

Planning, Memory, and Decision Making

Amanda Seed, Nicola Clayton, Peter Carruthers,
Anthony Dickinson, Paul W. Glimcher,
Onur Güntürkün, Robert R. Hampton, Alex Kacelnik,
Murray Shanahan, Jeffrey R. Stevens, and Sabine Tebbich

Introduction

Animals make a variety of choices, and it is a fair start to assume that the psychological mechanisms underpinning their choices will be adaptive in the sense of maximizing their net gain of resources, reproductive opportunities, predator avoidance and ultimately their fitness.¹ Even choices that initially appear simple can be complicated because adaptive decisions often involve trade-offs along multiple dimensions. All other things being equal, we would expect an animal to delay consumption of an immediately available green apple until it is ripe but not so long that it gets moldy, but all things are rarely equal, and an ideally optimal consumer is only a theoretical abstraction. Implementing ideally optimal choices may be seen as being tuned to how the fruit matures: to current temperature and humidity, to present and future needs, to competition with other consumers, etc. These are the complex trade-offs that cognitive mechanisms implement.

In this chapter we are motivated by the following sorts of questions: What sorts of information do animals use to make decisions, and what mechanisms underpin the actions of different animal species in different contexts? Can they act now to secure an outcome with value in the future, even if it has limited value in the present? Can they integrate disparate sources of information, and recognize when they do not have enough information to make a good choice?

¹ It is customary to frame biological research around the notion that traits are designed to maximize fitness. This is the approach we adopt because the psychological abilities that are at the center of this contribution are of course biological traits. The many virtues and vices of this approach have been discussed extensively so we will not belabor them here.

To what extent do these decisions result from deliberation, as opposed to the sculpting effect of consequences, either over the course of evolution or the life history, and how can we tell the difference from studying animal brains and behavior?

Although answers to these questions will vary for different species in different contexts, comparative study can help us to extract common evolutionary and psychological principles. Impressive progress toward answering these questions has been made in the past decades within several disciplines, but substantial disagreements about how to interpret findings persist, and there is plenty still to learn. Our aim has been to outline the state of the art, to clarify the points of disagreement, and to suggest future directions.

Decision Making: Who Decides, How, and Why?

What Is a Decision?

Stevens (this volume) broadly defines a decision as the results of an evaluation of possible options. This definition encompasses the study of decision making from several different perspectives, with the fundamental difference between them being who (or what) decides (i.e., evaluates) the options to determine the resulting choice (e.g., natural selection, the mind, neural networks). From an evolutionary perspective, decision making or “choice” simply describes the selection of one among a set of possible targets. Similarly, in economics, the processes by which decisions are reached do not figure among the list of priorities in the study of choices and preferences: preferences and decisions are what the subject does, not what it thinks about it or how it achieves it.

The psychological approach, however, is critically concerned with the process by which options are evaluated in the minds of individuals. Determining what counts as a “decision” from a psychological perspective is slightly more controversial. One possible definition of a decision is any process (at the psychological, algorithmic or neural level) that leads to choices and preferences. An alternative definition is that only a subclass of such processes qualifies as decision making. For example, Dickinson argues that if choices result from no additional mechanisms than those generating action when an option is presented without competition, then this process implies choice without decisions. Dickinson uses the sequential choice model (SCM) as an example (Shapiro et al. 2008). In this model, when the agent encounters a source of reward, its actions toward it tend to reflect the value of this source relative to its context. The strongest and most quantifiable measure of this value is the frequency distribution of latencies to respond to the stimulus. More valuable options produce, on average, shorter latencies. The SCM postulates that when two sources of reward in the environment are met simultaneously, the same distributions are elicited as when each is met on its own, and their outcomes

simply cross-sensor each other, since exerting one action removes the opportunity to use the alternative. In this case there is no special evaluation of the difference between the alternatives in the brain (or mind) of the animal. In other words, choices result from a horse race, and acting toward some option simply occludes the expression of any process addressed toward the competition. The SCM makes an important prediction: observed latencies toward any reward source in choice contexts should be shorter than those toward the same source when met alone, simply because only the left tail of each distribution is expressed and recorded by the observer. This is the opposite of what might be expected if choice involved an evaluation mechanism at the time of choice. Kacelnik and colleagues have provided empirical evidence in favor of SCM in several different paradigms in starlings, but whether this applies very generally across the animal kingdom (while interesting) is not crucial here (for data that SCM has difficulty explaining, see Mazur 2010).

The point is that, for Dickinson, if SCM were a correct description of the process underlying choice, then choice occurs without a decision, whereas for Kacelnik, SCM describes a mechanism for decision making that happens not to involve evaluation at the time of choice. This is to some extent a terminological distinction, but one that is important to keep in mind for interdisciplinary exchanges, as it exemplifies how different research programs need to toil to achieve mutual understanding.

From a neural perspective, the critical issue is how psychological choices (or decisions as broadly construed above) are physically realized. Before we explore the psychological processes that underpin decisions and actions, let us review the state of the art from this point of view for a particular taxon.

The Neural Mechanisms for Decision Making in Primates

Over the course of the last decade, significant advances have been made in our understanding of the basic architecture for decision making in humans and old world monkeys (for a review, see Glimcher 2011; Kable and Glimcher 2009). Current data strongly suggest that decision-making circuits in these animals can be described as being constructed from two sets of tightly interconnected networks. The first of these networks, located primarily in the frontal cortex and the basal ganglia, appears to be involved in learning and representing the values of the objects of choice. When, for example, a monkey repeatedly samples each of two food dispensing levers, regions in these areas of the brain come to represent the values of those two actions. While there is no doubt that many subsystems contribute to this valuation process, it is now widely accepted that the neurotransmitter dopamine participates in this process by encoding the difference between expected and obtained reward and broadcasting that signal throughout these areas. Multiple subsystems are now known to use these dopaminergic signals to compute and represent the values of action. Current evidence thus suggests the existence of interconnected sets

of valuation mechanisms that interact to yield an overall value for each action, good, or option. While significant debate has taken place over the structure of the interactions between these multiple systems, current evidence leans toward the suggestion that these interactions are largely additive in nature within the timescale of a single decision, although that interaction may be significantly more complex over longer timescales.

Current evidence then suggests that these learned values are projected to a decision-making network located in the frontal and parietal cortices which effectively selects, from among the currently available options, that single option which has the highest “value” (i.e., the highest firing rates in the frontal valuation networks). These data, gathered mostly from physiological recordings in monkeys, suggest that topographically organized maps encode the values of each available action or option as a firing rate at a distinct point in the topography in these frontal and parietal networks. Algorithmically, choice is proposed to occur when change in the intrinsic excitatory and inhibitory tone of these networks forces a competition between different points in one or more of these topographic maps. The result of this winner-take-all competition is the unique identification of the most highly valued option, a burst of neural activation which then gains access to the motor control circuitry through a biophysically fixed thresholding mechanism. It should also be noted, however, that the separation of the valuation and choice circuits appears to be one of degree rather than an absolute categorical boundary. Neurons in the frontal cortex and the basal ganglia clearly encode chosen actions, and there is good reason to believe that the winner-take-all process, though synchronized across many brain areas, might well be driven from more than one point in the network under different conditions. For this reason, the clear presence of mechanisms in frontal areas and in the basal ganglia must not be taken as excluding the possibility that choice, the winner-take-all process, involves these areas as well under some conditions.

Comparing Species

Historically, there has been little integration between the different approaches to the study of decision making, but the value of cross-fertilization of ideas, paradigms, concepts, and species models between approaches is increasingly being recognized. The comparative psychological approach to decision making allows for the exploration of relevant evolutionary pressures that can shape decision mechanisms. Though decision making is investigated in a variety of species, we have only a handful of cases in which socioecological factors are considered. To show the benefit of this approach, we outline two examples.

Just as cats can be skinned in many ways, but look the same once the deed is done, actions can be generated by many different neural and/or psychological processes yet appear the same when the animal behaves. If behaviors look the same, then how can we identify the different processes underlying them? If

different processes cause the same behavior, then why have divergent species evolved convergently toward these behaviors? Conversely, why do animals behave differently when faced with exactly the same choice, even if they are closely related and presumably using similar neural and cognitive processes to execute their choices? Can we use these similarities and differences among species to extract general evolutionary principles?

Different Mechanisms, Convergent Behavior

In many cases, two processes can cause the same behavior in one class of problems, but they would typically show different outcomes when circumstances are modified. If behavior is the same under all sorts of environmental transformations, then the underlying processes may still differ but not be distinguishable by behavioral experiments. In such cases, analysis at the neural level might be revealing. For instance, the same learning algorithm may account for data in mammals and insects, but surely the neural architecture underlying the process in each species would differ (albeit not necessarily at cellular level).

An interesting case is when distant species use clearly different mechanisms to generate similar behavior. In such cases one may conclude that there must be something important about the behavioral outcome that caused convergent evolution or selects for the maintenance of the trait. This may be illuminating when the behavior's function is not clear, as it may guide the functional analysis.

A concrete example, developed by Kacelnik and his colleagues, may help. These authors describe how choice can be controlled by the state-dependent value of reinforcement. For instance, if a subject learns about a food source when hungry and about another when satiated, the subject may overvalue the source found when in greater need, because the memory for the reinforcement experienced at the time of acquaintance is stronger. This may even occur if the animal possesses a veridical representation of the properties of the reward sources. For instance, in one study, starlings were trained with a blue key when hungry and with a red key when partly satiated (Pompilio and Kacelnik 2005). In one condition, pecking the blue key (in hungry sessions) resulted in food with a 15-second delay, while pecking at the red key (in partially satiated sessions) resulted in food after a 10-second delay. When the birds' preferences were tested at a later stage, they chose the blue key, even though their pecking behavior showed that they possessed accurate representations of both delays. In other words, preference was not due to the animals storing a distortedly short delay when hungry, because their pecking shows that the representation of the interval was accurate. Instead, a different valuation tag had been attached to the keys during training. Similarly, when locusts were trained with two odors signaling equally sized blades of grass when hungry or satiated, later they preferred the odor experienced in hungry sessions (Pompilio et al. 2006). In this case, however, the mechanism appears to be different at both neural and

algorithmic levels. Preexisting neurobiological information suggests that the modulation due to hunger is brought about by state-dependent differences in the odor receptors. The locusts seem to smell the odor associated with greater need more intensely. In summary, two distant species overprefer alternatives that have been met in the past when hungrier, but one does so in spite of remembering their properties accurately and the other by modulating the amplitude of the sensory receptors.

This being the case, the issue is why would such different species (the phenomenon is also observed in fish) end up choosing according to state at the time of learning, if this can cause the wrong outcome (as when starlings choose the more delayed reward)? Searching for the answer may be guided by the comparative observation itself, and especially by the divergence in underlying mechanisms: if locusts, fish, and starlings show a state-dependent valuation effect but they achieve it by different means, we infer that the net selective pressure is sufficiently strong to favor these mechanisms in the majority or more significant problems, even if they may occasionally be suboptimal. In this case, the comparative approach points to the likely existence of an adaptive explanation and helps us to infer what the decision process may be like in other species.

Different Decisions, Homologous Mechanisms

The comparative approach also allows us to explore how evolutionary pressures can shape decision mechanisms by testing decision making in closely related species that differ in key aspects of their ecology. This question has been tested in the study of risky and intertemporal choice.

Risky choices involve opting for alternatives that vary in the probability of receiving rewards of different sizes. In the laboratory, this is often tested by offering individuals a choice between a smaller, certain reward and a larger, more variable reward (e.g., 4 for sure vs. a 50% chance of 1 and a 50% chance of 7). Dozens of studies across a broad range of species, including insects, birds, and mammals, suggest that many species are averse to variance in reward amounts (Kacelnik and Bateson 1996). Heilbronner and colleagues (2008) tested whether foraging ecology could shape preferences in risky choice situations in chimpanzees and bonobos. Both species consume fruit. Chimpanzees, however, engage in cooperative hunting activities in which the group hunts monkeys or other small mammals. This foraging strategy involves a variance in whether the group is successful and whether an individual will receive any meat. Bonobos, in contrast, forage on terrestrial herbaceous vegetation, an abundant food resource. If foraging ecology shapes risk preferences over evolution, chimpanzees may be willing to accept more risks than bonobos. An experimental study on risky choice in both species is consistent with this hypothesis (Heilbronner et al. 2008). When given a choice between a fixed

and variable option, chimpanzees are more risk seeking than bonobos and even show a preference for the variable option.

In the related topic of intertemporal choice, individuals choose between options that vary in the timing and size of rewards, such as between smaller rewards available sooner and larger rewards available later. To test the effects of foraging ecology on intertemporal choice, Rosati et al. (2007) compared chimpanzees and bonobos in a choice task. Hunting in chimpanzees not only involves the risk of not capturing any prey but also involves a significant delay between making the decision to hunt and actually receiving the food. Bonobos face very little time delay when consuming vegetation because of its abundance in the environment. The foraging ecology hypothesis predicts, therefore, that chimpanzees will have evolved a greater willingness to wait for long delays than bonobos. Rosati et al. (2007) tested these species in an intertemporal choice task and again found results consistent with the foraging ecology hypothesis: chimpanzees waited longer than bonobos. A similar study on cotton-top tamarins and common marmosets also suggests that species differences in temporal preferences match temporal aspects of their foraging ecology (Stevens et al. 2005b). Thus, we have a number of cases in which comparisons across species can help resolve evolutionary questions about the selective pressures relevant for the evolution of decision mechanisms.

Let us now focus our attention on the different psychological mechanisms that underpin animal decision making; namely the processes by which information in the environment is translated into action.

Goal-Directed Behavior

As described above, animals can make choices based on past experience of the value of different options without knowing what they want, or how to get it; that is, behavior does not have to be goal directed in the psychological sense for animals to make good “decisions” from an evolutionary perspective. Nevertheless, it may be beneficial for animals to be able to adjust their behavior in accordance with up-to-date information about the value of different alternatives and the causal relationship between their efforts and the attainment of those goals. As Dickinson (this volume) explains, for animal behavior to be psychologically goal directed, the behavior must be shown (usually experimentally) to satisfy two criteria:

1. The goal criterion: Behavior must be immediately sensitive to changes in the value of the goal, such that changing the value of the goal in another context has immediate effects on actions that produce that goal, even when there has been no direct experience of the (updated) consequences of that action.
2. The instrumental criterion: Behavior must be sensitive to the causal relationship between an action and its consequences, such that if the

contingency between an action and its outcome is disrupted, behavior can be flexibly changed or ceased.

Taken together, the criteria were designed to bring about a 2 × 2 classification:

		<i>Instrumental Criterion</i>	
		Present	Absent
<i>Goal Criterion</i>	Present	Goal-directed behavior	Habits
	Absent	Pavlovian conditioning (some forms)	Pavlovian conditioning (other forms)

In contrast to habitual action, goal-directed behavior is rational and intelligent, as animals are shown to represent what they want and what they have to do to get it when they act. Does this framework, however, capture all cases of such behavior?

Are There Classes of Behavior that Defy This Classification?

Consider sophisticated forms of navigational behavior of the sort undertaken by honeybees, mediated by a cognitive map (see Menzel, this volume): a complex representation. A bee has become lost and wants to return to the hive, or to a particular feeder. When it locates a known landmark on its cognitive map, it calculates the direction and distance that it needs to fly to reach its goal and sets out accordingly. The resulting behavior is entirely novel and cannot be considered a conditioned response to a stimulus, nor explained in any merely associative way. Do cases in which navigation through space is mediated by a representation (e.g., cognitive map) and desire (outcome value) argue for a conferral of goal-directed status, according to a belief-desire psychology (see Dickinson, this volume)?

It seems unlikely that the bee knows that flying a particular vector will cause it to achieve its goal, for it seems unlikely that bees possess the concept of cause at all (nor indeed the action concept, flight). Rather, once the appropriate vector has been calculated, motor instructions for flight are issued, and the direction of flight is thereafter guided by matching the intended vector against current directional information calculated from the position of the sun or from polarized light. This causes the bee to fly in the correct direction. However, the bee itself does not represent its own actions as causing it to go in the right direction. It just acts, thereby, as a matter of fact, going in the right direction (given the correctness of its map-based representations and current calculations). One could thus argue that there are two distinct forms of goal-directed behavior represented in the animal kingdom: one in which the causal status of the animal’s own action is represented, and one in which it is left implicit as in navigation toward a goal.

Another potentially problematic example is sophisticated tool use, as carried out by humans and perhaps other animals. Sometimes individuals use a tool without knowing how it works (e.g., a television set's remote control). When the contingency between using the remote control and changing the channel breaks down, the mechanically naive individual stops trying to use it, or could feasibly learn to point it away from the television set to make it work. By responding flexibly to changes in the instrumental contingency, the action passes the instrumental criterion. However, other forms of tool use (e.g., raking food toward you with a stick) might be mediated not only by a representation of the causal relationship (raking causes goal attainment) but also of the specific generative mechanism (by contacting the object and exerting force on it). Such beliefs may be particularly robust to changes in the instrumental contingency; for example it would be hard to learn to push the reward away from you in order to pull it in. Some studies have provided evidence that some large-brained animals (such as corvids and apes) display knowledge of properties such as connection, weight, and solidity when solving problems, but whether this knowledge is integrated into a causal framework is an open question (Seed and Call 2009).

Dickinson (this volume) suggests that with these queries in mind, passing the instrumental criterion should be downgraded from being "necessary" for the assignment of goal-directed status to being "sufficient" (along with the goal criterion, which is necessary). In cases which fail the instrumental criterion, additional evidence would then have to be invoked to show that apparently maladaptive behavior (under what is in effect an "omission" contingency) is mediated by a belief about a generative causal mechanism to retain the goal-directed status. Developing paradigms that can provide convincing evidence of this is an important goal for future work.

Behaviors that Meet the Criteria Need Psychological Accounts

How can we explain goal-directed behavior at an algorithmic level? Dickinson provides two categories of explanation: one "rational" (belief-desire psychology), the other associative (the associative-cybernetic model, which is an example of simulation theory). One might ask if both classes of explanation are right at different levels of analysis. From a folk psychological perspective, the computation can be described as behavior driven by a belief (that acting causes a specific goal) and a desire (for that goal). At the algorithmic level, this can be cashed out in associative terms (e.g., simulation theory). An immediate problem with this idea is that it assumes that the current associative models can account for the computations for which animals are capable. Some findings, as described by Dickinson (this volume), such as causal reasoning about events in rats (Blaisdell et al. 2006) and some future oriented behavior in scrub jays (Raby et al. 2007) are not amenable to associative explanations,

as they are currently modeled, and therefore may demand recourse to a more sophisticated, representational architecture capable of inference and induction.

A parallel can be drawn with different approaches to planning deployed in cognitive robotics. The symbolic approach inherited from so-called classical artificial intelligence involves reasoning with sentences in a propositional language. The goal state, initial state, and effects of actions are represented in this language, and a reasoning system carries out deductive inference with these representations to find a sequence of actions that will achieve the goal. By contrast, a simulation approach uses so-called analogical representations, in which the structure of the representational medium reflects the structure of what is being represented. These spatially organized structures (such as two-dimensional arrays) are used to represent spatial structure (of an image, say), arranged over time, as in a movie. Using this approach, a simulation (based on past experience) is run forward from the initial state until a goal state is reached (Marques and Holland 2009; Shanahan 2006; Ziemke et al. 2005). One limitation of the use of a simulator for planning is that other forms of reasoning, such as explanatory reasoning, require additional mechanisms, whereas in the symbolic approach such reasoning is carried out within the same representational and inferential framework as planning (Shanahan 2006).

Further conversations between those trying to model artificial intelligence and real animal intelligence could be productive in trying to describe how rational, goal-directed action is algorithmically realized. It is fascinating to note that so far, some recourse to propositional representations is needed both to explain what animals do and to produce robots capable of doing the same.

Why Isn't All Behavior Goal Directed?

Rats and probably many other species show the capability for goal-directed behavior, in the sense defined by Dickinson (this volume). Given the additional flexibility afforded by goal-directed behavior compared to habits and conditioned responses, we might ask why, from an evolutionary perspective, it is limited to certain contexts, and why actions can become habitual (autonomous of the current value of the goal) in contexts that are extensive such as over-training (Adams 1982).

One possible advantage of habitual behavior under stable ecological circumstances might arise from the fact that it is evidence based. Animals may simply repeat what has worked in the past because this is a safe and normally predictive cue for what will follow. This may imply that behavior is less flexible than it would be if the animal were persistently evaluating the best route to achieve a goal, but it would run less risk of getting it wrong. Furthermore, because it is computationally simpler, it might be faster.

Choice behavior in very stable environments has been well studied in non-human animals. In a typical experiment, a rat, pigeon, or monkey faces a choice between two or more actions which offer different quantities of or delays to the

same resource (e.g., Herrnstein 1961; Lau and Glimcher 2005). The findings point toward conditioned rather than goal-directed behavior in such contexts, although note that in cases where the two options yield different outcomes (e.g., different types of food), behavior always appears to be goal directed. The general observation is that under certain conditions, subjects distribute their choice between the two or more actions in a way that matches the ratio of payoffs from the actions; this is the matching law. By itself, this is surprising because it implies not allocating all behavior to the action with higher payoff, and many authors differ in their view of how much the data really support matching when its outcome is clearly poorer than maximizing. Here, however, we stay out of this controversy and focus instead on how mechanisms that produce matching can be modeled. Modeling studies suggest that the subject is engaged in straightforward reinforcement learning (e.g., Corrado et al. 2005; Lau and Glimcher 2005). Interestingly, very little attention has been given to the question of how the learning rates for these behaviors are set (i.e., how quickly the distribution of choices is adjusted to the ratio of payoffs). One reason for this omission is that in the variable interval environments which have been studied most extensively, it is not possible to say what learning rate is optimal for a given set of environmental conditions.

One approach to this problem has been to study choices under conditions in which optimal learning rates can be defined precisely and then to ask whether different species produce behaviorally observed learning rates that are well correlated with these normative solutions. Studies of this kind now underway (Glimcher, pers. comm.), in which monkeys have to choose over water rewards, suggest that the learning rates observed in well-trained animals (animals with hundreds of thousands of trials of experience) do in fact approximate normative solutions. When environments are highly variable in reward magnitude or probability, then learning rates must be high to allow animals to track the rapidly changing environment. When environments are stable, the reverse is true: learning rates should be low to allow maximally precise estimates of the values of competing actions. In fact, not only is this broad pattern observed, but monkey subjects yield learning rates that are very close approximations of the precise optimal rates. The same cannot, however, be said of humans performing the same tasks for monetary rewards. Humans do, broadly speaking, adjust their learning rates in the same directions after thousands of trials of training, but they do so in a much less precise manner. One observes that under these conditions, recent rewards appear to influence the behavior of the subjects more strongly than they should; in effect, the humans show learning rates that are too high (Glimcher, pers. comm.).

Why is there this species difference? It could be due to the differences in degree of training or in reward type, but another possibility may be considered: a second system for valuation may be operating in humans which interacts with the (in this case) more normatively valuable reinforcement learning system. To

begin to test this latter possibility, one can require that human subjects, while performing tasks of this kind, also perform secondary concurrent tasks that may consume the time or resources of the hypothesized second system. To this end, DeWitt and Glimcher (pers. comm.) have human subjects perform a standard n-back number recall task while also performing the task described above. Under these conditions, it was found that the learning rates of the subjects were much closer to the normative rates observed in the monkeys. One possible conclusion that can be drawn from this observation would be that the n-back task effectively suppresses the output of a secondary system active under these conditions in the humans but not the monkeys.

In procedures employed by human cognitive psychologists (e.g., choice reaction time procedures), it is also observed that behavior appears to become automatic in the sense that responses can be performed in parallel (excluding, of course, peripheral interference) and are impervious to the imposition of a cognitive load. At present we do not know whether the behavioral autonomy (from current goal value) that characterizes habitual behavior and automaticity are the product of the same or different processes. Within the context of animal cognition, we need to develop paradigms for imposing an independent cognitive load while the animal is performing a goal-directed or habitual behavior. For example, the opportunity to perform a goal-directed or habitual (target) action could be provided during the retention interval of a delayed conditional discrimination so that the animal has to remember the identity of the stimulus while performing the target action. If the imposition of the memory task interferes with goal-directed but not habitual performance, we would demonstrate a concordance between automaticity and autonomy in the same behavior.

This line of reasoning suggests that goal-directed behavior, when held up as a contrast to automatic or autonomous behavior, is volitional and effortful. This contrast sounds remarkably like the distinction between conscious and unconscious processing made in humans. Later, in the section on consciousness, we will investigate the applicability of this contrast, which has been well studied in human psychology, to the study of animal consciousness.

Goal-directed actions are sensitive to the current value of an action's outcome. However, some animals go beyond this and take actions toward outcomes that will have value in the future, sometimes even when their current incentive value is low (food-caching birds and mammals such as squirrels can be full of acorns and nevertheless continue to forage for them to hoard the excess for the winter). This raises another controversial question: Are animals capable of imagining likely future events and needs?

Animal Future Planning

In classically associative models of valuation and choice, changes in the state of a chooser, such as changes in hunger state, can alter current valuation. In

goal-directed action, as discussed above, these changes in value can feed into action so that animals act appropriately given their current needs. However, in these models, anticipated future changes in the state of a chooser which have not yet been experienced cannot affect value or choice. Here we consider whether animals other than humans are able to forecast changes in the future state of the world (e.g., the anticipated future prevalence of a reward) and use this information to influence current decision making in a manner for which traditional associative models cannot account. This could be an adaptive ability for animals that experience large fluctuations in food availability, such as food-caching animals, which hide acorns and other foods that are only available at certain times of the year, and then live on those stores of food throughout the coming months.

The question of animal planning is controversial because it has been proposed that in humans, the ability to forecast future states of the world is intrinsically linked to episodic memory and that both stem from an ability of the individual to engage in “mental time travel” (Corballis, this volume). Episodic memory—the ability to recall specific prior states of the world visited by the individual—allows individuals to return effectively to those prior world states and to reexamine values, choices, and motivations at those prior times. The central idea of the mental time travel hypothesis is that future world states can also be sampled in this way, constructively, by creating future episodes and prospectively evaluating objects like values, choices and motivations under those hypothesized future conditions. Critically, Corballis maintains that, at least in its fully fledged form, this is a uniquely human ability. Studying nonhuman planning is, of course, crucial to evaluate this claim, and Clayton, Dickinson, and their colleagues have conducted a number of experiments on food-caching by western scrub jays to evaluate this claim. The first step along that path, however, does not come directly into conflict with the mental time travel hypothesis. As Dickinson (this volume) explains, we first need to demonstrate that animals act in the present in ways that cannot be accounted for by traditional associative accounts grounded in the animals’ present values and preferences. One can ask whether nonhuman animals, in their decision making, can take into account future internal and external states of the world even if they have never directly experienced the relationship between those states and the subject’s actions/decision. Only if the answer to this question is positive will an inquiry into exactly how such a computation is algorithmically realized (through mental time travel or some other mechanism) become pertinent.

Dickinson (this volume) reviews a number of recent experiments in which animals have been shown to take an action in the present (cache food, or select and transport a tool) for a future need (a lack of that food, or the opportunity to access an apparatus where a tool, if brought along, could be used to get food). Some of these experiments can be explained by an associative account supplemented with a memory mechanism to bridge the temporal gap between the action and the goal (mnemonic-associative theory), although not all recent

research is designed to be subjected to such an analysis. In particular, the work with primate tool retention is in its infancy, and we need to know more about the current motivational state of the animals, the incentive values of the food and tools on offer, and how value is generalized across tools that have similar perceptual or physical properties. This will allow the assumptions of the mnemonic-associative theory model to be fully tested. This is obviously an important direction for future research.

Controversies and Future Paradigms

Although the mnemonic-associative theory can account for many of the patterns of caching behavior arising from the scrub-jay experiments, the experiment conducted by Raby et al. (2007) on the ability of scrub jays to cache food for tomorrow's breakfast is a notable exception (see Dickinson, this volume). Having been trained that peanuts were available in one compartment at breakfast time, and kibbles in the other end compartment at breakfast time, the birds spontaneously cached kibbles in the compartment that served peanuts for breakfast and peanuts in the compartment that served kibbles for breakfast. In this experiment, the birds received a novel test of caching and therefore they could not have associated caching a particular food with either compartment. This experiment has been criticized (Shettleworth 2007) on the grounds that the birds may simply employ a heuristic to spread their caches of a particular food type rather than any cognitive plan of where they should cache the food tomorrow. According to this cache-spreading hypothesis, the jays will show a preference to cache peanuts away from sites previously associated with peanuts, an explanation that does not refer to the prospective aspect of mental time travel.

One way to test whether the jays can cache in the service of future needs while controlling for cache spreading is to ensure that all of the cache locations are associated with all food types, but at different times. For example, one could design a "kibble-for-breakfast, peanuts-for-tea" study in which the birds are given the opportunity to eat powdered kibbles in one compartment in the morning and then receive powdered peanuts in the same compartment in the afternoon, while also experiencing the reverse temporal pattern of feeding in the other end compartment. If the birds then at test cache with respect to the type of breakfast food available in each compartment, this preference would demonstrate that their behavior accorded with the temporal order of future needs in a way that could not be explained by cache spreading. In short, this would indeed be evidence of future planning.

A second test of such prospective cognition that transcends mnemonic-associative theory processes would be a test of novel rule integration. Consider the following two-phase experiment: Suppose that the jays are first trained that one food decays with time after caching, whereas the other food ripens—something that these jays readily learn (de Kort et al. 2005). In the second

phase they are given the opportunity to cache both a nonperishable and nonripening food, such as peanuts, in two visuospatially distinct caching trays. Thus, in phase 2 they learn that they can recover the caches they have made from one tray after a short delay, and from the other tray after a long delay. At issue is whether, when given the opportunity to cache the ripening and decaying foods from phase 1 in the two caching trays used in phase 2, the jays will prefer to cache the decaying food in the tray associated with a short delay and the ripening food in that associated with the long delay.

Such behavior on the novel test would be problematic for mnemonic-associative theory because the jays have never previously had the opportunity to associate the memory of caching the perishable and ripening food items in these trays with the ripened and decayed states of the foods. It does, however, call upon prospective processes in that the jays must integrate the anticipation of the future states of the food with those of the opportunity to recover from the two trays.

At present, the priority for research in the domain of animal planning will be to gather more evidence to bolster the claim that animals plan for the future in ways that transcend associative models, perhaps through experiments like those proposed above. For the present, therefore, researchers will continue to use terms such as “episodic-like memory” which sidestep the issue of whether the animals actually mentally experience or construct personal past or future events. In the next two sections, we will assess the evidence that animals reflect on their own memories and knowledge, and the extent to which any animal thinking resembles human conscious processing.

Metacognition: What We Know and What We Know We Don’t Know

As Hampton (this volume) explains, most paradigms in this research area require animals to discriminate between cases in which they know what they need to know in order to succeed at a particular test (e.g., discriminating stimuli, matching-to-sample, or locating a food reward) and when they don’t. Metacognition can be broadly defined as monitoring or controlling cognition by whatever mechanism works. Several dependent measures have been used to assess this ability, including the use of a “declining test” response to avoid making mistakes and ensure a low-value reward, “betting” on performance, or searching for more information. A brief overview of the main paradigms that animals can successfully solve is given below (for more details, see Hampton 2009):

- *Perceptual metacognition:* Animals are trained on a perceptual task, such as classification of fields of dots as either sparse or dense. The difficulty of this primary task can be varied from easy to difficult. Animals are simultaneously offered a “decline test” response which allows them to avoid particular tests and progress to another. Metacognition is

inferred when monkeys use the decline test response more on difficult than on easy trials. This pattern has been reported several times (for a review, see Smith 2009).

- *Retrospective betting*: Monkeys initially perform a perceptual task, much as described above. After completing tests, the monkeys are able to “give confidence ratings” by gambling either a large or small number of pellets on being correct. Metacognitive monkeys should gamble more food rewards following accurate responses and fewer on inaccurate responses. They do so, and they generalize to new tasks (for a review, see Kornell 2009).
- *Prospective memory*: Monkeys match to sample at delays long enough to produce considerable forgetting. At the end of memory delays, but before presentation of the test stimuli, monkeys choose between progressing to the memory test or declining the test. Declining is followed by a poor quality reward. Accurate matching is rewarded with a highly preferred food. Inaccurate test responses are followed by no reward. While most delay intervals end with a choice between taking the test and declining the test, sometimes only the choice to take the test is available. Monkeys demonstrate metacognition by being more accurate on trials they choose to take than on those they were forced to take.
- *Searching for information*: Subjects are presented with a set of opaque tubes in which a food reward is hidden. Subjects either witness the baiting or do not. At test, subjects can bend over and look down the length of the tubes to locate the food if they choose to, select a single tube and collect the reward, if correct. Subjects demonstrate metacognition by collecting information more often when they have not witnessed the baiting than when they have. Human children, chimpanzees, orangutans, and rhesus monkeys clearly showed this pattern of behavior, and while the case for capuchin monkeys was less clear, some capuchins made this differentiation under at least some conditions (Hampton 2009).

As Hampton (this volume) explains, metacognitive performance, as broadly defined, requires some type of discriminative stimulus that distinguishes between trials on which the animal knows the answer and trials on which it does not. Effective discriminative stimuli can take the form of either publicly available stimuli (such as the objective difficulty of the test) or private stimuli available only to the subject (e.g., a representation of one’s own knowledge, the strength of a memory trace, a “feeling of uncertainty,” or an emotional state such as anxiety). Carruthers (2008, 2009) has argued that there is no evidence to date that the discriminative stimulus used by subjects in the paradigms described above involves representations about one’s own mental states (do I know or do I remember some specific thing), which we could refer to as metacognition in the narrow sense, or meta-representation (Box 9.1). It will be

important for work in this field to move beyond demonstrations of metacognitive patterns of behavior to reductive analyses of the stimuli and processes controlling performance, in order to decide which of these two definitions is justified in a particular circumstance. For now, however, we will consider the adaptive value of metacognition, in its broadest sense, and consider hypotheses for the selective pressures that may have caused it to evolve.

Box 9.1 Defining metacognition.

Carruthers (2008, 2009) argues that most of the data alleged to support metacognition in the narrow sense admits to a common sort of anxiety-involving explanation. In each case the animal knows the structure of the experiment: it knows, for example, that it has to select the more dense of two patterns, or that it has to touch the longest of nine lines, or that it has to select the object that had previously been displayed on a screen. In cases where the patterns are hard to discriminate, the animal will experience anxiety at the thought of selecting either one of them, for it knows that the result is likely to be loss of a desired reward followed by a period of “time out.” The negative valence component of anxiety motivates the animal not to press either of the primary response keys. In contrast, the “opt out” key is known to move the animal on to the next trial without delay, and hence will not be negatively valenced. As a result, that is what the animal selects.

Likewise in the retrospective betting paradigm, animals know that if they select the “high stakes” symbol, they face large gains if they have already made the correct discrimination (e.g., touching the longest line), but large losses if they have not. In psychophysically difficult cases, the chances of gaining a large reward will be appraised as low and, again, the animals will experience anxiety at the thought of pressing the “high stakes” symbol. The “opt out” key, in contrast, will be mildly positively valenced, since it issues in a guaranteed small reward.

Similarly in the prospective memory paradigm, animals know that to gain the desired reward, they have to select the item that had originally been shown on the screen. In cases where that item has been forgotten (or is only imperfectly recalled), an appraisal of low chances of success, issuing in anxiety, should be expected. This need not mean that the animal knows (i.e., meta-represents) that it has forgotten. Rather, failing to find an answer to a non-meta-representational question (e.g., what was on the screen), the animal feels anxious at the thought of accepting the test.

Some instances of motivated search behavior admit of the same pattern of explanation. Lacking information that is needed for success in some task, the animal feels anxious before making a choice, and this then cues a fairly stereotyped “bending down to look” response. (Note that search, in general, cannot require metacognition in the narrow sense, or almost all creatures will turn out to be metacognitive. Rather, failing to find a representation of food, say, on its mental map, the animal automatically initiates a search.) However, experiments in which the animals not only bend to look, but on the first trial move around to adopt the appropriate position for looking into an unfamiliar container (Krachun and Call 2009) suggest that the animals have some understanding of visual perspective, as well as of the actions that they should take to achieve a required perspective. This might show metacognition for perceptual access, if not for belief or memory.

Adaptive Value of Metacognition

To address the potential adaptive value of metacognition, we should distinguish between situations in which metacognition is empirically demonstrable and situations in which it might be of adaptive value. To demonstrate the use of metacognition in nonverbal individuals, we must create situations in which metacognition is the only solution that would work. We must, however, bear in mind that metacognition (and other forms of cognition such as future planning, theory of mind, etc.) does not require these strict situations to evolve. Metacognition can evolve whenever it gives a selective advantage; to evolve, it does not require a problem for which metacognition is the only solution.

The putative biological significance of metacognition may be seen by relating it to choice under uncertainty. In a canonical test of metacognition in nonhumans, a subject is presented with a choice between a safe payoff versus taking a test whose outcome depends on knowledge. If we label the safe payoff as *F*(ixed) and the alternative as *R*(isky), the task is characterized by the relation $R_{low} < F < R_{hi}$, where the subscripts *low* and *hi* denote the payoffs in the test option if knowledge is low or high. A further condition that is helpful to add is that

$$F > \frac{R_{low} + R_{hi}}{2}. \quad (9.1)$$

This means that if the subject has no additional information other than the structure of the problem, when it has no sensitivity to its knowledge state, the two payoffs of the risky option are equiprobable, and then a higher expected payoff would result from choosing *F*.

To increase the accuracy of the comparison, let us add the probability of being in a high state of knowledge, (*p*) to the equation. Now, to maximize payoff, the subject must consider whether

$$F > p \times R_{hi} + ((1 - p) \times R_{low}). \quad (9.2)$$

If *p* is sufficiently small, this inequality will be true and then *F* should be chosen. The variable *p* is determined by the subject's knowledge about its state of low or high potential performance in the test. Tests of metacognition assume that the problem lies in the animal being well tuned to *p*, so as to choose whatever maximizes the expected payoff. Note, however, that this assumes that the expected payoff, and not its variance, is the only factor considered. If the subject values payoffs nonlinearly, then even in the absence of any additional knowledge about its knowledge state, it may prefer risky or fixed. If, in fact, the subject is sensitive to variance, in the sense that its behavior is not designed so as to maximize expected (average) payoff per se, then the results may be more difficult to interpret. For instance, if there is surviving threshold *T* so that $F < T < R_{hi}$, then the subject should choose the test whatever its state of knowledge.

This reflection can, of course, be made more rigorous and detailed, but suffice it to say that the putative biological value of metacognition is obvious when it is considered as an equivalent of p when the latter is a property of the distribution of probabilities in the world. Metacognition is as adaptive as knowledge of the probabilities of low or high outcome caused by factors other than knowledge.

Here we suggest a number of contexts in which being tuned in to the state of one's own knowledge might lead to a maximization of pay-offs, in the hope that this exercise might stimulate ideas for future research paradigms.

Knowing When to Stop

Much of the work on metacognition has focused on paradigms that require animals to be sensitive to known unknowns (so as to opt out or search for information). Yet knowing when you know enough could be an adaptive feature in contexts such as search, sampling, exploration and practice, so as to devote no more than an optimal amount of time to these activities. Consider the case of HM, the deeply amnesic patient who had debilitating deficits in episodic memory. Although he was unable to remember any declarative information (memories and facts), he was able to remember procedural information, such as how to mirror write. Yet, although he knew how to mirror write, he did not know that he knew how to do so (Milner 1962). This distinction between procedural and declarative memory is critical and has implications for metacognition: because HM lacked any awareness of his memory, he did not know when to stop practicing.

Avian brood parasites provide a nonhuman example. These are bird species that instead of building their own nests and raising their young, they lay eggs in the nests of other species and exploit their hosts' parental care. Successful parasitism depends on laying in nests that are in the right phase; namely, where the host is still on a laying period and has not yet started to incubate. This matters because otherwise the parasite hatchling could be born later than the hosts and be outcompeted for food. To achieve this synchronicity, parasites such as the shiny cowbird in South America seem to scout for suitable nests during daytime, so as to use this knowledge on the following dawn to visit one of them and lay her own egg. Sampling here is costly and it only pays up to a point: knowledge of one suitable nest is close to enough, although knowing a few more may be a suitable insurance, should the preferred option be predated overnight. Thus, once a female cowbird has located one or a few suitable nests, she should stop searching and focus on foraging, necessary to complete the formation of the eggs that she will lay in the future. Some amount of knowledge is "enough," and the bird should act as if it knows that it already knows what it needs, and then stop searching. Here, then, is a potential benefit for a response of behavior to knowledge.

Avoiding Temptation

Most discussion thus far has focused on sensitivity to knowledge about some facet of the environment (e.g., memory of a food location). Sensitivity to one's own desires could, however, also be beneficial, if desires for short-term or immediate payoffs get in the way of larger long-term benefits. For example, an individual that is trying to give up smoking may avoid going to a bar or buying a packet of cigarettes, because she knows that these actions will increase her desire to smoke. To take an example from nonhuman animal behavior, sitting close to a receptive female may increase the desire to try to mate, but doing so may lead to retribution from the dominant male. Some evidence for animal's engaging in these kinds of behavioral strategies comes from temporal discounting studies that require an animal to inhibit taking an immediate small reward in order to receive a larger reward some time later. Some capuchin monkeys engage in self-distraction activities, such as looking away from the small food item, or in paradigms in which the smaller item must be retained and exchanged for the larger one, holding it at arm's length (Dufour, pers comm). Similarly, sensitivity to the fragility of a memory trace might increase behavioral strategies that increase one's chances of retaining information, such as marking a location or avoiding distraction. As in the previous example, behavioral steps are taken in the present to influence a future mental state. Again, these behavioral strategies need not be based on metacognitive abilities, but a putative advantage can be envisaged.

Mind Reading

The debate about animal metacognition (in the narrow sense; see Box 9.1) is linked to two competing accounts of the evolution of meta-representational capacities. In one account, first-person forms of meta-representation evolved first, for purposes of metacognitive control. In the other, third-person forms of meta-representation evolved first, for "Machiavellian intelligence" or for social cognition more generally. In the first account, mind reading grows out of metacognition. In the second, metacognition results from turning one's mind-reading abilities on oneself. From this latter perspective, we might predict that animals already capable of simple forms of mind reading should display matching forms of metacognition (in the narrow sense). Selective pressure for the evolution of this ability would therefore come from the social environment, leading to the prediction that animals living in complex societies would be more likely to show metacognitive performance.

In this context, it is important to note that developmental psychologists often distinguish between two forms of mind reading. There is a kind of goal/perception/knowledge-ignorance psychology which emerges early in development, followed somewhat later by an understanding of false belief, pretence, and misleading appearances. Likewise, Call and Tomasello (2008) have argued

that apes understand goals, perception, and knowledge states, while being incapable of reasoning about false belief. Thus, even if animals cannot monitor their own beliefs and memories, they might be capable of reasoning about their own goals and their own perceptual access to objects. Examples of animals reasoning about their own goals might include the scrub jays who plan for what they will want for breakfast the next day (Raby et al. 2007), and the monkeys who use distraction from desired objects to maximize their longer-term rewards (Dufour, pers. comm.). An example of animals reasoning about their own perceptual access might be the study by Krachun and Call (2009) in which apes moved around to the appropriate place to look into an unfamiliar container.

Future Directions

The central question for future work will be to discover which discriminative cues are used by animals for metacognitive control. It will be particularly important to devise paradigms that can distinguish between different sorts of internal cues (anxiety-based or meta-representational). It might be possible to test the anxiety-involving explanations of metacognitive behavior proposed in Box 9.1 through manipulations of mood. Animals put into a good mood are likely to experience less anxiety, whereas those put into a bad one are likely to experience more. One might predict, then, that animals in the first group would use the “opt out” key less than usual, whereas animals in the second group would use it more. There is no reason why a metacognitive account (in the narrow sense) should predict such a result.

A possible advantage of meta-representation over an anxiety-based discriminative cue might be that an agent can monitor which information is missing and search for the relevant information more specifically. One could test this with a variation of the matching-to-sample memory paradigms presented by Hampton (this volume). In this task, an animal is required to match to sample after some delay. In contrast to Hampton’s example, the stimulus that has to be remembered has two dimensions (e.g., a certain odor and a certain color), but, as in Hampton’s tasks, the animal has two options: (a) to decline the task and receive a small fixed reward or (b) take the test, where accurate matching is rewarded with a highly preferred food and inaccurate matching with no reward. In the proposed task, normal trials are interspersed with trials in which the information about one stimulus dimension is missing. In these trials the animal has an additional choice to the two just described: to peck at a key that reveals additional information about one of the two stimulus dimensions, color or odor, which is followed by the task presentation. Under the assumption that animals can monitor which information they still need, they should only choose to decline if the key with the irrelevant information is presented and continue if the key for the relevant information is presented.

Part of what makes goal-directed behavior, future planning, and metacognitive performance fascinating is because when humans perform these tasks, they do it through conscious reflection and manipulation of representations. To study these topics in nonhuman animals we usually operationalize them in terms of measurable behavior and try to model them algorithmically in terms of inputs and representations and even neural circuitry. This has been very productive. However, when traditional models of animal learning and choice fall short, the question of whether or not the animal manipulates representations of different options to make a decision (causal beliefs, future values, or past events) becomes pertinent. Is this a tractable question for us to ask of a nonverbal creature? Next we consider how we study this question in human animals, to see if some lessons can be learned.

Animal Consciousness: Insights from Human Psychology

The scientific study of consciousness has come a long way since the conference on “Animal Mind – Human Mind” (Griffin 1982), at which time the very possibility of applying the scientific method to consciousness was viewed with widespread skepticism. The methodological progress that has been made, however, is largely confined to the human case. In this final section, we ask whether observations and theories about the difference between conscious and unconscious thought in humans can be usefully applied to animal psychology.

Experimental Study of Consciousness in Human Psychology

Numerous experiments in this area are based on the idea of contrastive analysis (Baars 1988, 1997). In a contrastive analysis paradigm, closely matching conscious and unconscious conditions are compared and contrasted. For example, in a backward visual masking experiment (Breitmeyer and Öğmen 2006), the subject is presented with a series of visual images, which includes a target image quickly followed by a second image. By manipulating the timing between the target and its successor, two conditions can be obtained. In the masked (or unconscious) condition, the subject is unable to report the target even though it can be shown to have had a priming effect on a subsequent task. In the unmasked (or conscious) condition, the subject can report the target image. Having established the requisite contrasting conditions, neural activity in the conscious and unconscious conditions can then be compared, for example, using MRI or MEG (for an overview, see Dehaene et al. 2006a).

One advantage of contrastive analysis is that it sidesteps many of the philosophical difficulties attendant on the amorphous subject of consciousness and moves directly to the study of a distinction—the distinction between conscious and unconscious conditions—that is more amenable to empirical study. The challenge in the present context is how to transfer such paradigms to the

nonhuman case, where we cannot rely on verbal reports as an index of the conscious condition. Candidate methods for indexing the conscious condition in a nonhuman animal include post-decision wagering and nonverbal reports (for an overview, see Edelman and Seth 2009).

On the theoretical front, a number of overarching frameworks have been proposed. One of the most widely accepted is the global workspace theory (Baars 1988, 1997). According to global workspace theory, the conscious/unconscious distinction is realized in a brain with an architecture that comprises a large cohort of parallel processes and a global workspace. According to the theory, the parallel processes, or coalitions of these processes, compete for control of the global workspace, which is in essence a communications infrastructure capable of disseminating influence and information throughout the brain. So the winning coalition of processes can be thought of as broadcasting out to the full cohort of parallel brain processes, and the overall dynamic comprises episodes of broadcast punctuated by periods of competition (Shanahan 2010, chapter 4). Against the backdrop of a global workspace architecture, the theory hypothesizes that the unconscious condition corresponds to localized activity in the parallel processes, whereas the conscious condition results when widespread brain activity arises, mediated by the global workspace.

A further postulate of global workspace theory is that the conscious condition is cognitively efficacious in that it confers simultaneous access to numerous neural resources, including those associated with language, working memory, episodic memory, and mental time travel in general. In addition, the conscious condition permits the formation of novel coalitions of processes whose membership might be drawn from diverse areas of micro-expertise. It can thus be thought of as an integrative condition which, in the terminology of a modular theory of mind, allows the strict boundaries of specialist modules to be transcended, promoting cognitive flexibility (Shanahan and Baars 2005).

Neural Correlates in Human Brains

An obvious question for advocates of global workspace theory concerns what the neural substrate of the putative global workspace might be. According to Dehaene and colleagues, the global neuronal workspace should be sought in the cerebral white matter of the human brain (Dehaene et al. 1998; Dehaene and Naccache 2001), which carries long-range corticocortical and thalamo-cortical fiber tracts. Recent diffusion imaging studies of human cerebral white matter have produced connectivity matrices that can be analyzed using the mathematical theory of networks (e.g., Sporns 2010). This allows us to explore the question of whether white matter connectivity can support the neurodynamics required of a global neuronal workspace at a theoretical level. In particular, Shanahan (2010, chapter 5) has suggested that a modular small-world network with connector hubs is the right topology to support both the integrative and communicative functions attributed to a global workspace.

The question for this discussion is whether these two lines of inquiry that have proved fruitful in the study of human consciousness—namely, experiments that yield behavioral signatures of consciousness and the topological signatures that can be observed in neuroanatomy—could be usefully applied to nonhuman animals.

Topological Signatures in Animal Brains

Insofar as the tenets of global workspace theory are applicable in the context of animal cognition, it should be possible to establish a related set of results for nonhuman neuroanatomy. If the relevant white matter tractography can be conducted and the necessary connectivity matrices extracted for these species, then it would be possible to test whether the topological features, which are hypothesized by the theory to underpin the conscious condition in humans, are present in the cognitively well-endowed nonmammalian species such as corvids, parrots, and perhaps the octopus.

Another issue of importance is the relevance of midline thalamic activation to consciousness. Significant neurological evidence now suggests that activation levels in certain thalamic circuits are tightly correlated with levels of consciousness. Artificial activation of these areas can even induce apparently conscious states in persistently vegetative patients. Conversely, damage to these areas appears to reduce the level of consciousness in human patients. This suggests that the neurobiological signature of consciousness will likely include these areas.

There is already some evidence for some of these features in mammalian and avian brains. A systematic investigation would do well to include as many variables as possible to look at the covariation between the communicative and integrative features of the network and its potential to provide associating emotional experience as evidenced by the type of connectivity.

Possible Future Experiments Based on Human Psychology

Backward Masking

Animals are obviously unable to report verbally on whether or not they have seen a cue. However, experiments by Cowey and Stoerig (1995) on monkeys with blindsight (lateralized lesions of primary visual cortex) gained “reports” from the monkeys about whether or not they had seen a presented cue using two keys: one for presence and one for absence. They also trained them on a spatial discrimination. In the intact field, monkeys successfully performed both the presence/absence and the spatial discriminations (i.e., they reported both that they had seen the cue, and where it had been located). In the “blind” field they reported “no stimulus” but correctly performed the spatial discrimination. An analog of the backward masking experiments described above could

use this procedure to present nonlesioned individuals with short exposures to visual cues and manipulate presentation time to see if there is a point at which, like humans, monkeys report “no stimulus” but can still locate it.

Directed Forgetting

Directed forgetting is a paradigm that has been used to test for working memory in nonhumans. The idea is that if working memory involves some active processing, rather than the passive decay of memory traces, the procedures that interrupt the active maintenance of memory will impair performance. In contrast, if memory involves only passive decay of traces resulting from study, then such interfering events should have little or no effect on performance. Several studies have addressed this using match-to-sample paradigms. In these studies animals are trained with a “remember” cue and a “forget” cue. Animals come to expect a test of memory following a remember cue and to expect no test following a forget cue. On probe trials, subjects are exposed to the forget cue but unexpectedly tested anyway. The evidence supporting active working memory is that performance is typically worse on tests that follow a forget cue.

These studies, however, have been vulnerable to a variety of criticisms that weaken the case for active memory. One problem is that tests following the forget cue are unexpected, and presenting tests when they are not anticipated may be disruptive to memory. A related concern is that in most paradigms, the forget cue signals nonreinforcement. Signaling nonreinforcement may cause animals to have negative emotional states or have other consequences that lead to behaviors incompatible with correct choice behavior on probe trials (Zentall et al. 1995). Some experiments may adequately control for these alternatives and still show directed forgetting, but they are a distinct minority (Kaiser et al. 1997). Nevertheless, improving on these paradigms might be a useful direction for future research.

Spatial Memory

Cook et al. (1985) assessed active use of working memory by testing whether rats can use working memory strategically by shifting from a retrospective to a prospective code to control working memory load. Rats working on a 12-arm radial maze were removed from the maze for 15 minutes after visiting various numbers of arms. The interpolated delay of 15 minutes impaired performance according to a U-shaped function. If the delay was interposed after a few arms had been visited, or after most arms had been visited, the impact was small and accuracy remained high. The delay had greatest effect when rats had visited about half of the arms. The interpretation is that the rats strategically regulated working memory load by initially remembering which arms it had visited (a retrospective code) but then switched to remember which arms it had yet to

visit (a prospective code) when the number of visited arms exceeded the number of to-be-visited arms.

Dual Tasking

As discussed in the section on goal-directed behavior, paradigms that require an animal to perform two tasks—one of which must be solved by holding information in working memory (e.g., a delayed match to sample) and one which can be shown to require goal-directed action—could potentially present animals with problems if they are holding and manipulating multiple representations “in mind.” Manipulating the amount of training on the instrumental task to shift the behavior from goal-directed action to autonomous habit would be hypothesized to improve overall performance if this is the case. Some tasks with multiple components have already been run with nonhumans, which perhaps have some of these characteristics. Chimpanzees that had previously learned to move food away from a trap using their hands over the course of over a hundred trials were better able to solve a similar task if they were able to use the same action to move the reward. Subjects with the same experience that had to use a tool on the transfer task did instead worse, but some still performed above chance levels. Subjects with no experience either of the type of discriminative task, or the tool-using action, could not find the solution at all (Seed et al. 2009).

The ability to manipulate representations consciously seems to be relevant for several of the abilities that we have discussed in the previous sections, as well as the abilities discussed in the contributions on mind reading and communication. The development of paradigms that contrast conscious and unconscious processing could allow us to make a comparative study of working memory capacity in different species. Evolutionary change in this capacity could have an important role in enabling or constraining psychological adaptations based on representations (Coolidge and Wynn 2001; Seed and Byrne 2010).

Conclusions

In this chapter we have discussed the nature and potential adaptive value of complex cognitive traits in the context of animals making choices and decisions in their physical environments. Repeatedly, we have raised the difficulty of distinguishing complex skills, such as meta-representation, future planning and conscious reflection, from simpler alternatives—skills that are thought by many to be uniquely human. Sometimes the behavioral output of humans and animals in an experimental task is the same, yet we cannot assume the underlying processes are similar. This statement seems obvious, but controversy arises because some scientists argue that when comparing humans with other apes, the most parsimonious Darwinian assumption is that the underlying

processes are homologous (e.g., de Waal 2009). We argue, however, that a truly Darwinian approach dictates that the best mechanism is not necessarily the most complex (nor the simplest) but one that optimizes costs and benefits. From this perspective, mechanisms that need a lot of processing power or expensive neural substrate should only evolve if they are subject to strong selection pressure. Consequently, apes should share a costly human cognitive trait only if they receive benefits that exceed the costs.

As we have seen in this chapter, experimental methods for identifying underlying cognition can render ambiguous results. Better knowledge about neural bases could provide a solution to these methodological limitations. We have outlined several cases in which neurological research can inform our understanding of cognitive mechanisms, such as choice under uncertainty, state-dependent learning, and contrastive analysis of conscious and unconscious processing. The study of cortical substrates has also informed our understanding of goal-directed action (Balleine and Dickinson 1998).

Another essential source of information for a meaningful assessment of adaptive value is to study animals in the wild to ascertain whether a certain species is likely to encounter problems that demand such cognitive abilities (Kamil 1987). An acknowledgment of the fact that similar selective pressures could have led to human-like abilities in distantly related taxa has recently led to a valuable expansion of the research field. A truly comparative approach includes both the comparison of distantly related species that have similar ecological demands, which may have led to a convergence in their cognitive abilities, and the comparison of closely related species that differ in their ecology and may therefore differ in their cognitive abilities.

If we are to approach a deeper understanding of how thought has evolved in both human and nonhuman species, we must continue to integrate ecological, cognitive, and neurological research across carefully planned species comparisons.