

ON BEING SIMPLE MINDED

Peter Carruthers

The question 'Do fishes think?' does not exist among our applications of language, *it is not raised*.—Wittgenstein

1. ON HAVING A MIND

How simple minded can you be? Many philosophers would answer: no more simple than a language-using human being. Many other philosophers, and most cognitive scientists, would allow that mammals, and perhaps birds, possess minds. But few have gone to the extreme of believing that very simple organisms, such as insects, can be genuinely minded.¹ This is the ground that the present paper proposes to occupy and defend. It will argue that ants and bees, in particular, possess minds. So it will be claiming that minds can be very simple indeed.

What does it take to be a minded organism? Davidson (1975) says: you need to be an interpreter of the speech and behavior of another minded organism. Only creatures that speak, and that both interpret and are subject to interpretation, count as genuinely thinking anything at all. McDowell (1994) says: you need to exist in a space of reasons. Only creatures capable of appreciating the normative force of a reason

for belief, or a reason for action, can count as *possessing* beliefs or engaging in intentional action. And Searle (1992) says: you need consciousness. Only creatures that have *conscious* beliefs and *conscious* desires can count as having beliefs or desires at all.

Such views seem to the present author to be ill-motivated. Granted, *humans* speak and interpret the speech of others. And granted, *humans* weigh up and evaluate reasons for belief and for action. And granted, too, humans engage in forms of thinking that have all of the hallmarks of consciousness. But there are no good reasons for insisting that these features of the human mind are necessary conditions of mindedness as such. Or so, at least, this paper will briefly argue now, and then take for granted as an assumption in what follows.

Common sense has little difficulty with the idea that there can be beliefs and desires that fail to meet these demanding conditions. This suggests, at least, that those conditions are not conceptually necessary ones. Most people feel pretty comfortable in ascribing simple beliefs and desires to non-language-using creatures. They will say, for example, that a particular ape acts as she does because she *believes* that the mound contains termites and *wants* to eat them. And our willingness to entertain such

thoughts seems unaffected by the extent to which we think that the animal in question can appreciate the normative implications of its own states of belief and desire. Moreover, most of us are now (post-Freud and the cognitive turn in cognitive science) entirely willing to countenance the existence of beliefs and desires that are not conscious ones.

It isn't only ordinary folk who think that beliefs and desires can exist in the absence of the stringent requirements laid down by some philosophers. Many cognitive scientists and comparative psychologists would agree. (Although sometimes, admittedly, the *language* of 'belief' and 'desire' gets omitted in deference to the sensibilities of some philosophical audiences.) There is now a rich and extensive body of literature on the cognitive states and processes of non-human animals (e.g., Walker, 1983; Gallistel, 1990; Gould and Gould, 1994). And this literature is replete with talk of information-bearing conceptualized states that guide planning and action-selection (beliefs), as well as states that set the ends planned for and that motivate action (desires).

True enough, it can often be difficult to say quite *what* an animal believes or desires. And many of us can, on reflection, rightly be made to feel uncomfortable when using a that-clause constructed out of our human concepts to describe the thoughts of an animal. If we say of an ape that she believes *that the mound contains termites*, for example, then we can easily be made to feel awkward about so doing. For how likely is it that the ape will have the concept *termite*? Does the ape distinguish between termites and ants, for instance, while also believing that both kinds belong to the same super-ordinate category (insects)? Does the ape really believe that termites are living things which excrete and reproduce? And so on.

These considerations give rise to an argument against the very possibility of non-linguistic thought, which was initially presented by Davidson (1975). The argument claims first, that beliefs and desires are content-bearing states whose contents must be expressible in a sentential complement (a that-clause). Then second, the argument points out that it must always be inappropriate to use sentential complements that embed *our* concepts when describing the thoughts of an animal (given the absence of linguistic behavior of the appropriate sorts). In which case (putting these two premises together) it follows that animals cannot be said to have thoughts at all.

The error in this argument lies in its assumption that thought contents *must* be specifiable by means of that-clauses, however. For this amounts to the imposition of a *co-thinking constraint* on genuine thoughthood. In order for another creature (whether human or animal) to be thinking a particular thought, it would have to be the case that someone *else* should also be capable of entertaining that very thought, in such a way that it can be formulated into a that-clause. But why should we believe this? For we know that there are many thoughts—e.g., some of Einstein's thoughts, or some of the thoughts of Chomsky—that *we* may be incapable of entertaining. And why should we nevertheless think that the real existence of those thoughts is contingent upon the capacity of someone *else* to co-think them? Perhaps Einstein had some thoughts so sophisticated that there is no one else who is capable of entertaining their content.

The common-sense position is that (in addition to being formulated and co-thought from the inside, through a that-clause) thoughts can equally well be characterized *from the outside*, by means of an indirect description. In the case of the ape dipping for termites, for example, most of us would,

on reflection, say something like this: we do not know how much the ape knows about termites, nor how exactly she conceptualizes them, but we do know that she believes *of* the termites in that mound that they are there, and we know that she wants to eat them. And on this matter common sense and cognitive science agree. Through careful experimentation scientists can map the boundaries of a creature's concepts, and can explore the extent of its knowledge of the things with which it deals (Bermúdez, 2003). These discoveries can then be used to provide an external characterization of the creature's beliefs and goals, even if the concepts in question are so alien to us that we couldn't co-think them with the creature in the content of a that-clause.

What does it take to be a minded organism, then? We should say instead: you need to possess a certain core cognitive architecture. Having a mind means being a subject of perceptual states, where those states are used to inform a set of belief states which guide behavior, and where the belief states in turn interact with a set of desire states in ways that depend upon their contents, to select from amongst an array of action schemata so as to determine the form of the behavior. This sort of belief/desire architecture is one that humans share. It is represented diagrammatically in Figure 1.

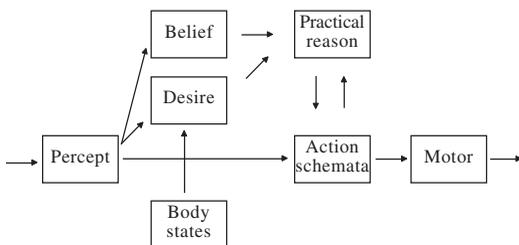


Figure 1: The Core Architecture of a Mind

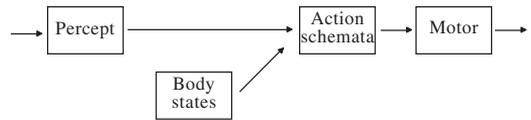


Figure 2: An Unminded Behavioral Architecture

The crucial components of this account are the beliefs and desires. For it is unlikely that possession of perceptual states alone, where those states are used to guide a suite of innate behavioral programs or fixed action schemata, could be sufficient for a creature to count as possessing a mind. Consider the architecture represented in Figure 2. The innate behaviors of many mammals, birds, reptiles, and insects can be explained in such terms. The releasing factors for an innate behavior will often include both bodily states (e.g., pregnancy) and perceptual information, with the perceptual states serving to guide the detailed performance of the behavior in question, too.² (Consider a butterfly landing neatly on a leaf to lay her eggs.) But engaging in a suite of innately coded action patterns isn't enough to count as having a mind, even if the detailed performance of those patterns is guided by perceptual information. And nor, surely, is the situation any different if the action patterns are not innate ones, but are, rather, acquired habits, learned through some form of conditioning.

This paper will assume that it isn't enough for an organism to count as having a mind, either, that the animal in question should be merely *interpretable as* possessing beliefs and desires. For as Dennett (1987) has taught us, we can adopt what he calls 'the intentional stance' in respect of even the most rigidly pre-programmed of behaviors. No, the architecture represented in figure 1 needs to be construed *realistically*. There needs to be a real distinction between the belief-states and the desire-states, in virtue of which they possess their distinctive causal roles (guiding

and motivating action, respectively). And these states must, in addition, be both discrete, and structured in a way that reflects their semantic contents. And their detailed causal roles, too (the ways in which particular belief states and particular desire states interact) must be sensitive to those structural features.

Dennett (1991) thinks that only language can provide these properties. It is only with the first appearance of the *Joycean machine*—the stream of inner verbalization that occupies so much of our waking lives—that humans come to have discrete structured semantically-evaluable states, where the interactions of those states are sensitive to their structures. So although Dennett might *say* that many animals possess simple minds, all he really means is that their behavior is rich enough to make it worth our while to adopt the intentional stance towards them. But he denies that non-human animals possess minds *realistically construed*, in the way that the present paper proposes to construe them.

To be minded means to be a thinker, then. And that means (it will hereafter be assumed) having distinct belief states and desire states that are discrete, structured, and causally efficacious in virtue of their structural properties. These are demanding conditions on mindedness. The question is: how simple can an organism be while still having states with these features?

2. MEANS/ENDS REASONING IN RATS

Dickinson and colleagues have conducted an elegant series of experiments, reported in a number of influential papers, providing evidence for the view that rats, at least, are genuine means/ends reasoners (Dickinson and Balleine, 1994; Dickinson and Shanks, 1995; Dickinson and Balleine, 2000). They argue that rats engage in goal-directed behavior, guided by beliefs about

causal contingencies, and by goals that are only linked to basic forms of motivation via learning. (We, too, have to monitor our own reactions to learn what we want; Damasio, 1994.) These arguments are convincing. But at the same time Dickinson advances a particular conception of what it takes to be a truly minded creature. The latter will need to be resisted if the position that insects have minds is to be defensible.

Here are some of the data (Dickinson and Balleine, 2000). Rats do not press a lever for food any more when hungry than when nearly satiated, *unless* they have had experience of *eating* that food when hungry. While the reward-value of the food is something that they know, the increased value that attaches to food when hungry is something that they have to learn. Similarly, rats caused to have an aversion to one food rather than another (via the injection of an emetic shortly after eating the one food but not the other) will thereafter stop performing a trained action causally paired with just that food, in the absence of feedback resulting from its actions (i.e., without receiving any rewards). However, they will do so *only* if they have been re-presented with the relevant food in the interval. They have to *learn* that they now have an aversion to the food in question. But once learned, they make appropriate decisions about which of two previously learned actions to perform—they know which action will make them sick, and choose accordingly.

Dickinson thinks that the data warrant ascribing to rats a two-tier motivational system much like our own. There is a basic level of biological drives, which fixes the reward value of experiences received. And there is an intentional level of represented goals (e.g., to eat *this* stuff rather than *that* stuff), which has to be linked up to the basic level via learning to achieve

its motivating status. Others have made similar proposals to explain aspects of the human motivational system (Damasio, 1994; Rolls, 1999).

In other experiments Dickinson and colleagues have shown that rats are sensitive to the degree of causal efficacy of their actions (Dickinson and Charnock, 1985; Dickinson and Shanks, 1995). In particular, the rat's rate of action drops as the causal connection between act and outcome falls. It even turns out that rats display exactly the same *illusions* of causality as do humans. The set-up in these experiments is that the probability of an event occurring (e.g., a figure appearing on a TV monitor, for the humans) or a reward being delivered (for the rats) is actually made independent of the action to be performed (pressing the space-bar, pressing a lever), while sometimes occurring in a way that happens to be temporally paired with that action. If (but only if) the unpaired outcomes are signaled in some way (by a coincident sound, say), then both rats and humans continue to believe (and to behave) as if the connection between act and outcome were a causal one.

The conclusion drawn from these and similar studies, then, is that rats are genuine means/ends reasoners. They possess learned representations of the goals they are trying to achieve (e.g., to receive a particular type of food). And they have acquired representations of the relative causal efficacy of the actions open to them. Taken together, these representations will lead them (normally, when not fooled by cunning experimenters) to act appropriately. However, Dickinson and colleagues claim that *only* creatures with these abilities can count as being genuinely minded, or as possessing a belief/desire cognitive architecture of the sort depicted in figure 1 (Heyes and Dickinson, 1990; Dickinson

and Balleine, 2000). Their reasoning is that otherwise the animal's behavior will be explicable in terms of mere innate motor programs or learned habits created through some form of associative conditioning. These further claims are unwarranted, however, as we shall see.

3. NON-ASSOCIATIVE LEARNING AND NON-CAUSAL INSTRUMENTAL REASONING

Dickinson assumes that if behavior isn't caused by means/ends reasoning in the above sense, then it must either be innate or the product of associative conditioning. But this assumption is false. The animal world is *rife* with non-associative forms of learning (Gallistel, 1990). Many animals (including the Tunisian desert ant) can navigate by dead reckoning, for example. This requires the animal to compute the value of a variable each time it turns—integrating the direction in which it has just been traveling (as calculated from the polarization of the sun's light in the sky; Wehner, 1994), with an estimate of the distance traveled in that direction, to produce a representation of current position in relation a point of origin (home base, say). This plainly isn't conditioned behavior of any sort; and nor can it be explained in terms of associative mechanisms, unless those mechanisms are organized into an architecture that is then tantamount to algorithmic symbol processing (Marcus, 2001).

Similarly, many kinds of animal will construct mental maps of their environment which they use when navigating; and they update the properties of the map through observation without conditioning (Gould and Gould, 1994).³ Many animals can adopt the shortest route to a target (e.g., a source of food), guided by landmarks and covering ground never before traveled. This warrants ascribing to the animals a mental map of their environment. But they

will also update the properties of the map on an on-going basis.

Food-caching birds, for example, can recall the positions of many hundreds or thousands of hidden seeds after some months; but they generally won't return to a cache location that they have previously emptied. Similarly, rats can be allowed to explore a maze on one day, finding a food reward in both a small dark room and a large white one. Next day they are given food in a (distinct) large white room, and shocked in a small dark one. When replaced back in the maze a day later they go straight to the white room and avoid the dark one (Gould and Gould, 1994). Having learned that dark rooms might deliver a shock, they have updated their representation of the properties of the maze accordingly.⁴

Many animals make swift calculations of relative reward abundance, too, in a way that isn't explicable via conditioning. For example, a flock of ducks will distribute 1:1 or 1:2 in front of two feeders throwing food at rates of 1:1 or 1:2 within one minute of the onset of feeding, during which time many ducks get no food at all, and very few of them experience rewards from both sources (Harper, 1982). Similarly, both pigeons and rats on a variable reward schedule from two different alcoves will match their behavior to the changing rates of reward. They respond *very* rapidly, closely tracking random variations in the immediately preceding rates (Dreyfus, 1991; Mark and Gallistel, 1994). They certainly are not averaging over previous reinforcements, as associationist models would predict.

Gallistel and colleagues have argued, indeed, that even classical conditioning—the very heartland of associationist general-purpose learning models—is better explained in terms of the computational operations of a specialized rate-estimating foraging system (Gallistel, 1990, 2000;

Gallistel and Gibbon, 2001). One simple point they make is that animals on a delayed reinforcement schedule in which the rewards only become available once the conditioned stimulus (e.g., an illuminated panel) has been present for a certain amount of time, will only respond on each occasion after a fixed proportion of the interval has elapsed. This is hard to explain if the animals are merely building an association between the illuminated panel and the reward. It seems to require, in fact, that they should construct a representation of the reinforcement intervals, and act accordingly.

Moreover, there are many well-established facts about conditioning behaviors that are hard to explain on associationist models, but that are readily explicable within a computational framework. For example, *delay* of reinforcement has no effect on rate of acquisition so long as the intervals between trials are increased by the same proportions. And the number of reinforcements required for acquisition of a new behavior isn't affected by interspersing a significant number of unreinforced trials. This is hard to explain if the animals are supposed to be building associations, since the unreinforced trials should surely *weaken* those associations. But it can be predicted if what the animals are doing is estimating relative rates of return. For the rate of reinforcement per stimulus presentation *relative to* the rate of reinforcement in background conditions remains the same, whether or not significant numbers of stimulus presentations remain unreinforced.

There are many forms of learning in animals that are not simply learned associations, then. And consequently there are many animal behaviors that are not mere habits, but that do not involve representations of causality, either. A bird who navigates to a previously established food

cache, using landmarks and a mental map on which the location of the cache is represented, certainly isn't acting out of habit. But then nor does the action involve any explicit representation of the causality of the bird's own behavior. The bird doesn't have to think, 'Flying in that direction will *cause* me to be in that place.' It just has to integrate its perception of the relevant landmarks with the representations on its mental map, then keying into action the flying-in-that-direction action schema. The causation can be (and surely is) left implicit in the bird's action schemata and behaviors, not explicitly represented in the bird's reasoning.⁵

But for all that, why should such animals not count as exemplifying the belief/desire architecture depicted in Figure 1? If the animal can put together a variety of goals with the representations on a mental map, say, and act accordingly, then why should we not say that the animal behaves as it does because it *wants* something and *believes* that the desired thing can be found at a certain represented location on the map? There seems to be no good reason why we should not. (And nor is there any good reason to insist that an animal only has genuine desires if its goals are acquired through learning.) Dickinson is surely misled in thinking that means/ends reasoning has to involve representations of causal, rather than merely spatial, 'means.'⁶

The difference between rats and many other animals is just that rats (like humans) are generalist foragers and problem solvers. For this reason they have to learn the value of different foods, and they are especially good at learning what acts will cause rewards to be delivered. But unlearned values can still steer intentional behavior, guided by maps and other learned representations of location and direction.

And so there can, surely, be minds that are simpler than the mind of a rat.

4. INSECTS (1):

INFLEXIBLE FIXED ACTION PATTERNS

How simple minded can you be? Do insects, in particular, have beliefs and desires? Few seem inclined to answer, 'Yes.' In part this derives from a tradition of thought (traceable back at least to Descartes) of doubting whether even higher mammals such as apes and monkeys have minds. And in part it derives from the manifest rigidity of much insect behavior. The tradition will here be ignored. (Some aspects of it have been discussed briefly in section 1 above.) But the rigidity requires some comment.

We are all familiar with examples of the behavioral rigidity of insects. Consider the tick, which sits immobile on its perch until it detects butyric acid vapor, whereupon it releases its hold (often enough falling onto the bodies of mammals passing below, whose skins emit such a vapor); and then when it detects warmth, it burrows. Or there are the caterpillars who follow the light to climb trees to find food, in whom the mechanism that enables them to do this is an extremely simple one: when more light enters one eye than the other, the legs on that side of its body move slower, causing the animal to turn towards the source of the light. When artificial lighting is provided at the bottom of the trees, the caterpillars climb downwards and subsequently starve to death. And when blinded in one eye, these animals will move constantly in circles. How dumb can you be! Right?

Even apparently sophisticated and intelligent sequences of behavior can turn out, on closer investigation, to be surprisingly rigid. There is the well-known example of the SpheX wasp, that leaves a paralyzed cricket in a burrow with its eggs, so that

its offspring will have something to feed on when they hatch. When it captures a cricket, it drags it to the entrance of the burrow, then leaves it outside for a moment while it enters, seemingly to check for intruders. However, if an interfering experimenter moves the cricket back a few inches while the wasp is inside, she repeats the sequence: dragging the insect to the burrow's entrance, then entering briefly once more alone. And this sequence can be made to 'loop' indefinitely many times over.

Or consider the Australian digger wasp, which builds an elaborate tower-and-bell structure over the entrance of the burrow in which she lays her eggs (Gould and Gould, 1994). (The purpose of the structure is to prevent a smaller species of parasitic wasp from laying *her* eggs in the same burrow. The bell is of such a size and hung at such an angle, and worked so smooth on the inside, that the smaller wasp cannot either reach far enough in, or gain enough purchase, to enter.) She builds the tower three of her own body-lengths high. If the tower is progressively buried while she builds, she will keep on building. But once she has finished the tower and started on the bell, the tower can be buried without her noticing—with disastrous results, since the bell will then be half on the ground, and consequently quite useless. Similarly, if a small hole is drilled in the neck of the tower, she seems to lack the resources to cope with a minor repair. Instead she builds *another* tower and bell structure, constructed on top of the hole.

In order to explain such behaviors, we do not need to advance beyond the architecture of Figure 2. The digger wasp would seem to have an innately represented series of nested behavioral sub-routines, with the whole sequence being triggered by its own bodily state (pregnancy). Each sub-routine is guided by perceptual input, and

is finished by a simple stopping-rule. But once any given stage is completed, there is no going back to make corrections or repairs. The wasp appears to have no conception of the overall goal of the sequence, nor any beliefs about the respective contributions made by the different elements. If this were the full extent of the flexibility of insect behaviors, then there would be no warrant for believing that insects have minds at all.

It turns out that even flexibility of behavioral strategy isn't really sufficient for a creature to count as having a mind, indeed. For innate behavioral programs can have a conditional format. It used to be thought, for example, that all male crickets sing to attract mates. But this isn't so; and for good reason. For singing exposes crickets to predation, and also makes them targets for parasitic flies who drop their eggs on them. Many male crickets adopt the alternative strategy of waiting silently as a satellite of a singing male, and intercepting and attempting to mate with any females attracted by the song. But the two different strategies are not fixed. A previously silent male may begin to sing if one or more of the singing males is removed (Gould and Gould, 1994).

Admittedly, such examples suggest that something *like* a decision process must be built into the structure of the behavioral program. There must be some mechanism that takes information about, for example, the cricket's own size and condition, the ratio of singing to non-singing males in the vicinity, and the loudness and vigor of their songs, and then triggers into action one behavioral strategy or the other. But computational complexity of this sort, in the mechanism that triggers an innate behavior, isn't the same as saying that the insect acts from its beliefs and desires. And the

latter is what mindedness requires, we are assuming.

5. INSECTS (2):

A SIMPLE BELIEF/DESIRE PSYCHOLOGY

From the fact that *many* insect behaviors result from triggering of innately represented sequences of perceptually guided activity, however, it doesn't follow that *all* do. And it is surely no requirement on mindedness that *every* behavior of a minded creature should result from interactions of belief states with desire states. Indeed, some of our own behaviors are triggered fixed action sequences—think of sneezing or coughing, for example, or of the universal human disgust reaction, which involves fixed movements of the mouth and tongue seemingly designed to expel noxious substances from the oral cavity. So it remains a possibility that insects might have simple minds *as well as* a set of triggerable innately represented action sequences.

In effect, it remains a possibility that insects might exemplify the cognitive architecture depicted in figure 1, only with an arrow added between 'body states' and 'action schemata' to subserve a dual causal route to behavior (see Figure 3).⁷ In what follows it will be argued that this is indeed the case, focusing on the minds of honey bees in particular. While this paper won't make any attempt to demonstrate this, it seems likely that the conclusion we reach in the case of honey bees will generalize to all navigating insects. In which case

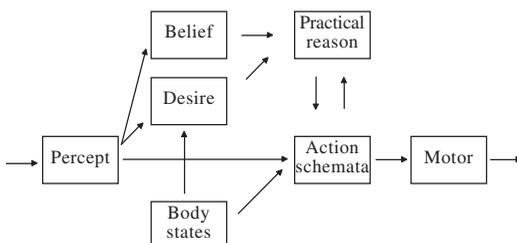


Figure 3: A Mind with Dual Routes to Action

belief/desire cognitive architectures are of very ancient ancestry indeed.

Like many other insects, bees use a variety of navigation systems. One is dead reckoning (integrating a sequence of directions of motion with the velocity traveled in each direction, to produce a representation of one's current location in relation to the point of origin). This in turn requires that bees can learn the expected position of the sun in the sky at any given time of day, as measured by an internal clock of some sort. Another mechanism permits bees can recognize and navigate from landmarks, either distant or local (Collett and Collett, 2002). And some researchers have claimed that bees will, in addition, construct crude mental maps of their environment from which they can navigate. (The maps have to be crude because of the poor resolution of bee eyesight. But they may still contain the relative locations of salient landmarks, such as a large free-standing tree, a forest edge, or a lake shore.)

Gould (1986) reports, for example, that when trained to a particular food source, and then carried from the hive in a dark box to new release point, the bees will fly *directly* to the food, but only if there is a significant landmark in their vicinity when they are released. (Otherwise they fly off on the compass bearing that would previously have led from the hive to the food.) While other scientists have been unable to replicate these experiments directly, Menzel et al. (2000) found that bees that had never foraged more than a few meters from the nest, but who were released at random points much further from it, were able to return home swiftly. They argue that this either indicates the existence of a map-like structure, built during the bees' initial orientation flights before they had begun foraging, or else the learned association of vectors-to-home with local landmarks. But

either way, they claim, the spatial representations in question are allocentric rather than egocentric in character.

As is well known, honey bees dance to communicate information of various sorts to other bees. The main elements of the code have now been uncovered through patient investigation (Gould and Gould, 1988). They generally dance in a figure-of-eight pattern on a vertical surface in the dark inside the hive. The angle of movement through the center of the figure of eight, as measured from the vertical, corresponds to the angle from the expected direction of the sun for the time of day. (E.g., a dance angled at 30° to the right of vertical at midday would represent 30° west of south, in the northern hemisphere.) And the number of 'waggles' made through the center of the figure of eight provides a measure of distance. (Different bee species use different innately fixed measures of waggles-to-distance.)

Honey bees have a number of innately structured learning mechanisms, in fact. They have one such mechanism for learning the position of the sun in the sky for the time of day. (This mechanism—like the human language faculty—appears to have an innate 'universal grammar.' All bees in the northern hemisphere are born knowing that the sun is in the east in the morning, and in the west in the afternoon; Dyer and Dickinson, 1994.) And they have another such mechanism for learning, by dead reckoning, where things are in relation to the hive. (Here it is 'visual flow' that seems to be used as the measure of distance traveled; Srinivasan et al., 2000.) They *may* have yet another mechanism for constructing a mental map from a combination of landmark and directional information. (At the very least they have the capacity for learning to associate landmarks with vectors pointing to the hive;

Menzel et al., 2000.) And they have yet another mechanism again for decoding the dances of other bees, extracting a representation of the distance and direction of a target (generally nectar, but also pollen, water, tree sap, or the location of a potential nest site; Seeley, 1995).

Although basic bee motivations are, no doubt, innately fixed, the goals they adopt on particular occasions (e.g., whether or not to move from one foraging patch to another, whether to finish foraging and return to the hive, and whether or not to dance on reaching it) would appear to be influenced by a number of factors (Seeley, 1995). Bees are less likely to dance for dilute sources of food, for example; they are less likely to dance for the more distant of two sites of fixed value; and they are less likely to dance in the evening or when there is an approaching storm, when there is a significant chance that other bees might not be capable of completing a return trip. And careful experimentation has shown that bees scouting for a new nest site will weigh up a number of factors, including cavity volume, shape, size and direction of entrance, height above ground, dampness, draftiness, and distance away. Moreover, dancing scouts will sometimes take time out to observe the dances of others and check out their discoveries, making a comparative assessment and then dancing accordingly (Gould and Gould, 1988).

Bees do not just accept and act on any information they are offered, either. On the contrary, they evaluate it along a number of dimensions. They check the nature and quality of the goal being offered (normally by sampling it, in the case of food). And they factor in the distance to the indicated site before deciding whether or not to fly out to it. Most strikingly, indeed, it has been claimed that bees will also integrate communicated information with the rep-

representations on their mental map, rejecting even rich sources of food that are being indicated to exist in the middle of a lake, for example.⁸

How should these bee capacities be explained? Plainly the processes in question cannot be associative ones, and these forms of bee learning are not conditioned responses to stimuli. Might the bee behaviors be explained through the existence of some sort of ‘subsumption architecture’ (Brooks, 1986)? That is, instead of having a central belief/desire system of the sort depicted in figure 3, might bees have a suite of input-to-output modular systems, one for each different type of behavior? This suggestion is wildly implausible. For (depending on how one counts behaviors) there would have to be at least five of these input-to-output modules (perhaps dozens, if each different ‘goal’ amounts to a different behavior), each of which would have to duplicate many of the costly computational processes undertaken by the others. There would have to be a scouting-from-the-hive module, a returning-to-the-hive module, a deciding-to-dance-and-dancing module, a returning-to-food-source module, and a perception-of-dance-and-flying-to-food-source module. Within each of these systems essentially the same computations of direction and distance information would have to be undertaken.

The only remotely plausible interpretation of the data is that honey bees have a suite of information-generating systems that construct representations of the relative directions and distances between a variety of substances and properties and the hive,⁹ as well as a number of goal-generating systems taking as inputs body states and a variety of kinds of contextual information, and generating a current goal as output. Any one of these goal states can then in principle interact with any one of

the information states to create a potentially unique behavior, never before seen in the life of that particular bee. It appears, indeed, that bees exemplify the architecture depicted in figure 3. In which case, there can be minds that are capable of just a few dozen types of desire, and that are capable of just a few thousand types of belief.¹⁰ How simple minded can you be? Pretty simple.

6. STRUCTURE-DEPENDENT INFERENCE

Recall, however, that the conditions on genuine mindedness that we laid down in section 1 included not just a distinction between information states and goal states, but also that these states should interact with one another to determine behavior in ways that are sensitive to their compositional structures. Now, on the face of it this condition is satisfied. For if one and the same item of directional information can be drawn on both to guide a bee in search of nectar and to guide the same bee returning to the hive, then it would seem that the bee must be capable of something resembling the following pair of practical inferences (using BEL to represent belief, DES to represent desire, MOVE to represent action—normally flight, but also walking for short distances—and square brackets to represent contents).

- (1) BEL [nectar is 200 meters north of hive]
 BEL [here is at hive]
 DES [nectar]
 MOVE [200 meters north]
- (2) BEL [nectar is 200 meters north of hive]
 BEL [here is at nectar]
 DES [hive]
 MOVE [200 meters south]

These are inferences in which the conclusions depend upon structural relations amongst the premises.¹¹

It might be suggested that we have moved too swiftly, however. For perhaps there needn't be a representation of the goal substance built explicitly into the structure of the directional information-state. To see why this might be so, notice that bees do not represent what it is that lies in the direction indicated as part of the content of their dance; and nor do observers acquire that information from the dance itself. Rather, dancing bees *display* the value on offer by carrying it; and observing bees know what is on offer by sampling some of what the dancing bee is carrying.

It might be claimed, then, that what really happens is this. An observing bee samples some of the dancing bee's load, and discovers that it is nectar, say. This keys the observer into its fly-in-the-direction-indicated sub-routine. The bee computes the necessary information from the details of the dance, and flies off towards the indicated spot. If it is lucky, it then discovers nectar-bearing flowers when it gets there and begins to forage. But at no point do the contents of goal-states and the contents of the information-states need to interact with one another.

This idea won't wash, however. Although the presence of nectar isn't explicitly represented in the content of the dance, it *does* need to be represented in the content of both the dancer's and the observer's belief-states. For recall that bees do not dance even for a rich source of nectar that is too far away (Gould and Gould, 1988). The distance-information therefore needs to be integrated with the substance-information in determining the decision to dance. Equally, observers ignore dances indicating even rich sources of nectar if the indicated distances are too great. So again, the distance information derived from the dance needs to be integrated with the value

information before a decision can be reached. So we can conclude that not only do bees have distinct information states and goal states, but that such states interact with one other in ways that are sensitive to their contents in determining behavior. In which case bees really do exemplify the belief/desire architecture depicted in Figure 3, construed realistically.

One final worry remains, however. Do the belief and desire states in question satisfy what Evans (1983) calls 'the Generality Constraint'? This is a very plausible constraint on genuine (i.e., compositionally structured) concept possession. It tells us that any concept possessed by a thinker must be capable of combining appropriately with any other concept possessed by the same thinker. If you can think that *a is F* and you can think that *b is G*, then you must also be capable of thinking that *a is G* and that *b is F*. For these latter thoughts are built out of the very same components as the former ones, only combined together with one another differently.

Now, bees can represent the spatial relationships between nectar and hive, and between pollen and hive; but are they capable of representing the spatial relationships between nectar and pollen? Are bees capable of thoughts of the following form?

BEL [nectar is 200 meters north of pollen]

If not, it may be said, then bees cannot be counted as genuine concept-users; and so they cannot count as genuine believer/desirers, either.

It is possible that this particular example isn't a problem. Foragers returning to a nectar site to find it almost depleted might fly directly to any previously discovered foraging site that is nearby; including one containing pollen rather than nectar, if pollen is in sufficient demand back at the hive. But there will, almost certainly, be other relationships that are never explicitly rep-

resented. It is doubtful, for example, that any scout will ever explicitly represent the relations between one potential nest site and another (as opposed to some such belief being implicit in the information contained in a mental map). So it is doubtful whether any bee will ever form an explicit thought of the form:

BEL [cavity A is 200 meters north of cavity B]

Rather, the bees will form beliefs about the relations between each site and the colony.

Such examples are not really a problem for the Generality Constraint, however. From the fact that bees never form beliefs of a certain kind, it doesn't follow that they *cannot*. (Or at least, this doesn't follow in such a way as to undermine the claim that their beliefs are compositionally structured.) Suppose, first, that bees do construct genuine mental maps of their environment. Then it might just be that bees are only ever *interested in* the relationships amongst potential new nest sites and the existing colony, and not between the nest sites themselves. But the same sort of thing is equally true of human beings. Just as there are some spatial relationships that might be implicit in a bee's mental map, but never explicitly believed; so there are some things implicit in our beliefs about the world, but never explicitly entertained, either, because they are of no interest. Our beliefs, for example, collectively entail that *mountains are less easy to eat than rocks*. (We could at least pound a rock up into powder, which we might have some chance of swallowing.) But until finding oneself in need of a philosophical example, this isn't something one would ever have bothered to think. Likewise with the bees. The difference is just that bees do not do philosophy.

Suppose, on the other hand, that bees *do not* construct mental maps; rather they learn a variety of kinds of vector information linking food sources (or nest sites) and the hive, and linking landmarks and the hive. Then the reason why no bee will ever come to believe that one nest-cavity stands in a certain relation to another will have nothing to do with the alleged absence of genuine compositional structure from the bees' belief states, but will rather result from the mechanisms that give rise to new bee beliefs, combined with the bees' limited inferential abilities.

For here once again, part of the explanation is just that bees are only interested in a small sub-set of the spatial relations available to them. They only ever compute and encode the spatial relationships between desired substances and the hive, or between landmarks and the hive, not amongst the locations of those substances or those landmarks themselves. And nor do they have the inferential abilities needed to work out the vector and distance relations amongst landmarks from their existing spatial beliefs. But these facts give us no reason to claim that bees do not really employ compositionally structured belief states, which they integrate with a variety of kinds of desire state in such a way as to select an appropriate behavior. And in particular, the fact that the bees lack the ability to draw inferences freely and promiscuously amongst their belief states should not exclude them from having any belief states at all. Which is to say: these facts about the bees' limitations give us no reason for denying that bees possess simple minds.

How simple minded can you be, then? Pretty simple; and a good deal more simple than most philosophers seem prepared to allow.¹²

University of Maryland

NOTES

1. One notable philosophical exception is Tye (1997). While the present paper shares some of Tye's conclusions (specifically, that honey bees have beliefs and desires) it offers different arguments. And the main focus of Tye's paper is on the question whether insects have phenomenally conscious experiences. This is quite another question from ours. Whether bees are belief/desire reasoners is one thing; whether they are phenomenally conscious is quite another. For a negative verdict on the latter issue, see Carruthers, 2000. It should also be stressed that the main question before us in this paper is quite different from the ones that formed the focus of Bennett's famous discussion of honey bee behavior (Bennett, 1964). Bennett argued that bee signaling systems are not a genuine language, and that honey bees are not genuinely rational in the fully-fledged sense of 'rationality' that is distinctive of human beings. (His purpose was to build towards an analysis of the latter notion.) Our concern, rather, is just with the question whether bees have beliefs and desires. (On this question, Bennett expressed no clear opinion.)
2. Do informational states that do not interact with beliefs and desires deserve to be counted as genuine *perceptions*? Common sense suggests that they do. Ask someone whether they think that the fish *sees* the net sweeping towards it through the water and they will answer, 'Of course, because it swims deftly in such a way as to avoid it.' But ask whether the fish has *thoughts* about the net or its own impending capture, and many will express skepticism. Moreover, the 'dual systems' theory of vision (Milner and Goodale, 1995; Clark, 2002) suggests that the perceptual states that guide our own movements on-line are not available for belief or desire formation, either. Yet we wouldn't deny, even so, that our detailed movements occur as they do because we *see* the shapes and orientations of the objects surrounding us.
3. Other animals navigate by the stars, or by using the Earth's magnetic field. Night-migrating birds study the sky at night when they are chicks in the nest, thereby extracting a representation of the center of rotation of the stars in the night sky. When they later leave the nest, they use this information to guide them when flying south (in fall in the northern hemisphere) and again when flying north (in spring in the northern hemisphere). The representations in question are learned, not innate, as can be demonstrated by rearing chicks in a planetarium where they observe an artificially generated center of night-sky rotation.
4. Similarly, western scrub jays will update the representations on their map of cached foods, and behave accordingly, once they learn independently of the different decay rates of different types of food. Thereafter they access the faster-decaying caches first. See Clayton et al., 2003.
5. The same is surely true of humans. When one wants a beer and recalls that there is a beer in the fridge, and then sets out for the kitchen, one doesn't explicitly represent one's walking as the *cause of* one getting the beer. Rather, once one knows the location, one just starts to walk.
6. Bermúdez (2003), too, claims that there can be no genuine decision-making (and so no real belief/desire psychology) in the absence of representations of instrumental causality. And this seems to be because he, too, assumes that there are no forms of learning between classical kinds of conditioning and genuine causal belief. But given the reality of a variety of forms of spatial learning (reviewed briefly above), it seems unmotivated to insist that sophisticated navigation behaviors are not really guided by decision-making, merely on the grounds that there are no causal beliefs involved.
7. Note that if the 'dual visual systems' hypothesis of Milner and Goodale (1995) generalizes to other perceptual modalities and to other species, then there should really be two distinct 'percept boxes' in figure 3, one feeding into conceptual thought and decision making, and another feeding into the action schemata so as to provide fine-grained on-line guidance of movement.

8. In these experiments two groups of bees were trained to fly to weak sugar solutions equidistant from the hive, one on a boat in the middle of a lake, and one on the lake shore. When both sugar solutions were increased dramatically, both sets of bees danced on returning to the hive. None of the receiving bees flew out across the lake. But this wasn't just a reluctance to fly over water. In experiments where the boat was moved progressively closer and closer to the far lake shore, more and more receiving bees were prepared to fly to it. See Gould and Gould, 1988.

9. Note that this satisfies the two criteria laid down by Bennett (1964, §4) for a languageless creature to possess beliefs. One is that the creature should be capable of learning. And the other is that the belief-states should be sensitive to a variety of different kinds of evidence.

10. The bee's capacity for representing spatial relations is by no means unlimited. There is probably an upper limit on distances that can be represented. And discriminations of direction are relatively crude (at least, by comparison with the almost pin-point accuracy of the Tunisian desert ant; see Wehner and Srinivasan, 1981). However, bees are also capable of forming a limited range of other sorts of belief, too. They come to believe that certain odors and colors signal nectar or pollen, for example (Gould and Gould, 1988).

11. Is there some way of specifying in general terms the practical inference rule that is at work here? Indeed there is. The rule might be something like the following: BEL [here is at x , F is m meters and n° from x], DES [F] @ MOVE [m meters at n°]. This would require the insertion of an extra premise into argument (2) above, transforming the first premise into the form, BEL [hive is 200 meters south of nectar].

12. The author is grateful to his students from a recent graduate seminar on the architecture of the mind for their skeptical discussion of this material, and to Ken Cheng, Anthony Dickinson, Keith Frankish, Robert Kirk, and Mike Tetzlaff for comments on an earlier draft.

REFERENCES

- Bennett, J. 1964. *Rationality: An Essay Towards an Analysis*. London: Routledge.
- Bermúdez, J. 2003. *Thinking without Words*. Oxford: Oxford University Press.
- Brooks, R. 1986. "A Robust Layered Control System for a Mobile Robot." *IEEE Journal of Robotics and Automation*, RA-2: 14–23.
- Carruthers, P. 2000. *Phenomenal Consciousness: A Naturalistic Theory*. Cambridge: Cambridge University Press.
- Clark, A. 2002. "Visual Experience and Motor Action: Are the Bonds Too Tight?" *Philosophical Review*, vol. 110, pp. 495–520.
- Clayton, N., N. Emory, and A. Dickinson. 2003. "The Rationality of Animal Memory: The Cognition of Caching." In *Animal Rationality*, ed. S. Hurley. Oxford: Oxford University Press.
- Collett, T., and M. Collett. 2002. "Memory Use in Insect Visual Navigation." *Nature Reviews: Neuroscience*, vol. 3, pp. 542–552.
- Damasio, A. 1994. *Descartes' Error*. Picador Press.
- Davidson, D. 1975. "Thought and Talk." In *Mind and Language*, ed. S. Guttenplan. Oxford: Oxford University Press.
- Dennett, D. 1987. *The Intentional Stance*. Cambridge, Mass.: MIT Press.
- _____. 1991. *Consciousness Explained*. Allen Lane.
- Dickinson, A., and B. Balleine. 1994. "Motivational Control of Goal-Directed Action." *Animal Learning and Behavior*, vol. 22, pp. 1–18.

- _____. 2000. "Causal Cognition and Goal-Directed Action. In *The Evolution of Cognition*, ed. C. Heyes and L. Huber. Cambridge, Mass.: MIT Press.
- Dickinson, A., and D. Charnock. 1985. "Contingency Effects with Maintained Instrumental Reinforcement." *Quarterly Journal of Experimental Psychology*, vol. 37B, pp. : 397–416.
- Dickinson, A., and D. Shanks. 1995. "Instrumental Action and Causal Representation." In *Causal Cognition*, ed. D. Sperber, D. Premack, and A. Premack. Oxford: Oxford University Press.
- Dreyfus, L. 1991. "Local Shifts in Relative Reinforcement Rate and Time Allocation on Concurrent Schedules." *Journal of Experimental Psychology*, vol. 17, pp. 486–502.
- Dyer, F., and J. Dickinson. 1994. "Development of Sun Compensation by Honeybees." *Proceedings of the National Academy of Science*, vol. 91, pp. 4471–4474.
- Evans, G. 1982. *The Varieties of Reference*. Oxford: Oxford University Press.
- Gallistel, R. 1990. *The Organization of Learning*. Cambridge, Mass.: MIT Press.
- _____. 2000. "The Replacement of General-Purpose Learning Models with Adaptively Specialized Learning Modules." In *The New Cognitive Neurosciences* (second edition), ed. M. Gazzaniga. Cambridge, Mass.: MIT Press.
- Gallistel, R., and J. Gibson. 2001. "Time, Rate and Conditioning." *Psychological Review*, vol. 108, pp. 289–344.
- Gould, J. 1986. "The Locale Map of Bees: Do Insects Have Cognitive Maps?" *Science*, vol. 232, pp. 861–863.
- Gould, J., and C. Gould. 1988. *The Honey Bee*. Scientific American Library.
- _____. 1994. *The Animal Mind*. Scientific American Library.
- Harper, D. 1982. "Competitive Foraging in Mallards." *Animal Behavior*, vol. 30, pp. 575–584.
- Heys, C., and A. Dickinson. 1990. "The Intentionality of Animal Action." *Mind and Language*, vol. 5, pp. 87–104.
- Marcus, G. 2001. *The Algebraic Mind*. Cambridge, Mass.: MIT Press.
- Mark, T., and R. Gallistel. 1994. "The Kinetics of Matching." *Journal of Experimental Psychology*, vol. 20, pp. 1–17.
- McDowell, J. 1994. *Mind and World*. Cambridge, Mass.: MIT Press.
- Menzel, R., R. Brandt, A. Gumbert, B. Komischke, and J. Kunze. 2000. "Two Spatial Memories for Honeybee Navigation." *Proceedings of the Royal Society: London B*, vol. 267, pp. 961–966.
- Milner, D., and M. Goodale. 1995. *The Visual Brain in Action*. Oxford: Oxford University Press.
- Rolls, E. 1999. *Emotion and the Brain*. Oxford: Oxford University Press.
- Searle, J. 1992. *The Rediscovery of the Mind*. Cambridge, Mass.: MIT Press.
- Seeley, T. 1995. *The Wisdom of the Hive: The Social Physiology of Honey Bee Colonies*. Cambridge, Mass.: Harvard University Press.
- Tye, M. 1997. "The Problem of Simple Minds." *Philosophical Studies*, vol. 88, pp. 289–317.
- Walker, S. 1983. *Animal Thought*. London: Routledge.
- Wehner, R. 1994. "The Polarization-Vision Project." In *The Neural Basis of Behavioral Adaptations*, ed. K. Schildberger and N. Elsner. Gustav Fischer.
- Wehner, R., and M. Srinivasan. 1981. "Searching Behavior of Desert Ants." *Journal of Comparative Physiology*, vol. 142, pp. 315–338.