functional states, but to the gauzy experience itself. He had ample opportunity to deny this when, in an exchange of "Letters to the Editor," Harnad asked Griffin: "As to studying consciousness, how is one to study an entity that (unlike say, quarks) is causally superfluous in any theory." Instead, Griffin took this as an allegation that he denies that consciousness has physical *causes*, and reaffirmed his belief that it also has physical *effects*:

> Some of our conscious thinking (and perhaps that of other species) entails comparisons of, and choices among, the likely outcomes of various actions. Such thinking may well be not more causally superfluous than contractions of cardiac muscle are superfluous for the transport of oxygen from lung to brain. Of course thoughts, hearts and hemoglobin have all resulted from antecedent causes. But why should this be selectively taken as a reason to disparage the significance of conscious thinking?

INTENTIONALITY AND UTILITY

If you were to ask Griffin: "Why should we flout scientific tradition and say of an animal that it has a certain belief?" he would probably reply "Because it is *true*." Dennett on the other hand would tell you that, under some conditions, it is *useful*.

Dennett (1983) addresses himself to the same audience as Griffin, in his words, "the new ethologists [who] having cast off the straightjacket of behaviorism and kicked off its weighted overshoes, are looking about somewhat insecurely for something presentable to wear."Griffin cites Dennett favorably, and cognitive ethologists are apt to bracket the two authors together, but in fact they are arguing very different points. Griffin is suggesting that when animal behavior can be plausibly explained in intentional or folk psychological terms then the animal is conscious, and that this consciousness *explains* the behavior. Dennett on the other hand is suggesting that when intentional terms are employed in a certain way to produce an account of a system which accurately predicts its behavior, then that account will *ultimately* be useful in explaining the behavior. According to Dennett, an intentional account is a useful *description*, not an explanation, and its success in predicting behavior tells us nothing about whether or not the system is conscious. Perhaps cognitive ethologists can be forgiven for having misunderstood this position since it appears that some of Dennett's closest colleagues, including Fodor and Dretske, have also failed to bite the bullet; to realize that, for Dennett, there is only one kind of intentionality not two, an "as if" and a "genuine" variety (Dennett, 1986).

Despite the conceptual gulf that separates their claims, the empirical programs recommended by Griffin and Dennett for cognitive ethology are similar in many respects. Dennett also stresses the importance of accumulating anecdotal reports of phyletically diverse animals' behavior under free-living conditions, and does not regard this as a means of *disproving* hypotheses which might be proffered by a hypothetical behaviorist. However, while the documentation of a range of behaviors that can be explained in intentional terms is regarded by Griffin as *sufficient* evidence in support of his claim, the same data are only a

starting point for the research that Dennett recommends. He is not content to eyeball a piece of behavior and decide casually whether *any* intentional statement can describe it; he wants cognitive ethologists to generate hypotheses concerning *which* intentional states, and *what order* of intentionality (in a Gricean sense), are appropriate. It is very important to Dennett that the hypotheses should be empirically testable. He stresses that,

> The question is empirical. The tactic of adopting the intentional stance is not a matter of *replacing* empirical investigations with aprioristic ('armchair') investigation, 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. (Dennett, 1986, p. 347)

There are several things worth noting about Dennett's recommendations concerning the use of empirical data to choose among intentional accounts. First, he does not imagine that any of the relevant types of hypotheses can be readily swept away by crucial experiments; the behaviorist and the romantic ethologist alike could "save the data" by postulating additional contextual variables or complex beliefs. In consequence, Dennett suggests that we augment the selection procedure with two meta-empirical principles: the rationality assumption, and what I shall refer to as a principle of limited parsimony. That is, when more than one hypothesis is consistent with the data, we should prefer those hypotheses which invoke lower orders of intentionality; and the rationality assumption qualifies this by specifying that the favored hypotheses should portray the animal as rational. However, we are not given any clear idea of what it is to be rational.

What then is Dennett's purpose in making these recommendations? He is not, at least in any straightforward way, advocating that we choose among hypotheses according to their likely correspondence with events inside animals' heads. For him, the function of the intentional stance within cognitive ethology is to provide descriptions of real life animal competences that will facilitate the tasks of the AI and information processing specialists. It is incumbent upon them, and not the "mentalist," to explain behavior in terms of (to borrow a phrase from Patricia Churchland, 1983) "causally active internal states." Exactly *how* intentional accounts will help AI specialists; or rather, the *unique* advantages of the intentional stance, are not too clear. Dennett writes:

> 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. (Dennett, 1986, p. 349)

I would like to consider two kinds of claim that Dennett could be making here. The first is weaker. It says that the intentional stance is useful because it is so familiar to us from everyday life, and yet relatively novel in the context of ethology. Therefore, it gives the cognitive ethologist a "handle" on the problem, it heightens his motivation by making his research exciting. (In this case it may be to the advantage of Dennett's cause for him to allow his position to be mistaken, so that cognitive ethologists believe that they are *really* discovering whether animals have mental states.) The result of all this is that cognitive ethologists churn out a lot of data, flesh out the descriptions of competences that are already dimly known, and discover new ones. The intentional character of these descriptions is not important in itself. "Specs" which describe animals in mentalistic terms and portray them as perfectly rational are no more useful to the information processing theorist than other

functional descriptions. On this reading, it does not matter whether Dennett's rationality assumption is incorrect (as, for example, Elster suggested in 1983), or whether his principle of limited parsimony is without theoretical foundation. The function of these criteria is merely to project upon cognitive ethologists the illusion that they are real scientists, involved in the explanation as well as the description of behavior, and thereby to increase their output.

The second and stronger version of Dennett's claim is more interesting, and more palatable to the intellectual pride of cognitive ethologists, but it seems to reveal some omissions and inconsistencies in his thesis. Dennett may be suggesting that, in addition to providing a "concert party for the troops," disciplined application of the intentional stance in cognitive ethology will yield descriptions of animal behavior that are especially useful to the information processing theorist. In telling the latter which behaviors can be construed as rational and, among these, which mental states and relations among those states are appropriate (in accordance with a principle of limited parsimony), the products of the intentional stance have unique advantages over other functional accounts. The following passage suggests that Dennett is, in fact, committed to this view.

> The intentional stance, however, provides just the right interface between specialities: a 'black box' characterization of behavioral and cognitive competences observable in the field, but couched in language that (ideally) heavily constrains the design of machinery to put in the black box. (Dennett, 1986, p. 350)

"Ideally" yes, but *how*?

At this point it should be clear that I think we have been short-changed on the distinction between the intentional stance and what Dennett (1987) has called the "design" stance; that Dennett has not told us enough about how intentional descriptions differ from other functional descriptions, and what it is about the difference that makes the former so useful. Teleonomic language (Pittendrigh, 1958), that which accounts for behavior in terms of goals and their means of attainment, is the first port of call for any student of animal behavior who has become exasperated by the complexity of appropriate causal/mechanistic/"physical" stance description. Contrasting intentional accounts only with those of a hypothetical behaviorist, Dennett ignores this important middle ground, such that we are uncertain as to whether, in the context of cognitive ethology, these functional accounts are equivalent or inferior to intentional accounts in their potential to describe what is represented, and if the latter, precisely what is that they lack.

Gleaning other sources (Dennett, 1978, 1981), it would appear that the intentional stance is distinctive in two respects: It assumes that the system is optimally designed, and employs a vocabulary with an age-old pedigree in the explanation of human action. If this is so, then I remain a little mystified and a little worried. How can an *optimal* characterization of a system *constrain* the AI specialist's task? And, isn't there a danger that use of belief-desire vocabulary will, to borrow another of Dennett's insights, "pump" some rather undesirable intuitions (Dennett, 1985); put us in a mental set which prevents us from remembering some important differences between our animal subjects and ourselves. For example, even when using the slightly strained language of goals and methods, it is difficult to keep in mind that the sensory systems of some animals are very different from our own. One consequence of forgetting this is the received wisdom in comparative psychology that "observational" learning must occur through the visual modality, even when the rodent learners are nearly blind.

I feel slightly apologetic about referring to the old literature on mental set formation, but I do believe that it can be interpreted as an interesting forerunner to the current literature on

human irrationality (e.g. Kahneman et al., 1982). These studies in the 1930s and '40s showed that people are slow to realize that a familiar object can be used in a novel way. Given an array of objects and asked to construct a pendulum, subjects would take forever to realize that a pair of ordinary pliers can substitute for a carpenter's plumb bob. However, when the pliers were labelled "HEAVY OBJECT," subjects had the pendulum set up in no time (Maier, 1940, 1945). The cumbersome vocabulary which cognitive ethologists use when they are trying to avoid words like "believe" and "desire" may be valuable; it may act like a label stuck to each animal saying something like "POOR VISION, GOOD SMELL, NO HANDS," without calling for painful interactions between glue and fur.

My worries focus on the extension of the intentional stance to animal, but there are several authors, such as Rosenberg (1980) and the Churchlands (1979, 1986) who regard the folk psychology from which the intentional stance comes as a very unsatisfactory way of accounting even for human behavior. They regard it as a theory of persons which is deeply entrenched in all human societies but which lacks explanatory and predictive power in everyday life, let alone in science. I find this position curious given the naturalistic commitments of these authors (to paraphrase an oft repeated battle-cry, it seems unlikely that natural selection would leave us with an habitual method of construing conspecific action that regularly deceives us), but it clearly presents a significant challenge to Dennett's approach. In brief, Rosenberg (1980) argues that the terms of any potential law covering intentional explanations (beliefs, desires etc.), cannot be defined independently of that law. For him, the most significant consequence of this is not the preclusion of non-circular tests of potential laws, but the isolation of those laws from generalizations of other kinds which might explain them, or be explained by them. The other kinds of generalization that Rosenberg considers concern physiological states of the brain, but his argument equally implies that intentional terms are isolated from information theory. The upshot of the argument is that intentional explanation must be *replaced*, it cannot be used as a pathway to, or be encompassed by, accounts that are consistent with a materialist ontology.

I have suggested that Dennett has a lot more explaining to do if he is going to convince us that the intentional stance has more than a recreational function in ethology; that ethologists' potential to contribute to AI in this way is great enough to make it worthwhile risking tunnel-vision anthromorphism. However, even if we accept Dennett's I. O. U. on how intentional accounts can be unpacked into AI models, we would still have problems in deciding *when* to apply the intentional stance. Dennett certainly does not recommend that it be used for all behavior in all animals. Indeed, he portrays it as a curious device—an invaluable last resort. That is, Dennett provides two guidelines for when to apply the intentional stance and they do not always converge on the same behaviors. The first is "realistic;" it concerns characteristics of the animal under investigation. Dennett tells the ethologist to use the intentional stance whenever there is reason to believe that the system in question is optimally designed and, for him, this is the case whenever there is evidence to suggest that it has been formed by selection processes. The other guideline is essentially "pragmatic;" it relates to the scientist's explanatory and descriptive resources. In this vein, Dennett tells us that as a source of explanations of intelligence the intentional stance is inherently undesirable because it *presupposes* intelligence. Therefore, since intentional characterizations of a system will eventually have to be converted into some other kind, ethologists should save themselves some work and use other, truly explanatory, devices whenever possible.

Problems arise when a behavior qualifies for intentional description on one of these counts but not on the other. Many simple innate behaviors are clearly the products of a selection process

(natural selection), but they can also be accounted for neatly in mechanistic terms. An example would be the behavior of bees that remove dead conspecifics from the hive. Should we say that they *recognize* that their sister is dead, *believe* that her corpse is a health hazard, and *desire* to avoid their own destruction; or simply that they are responding to oleic acid secreted by the corpse. In this particular case Dennett approves of the latter characterization. However, he also insists that it is inappropriate to give intentional descriptions of complex laboratory trained behavior, even though this behavior is both a product of selection processes and very often difficult to explain in traditional learning theoretic terms, let alone from the physical stance. Why? These are isolated judgements on Dennett's part. He seems to have imported several assumptions or inferential rules-of-thumb from the early comparative psychologist's realist approach to the issue of animal consciousness and set them down intact within his relatively instrumentalist framework. Just as Griffin failed to indicate under what conditions the analogy from introspection is valid, Dennett neglects to synthesize his two guidelines into a workable scheme for deciding when ethologists should throw in the towel on their attempts at mechanism and adopt the intentional stance.

In summary, Dennett's contribution provides a much needed, and very entertaining, lesson on the subtleties of scientific realism for cognitive ethologists like myself, and, pending a more explicit account of how intentional descriptions are to be unpacked into information processing models, it may predict very valuable cooperation between ethology and AI, two much divided disciplines. However, in its present form it does not achieve its purpose of alleviating our intellectual insecurity. Indeed, it may aggrevate the condition. Having been assured that intentional explanations are legitimate, we are uncertain as to why this is the case. Is it just that they will keep us working, or do they have a more specific function? If, ignoring our suspicion of a snub, we start generating intentional descriptions for a given behavior, we soon discover that they breed like toads and are beset by a plague. Remembering that the search for reductionist explanations is no longer interesting, *and* that lower order intentions are to be preferred to higher ones, we are in a quandary about where to begin culling the beasts. As in our confusion we reach out for the axe of the rationality assumption, we find that it was made of rubber all along, and may only have been put there to give us the illusion of strength.

REFERENCES

Burghardt, G. M. (1985). Animal awareness: Current perceptions and historical perspective. *American Psychologist, 40*, 905-919.

Campbell, D. T. & Fiske, W. (1959). Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin, 56*, 81-105.

Churchland, Patricia S. (1983). A frog at the bottom of a mug. Commentary on Dennett. *Behavioral and Brain Sciences, 6*, 343-390.

Churchland, Patricia S. (1986). *Neurophilosophy: Toward a unified theory of the mind-brain*. Cambridge, MA: Bradford Books, MIT Press.

Churchland, Paul M. (1979). *Scientific realism and the plasticity of mind*. Cambridge, MA: Cambridge University Press.

Dennett, D. C. (1978). *Brainstorms*. Cambridge, MA: Bradford Books, MIT Press.

Dennett, D. C. (1981). True believers: The intentional strategy and why it works. In A. F. Heath (Ed.), *Scientific explanation*. Oxford: Oxford University Press.

Dennett, D. C. (1983). Intentional systems in cognitive ethology: The "Panglossian paradigm" defended. *Behavioral and Brain Sciences, 6*, 343-390.

Dennett, D. C. (1985). *Elbow room: The varieties of free will worth wanting*. Cambridge, MA: Bradford Books, MIT Press.

Dennett, D. C. (1986). *Evolution, error and intentionality*. Duplicated manuscript, Tufts University, Center for Cognitive Studies, Medford, MA.

Elster, J. (1983). *Explaining technical change*. Cambridge: Cambridge University Press.

Felson, R. B. & Gmelch, G. (1979). Uncertainty and the use of magic. *Current Anthropology, 20*, 587-589.

Griffin, D. R. (1976). *The question of animal awareness*. New York: Rockefeller University Press.

Griffin, D. R. (1978). Prospects for a cognitive ethology. *Behavioral and Brain Sciences, 1*, 527-538.

Griffin, D. R. (Ed.). (1982). *Animal mind—human mind*, introduction. New York: Springer-Verlag.

Griffin, D. R. (1984). *Animal thinking*. Cambridge, MA: Harvard University Press.

Hansell, M. H. (1972). Case building behaviour of the caddis-fly larva *Lepidostoma hirtum*. *Journal of the Zoological Society (London)*, 167, 179-192.

Heider, R. & Simmel, M. (1944). An experimental study of apparent behavior. *American Journal of Psychology, 57*, 243-259.

Heyes, C. M. (1987). Cognisance of consciousness in the study of animal knowledge. In W. Callebaut & R. Pinxten (Eds.), *Evolutionary epistemology: a multiparadigm program*. Dordrecht: Reidel.

Humphrey, N. K. (1978). Nature's psychologists. *New Scientist, 78*, 900-903.

Humphrey, N. K. (1982). Consciousness: a just-so story. *New Scientist, 95*, 474-477.

Jaynes, J. (1976) *The origins of consciousness in the breakdown of the bicameral mind*. Boston: Houghton Mifflin.

Kahneman, D., Slovic, P. & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and their biases*. Cambridge: Cambridge University Press.

Lewis, L. (1963). Knowledge, danger, certainty, and the theory of magic. *American Journal of Sociology, 69*, 7-12.

Lorenz, K. (1963). Haben Tiere ein subjectives Erleben? *Jahr. Techn. Hochs. Munchen*. (English translation reprinted in *Studies in animal and human behavior* (Vol. II). Cambridge, MA: Harvard University Press.)

Maier, N. R. F. (1940). The behavior mechanisms concerned with problem solving. *Psychological Review, 47*, 43-58.

Maier, N. R. F. (1945). Reasoning in humans III. The mechanisms of equivalent stimuli and of reasoning. *Journal of Experimental Psychology, 35*, 349-360.

Nagel, T. (1974). What is it like to be a bat? *Philosophical Review, 83*, 435-451.

Oakley, D. A. (1985). Animal awareness, consciousness and self-image. In D. A. Oakley (Ed.), *Brain and mind*. London: Methuen.

Pittendrigh, C. S. (1958). Adaptation, natural selection, and behavior. In A. Roe & G. Simpson (Eds.), *Behavior and evolution*. New Haven: Yale University Press.

Romanes, G. J. (1883). *Mental evolution in animals*. London: Kegan Paul Trench.

Rosenberg, A. (1980). Obstacles to the nomological connection of reasons and actions. *Philosophy of Social Science, 10*, 79-91.

Silverman, P. S. (1983). Attributing mind to animals: The role of intuition. *Journal of Social & Biological Structure, 6*, 231-247.

Walker, S. (1983). *Animal thought*. London: Routledge & Kegan Paul.

Wimsatt, W. C. (1981). Robustness, reliability, and overdetermination. In M. Brewer & B. Collins (Eds.), *Scientific inquiry and the social sciences*. San Francisco: Jossey-Bass.