

This PDF includes a chapter from the following book:

# **The Cognitive Animal**

## **Empirical and Theoretical Perspectives on Animal Cognition**

© 2002 Massachusetts Institute of Technology

### **License Terms:**

Made available under a Creative Commons  
Attribution-NonCommercial-NoDerivatives 4.0 International Public License

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

### **OA Funding Provided By:**

The open access edition of this book was made possible by generous funding from Arcadia—a charitable fund of Lisbet Rausing and Peter Baldwin.

The title-level DOI for this work is:

[doi:10.7551/mitpress/1885.001.0001](https://doi.org/10.7551/mitpress/1885.001.0001)

Peter Godfrey-Smith

---

### Basic Principles

What are animal minds, including human minds, *for*? Although some versions of this question are teleological in a sense that has no place within an evolutionary world view, other versions of the question can be coherently asked. We can ask: Is it possible to make a general statement about the kinds of selective pressures and advantages that have been responsible for the evolution of cognitive mechanisms? Why has the expensive and delicate biological machinery underlying mental life evolved?

I suggest that the way to approach this question is not only to stress the continuities between human mental capacities and cognition in non-human animals, but also to recognize continuities between cognition and a wider class of “protocognitive” mechanisms. Cognitive mechanisms are mechanisms for behavioral control. And behavioral control mechanisms comprise one subset of a larger class of mechanisms that have the function of enabling organisms to adapt to changing problems and opportunities in their environments.

My work attempts to defend and develop tools for exploring a view of the mind based on an evolutionary perspective of this kind. This involves a combination of philosophical argumentation, commentary on empirical research, and some modeling. The chief goals of this work are foundational; the aim is to formulate general principles that unify diverse projects of empirical work on simple forms of cognition, and make explicit the connections between this empirical work and philosophical questions about the place of mind in nature. It is also hoped that discussions of this kind might sometimes help those engaged in the empirical research.

In this essay I sketch the basic framework used and some simple mathematical models that

illustrate this framework. A range of empirical examples are discussed in Godfrey-Smith (2002).

I begin with the following principle:

*Environmental complexity thesis (ECT)*: The function of cognition (and of a range of proto-cognitive capacities) is to enable an agent to deal with environmental complexity.

Each of the key terms in the ECT requires clarification (see also Godfrey-Smith 1996). “Function” is understood here in a strong sense; the function of a trait or structure is the effect or capacity it has which has been responsible for its success under a regime of natural selection (see Allen et al. 1998). Cognition, as I said earlier, is understood very broadly. We can think of cognition as a biological toolkit used to control behavior; a collection of capacities which, in combination, allow organisms to achieve various kinds of adaptive coordination between their actions and the world. This toolkit typically includes the capacities for perception, internal representation, memory, learning, decision making, and the production of behavior.

As the term “toolkit” suggests, we need not expect to find some single set of tools across all the organisms with cognitive capacities; different organisms have different collections of behavior control devices, according to their circumstances and evolutionary history. Furthermore, the list I gave of some core elements of the toolkit (perception, internal representation of the world, memory, learning) should not be seen as describing a set of recognizable and distinct “modules” found in the same form in all cognitive systems that have them. Rather, this is a set of capacities realized in different ways in different organisms, capacities that shade into each other and off into other, noncognitive parts of the biological machinery.

There is no sharp line between “real” cognition and a range of processes that we can call “protocognitive.” By any normal standard,

plants and bacteria (for example) do not have minds and do not exhibit cognition. But plants and bacteria do exhibit some capacities for flexible responses to environmental conditions, using environmental cues to control development and metabolism. Bacteria, for example, can modify their metabolism in order to take advantage of changes in the local food supply. In the case of many plants, a wide range of developmental processes are sensitive to the details of their local access to light. In plants, there are also some rapid and reversible changes that merit the title of “behavior.” [See also Silvertown and Gordon (1989) although I think they might err on the side of including too much in the category of plant behavior.] Cognition shades off into other kinds of biological processes, and there is no point in trying to draw an absolute line between them.

The simplest biological capacities that we might consider protocognitive are cases of adaptive flexibility in behavior, development, or metabolism that are controlled by a fixed response to a simple environmental cue. This category includes adaptive phenotypic plasticity of all kinds (Schlichting and Pigliucci 1998). As we add different types of flexibility of response and different kinds of inner processing of the output of perceptual mechanisms, we reach clearer and clearer cases of cognition. However, there is no single path that takes us from the simplest cases to the most elaborate. There are various ways of adding sophistication to mechanisms of behavioral control, ways that will be useful to different organisms according to their circumstances. The ability to expand or contract the range of stimuli coupled to a given response is one important sophistication (Sterelny 2001). The ability to learn through mechanisms such as classical conditioning, reinforcement, and imitation is another. Yet another is the ability to construct a “cognitive map” of spatial structures in the environment (Roberts 1998). It is an error to try to describe a single hierarchy of cognitive skills, from the simplest to the most complex.

The ECT asserts a functional connection between this broad sense of cognition and environmental complexity. But what is environmental complexity? I suggest that the most useful concept of complexity here is a simple one. Complexity is *heterogeneity*. Complexity is variety, diversity, doing a lot of different things or having the capacity to occupy a lot of different states.

There are many different kinds of complexity; as with skills, it is not just unnecessary, but positively mistaken, to try to devise a single scale to order all environments from the least to the most complex. Any environment will be heterogeneous in some respects and homogeneous in others. Whether a particular type of complexity is relevant to an organism will depend on what the organism is like—on the organism’s size, physiology, needs, and habits. The heterogeneity properties of environments are objective, organism-independent properties, but among the countless ways in which an environment is structured and patterned, only some will be relevant to any given organism. [See Levins (1968) for a classic discussion of these issues.] According to the ECT, the point of acquiring complex systems for behavioral control is to enable an organism to deal with heterogeneity in the challenges the environment presents and in the opportunities it offers.

Environmental complexity itself has many sources. It should not be thought that the only kind of patterns that constitute environmental complexity are those that are causally independent of the activities of the organisms under consideration. On the contrary, in many cases the environmental complexity that organisms must deal with is either a causal product of, or is constituted by, the activities of other organisms within the same population. Then we have a situation that can exhibit feedback, or a coupling of organism and environment (Lewontin 1985; Laland et al. 2000). The most graphic examples in this class—and perhaps the most important examples for hominid cognitive evolution—are

those that involve social interactions among individuals within a population. In these cases, the behavioral complexity of other organisms comprises part of the environmental complexity that each individual must deal with. Many hypotheses stressing the role of social or “Machiavellian” intelligence in hominid evolution (Byrne and Whiten 1988) are examples of the ECT.

Some of the original reasoning behind the social intelligence hypothesis illustrates ways in which discussion of general principles like the ECT might help guide empirical work. Nicholas Humphrey and other early advocates of the social intelligence hypothesis were motivated in part by the idea that human and other primate cognitive capacities are too complex for their evolution to be explained in terms of the demands associated with tasks like foraging and predator avoidance (Humphrey 1976). Problems of social living were seen as more likely to have the kind of complexity that would propel the evolution of high-powered mental machinery. This reasoning can be seen as making an implicit appeal to a principle about the overall function of cognition in dealing with environmental complexity. Explicit discussion of this kind of principle might lead to other insights about the kinds of selection pressures responsible for different aspects of the evolution of cognition. Can we make generalizations about the kinds of environmental complexity that select for imitative learning as opposed to simpler kinds of associative learning, for example, or planning intelligence as opposed to simpler kinds of goal-directed behavior (Sterelny 2001)?

---

## Models

Cognitive mechanisms provide a way to track the differences between environmental situations and adjust behavior accordingly. But when is it worth trying to track what is going on, rather than just producing a fixed behavior and hoping for the best? Intuitively, we might expect the

value of a flexible approach to depend on the relationship between a number of factors. Are the kinds of environmental heterogeneity that an organism confronts of such a kind that different responses are *worth* producing? Does the organism have some sufficiently reliable way to *track* which state of the world it is confronting or is likely to confront? Are the *costs* of cognitive mechanisms and information processing outweighed by the advantages of flexibility?

These issues can be investigated and illuminated with some simple mathematical models. The models make a huge number of idealizations, but they do give us a way of moving beyond vague arguments and intuitions about when tracking the world is likely to be adaptive. The first model I discuss was proposed independently by Nancy Moran (1992) and Elliot Sober (1994), and resembles various earlier discussions, such as that by Lively (1986). More complex models that explore similar themes have been developed by David Stephens (1989, 1991).

This first model is the most simple, general, and “ground-floor” model of the adaptive advantages of plasticity or flexibility. The model describes optimal phenotype or strategy, but does not address the dynamics of evolution itself. So it would have to be augmented to function as a genuine evolutionary model. The model applies not only to behavioral flexibility but to phenotypic plasticity in general.

We assume there are two alternative states of the world,  $S_1$  and  $S_2$ . These states are encountered by the organism with the probabilities  $P$  and  $(1 - P)$ , respectively. The organism has available two phenotypic states, or behaviors,  $C_1$  and  $C_2$ . The payoff for producing behavior  $C_i$  when the world is in state  $S_j$  is  $V_{ij}$ . When the world is in  $S_i$ , the behavior with the highest payoff is  $C_i$ . The behaviors,  $C_1$  and  $C_2$ , must be distinguished from the *strategies* available to the organism. The three possible strategies are

All-1: always produce  $C_1$

All-2: always produce  $C_2$

Flex: produce  $C_1$  or  $C_2$ , depending on the state of an environmental cue

The point of the model is to describe the situations in which Flex is the best strategy.

The cue used by the flexible strategy is the state of an environmental variable that provides some information about whether the world is in  $S_1$  or  $S_2$ . Given that the world is in some particular state, the reliability properties of the flexible strategy based on the cue can be expressed in a matrix of the probabilities that a specific response will be made:

Response likelihoods,  $\Pr(C_i | S_j)$

	$S_1$	$S_2$
$C_1$	$a_1$	$(1 - a_2)$
$C_2$	$(1 - a_1)$	$a_2$

Here are the expected payoffs of the three strategies:

$$E(\text{All-1}) = PV_{11} + (1 - P)V_{12} \tag{1}$$

$$E(\text{All-2}) = PV_{21} + (1 - P)V_{22} \tag{2}$$

$$E(\text{Flex}) = P[a_1V_{11} + (1 - a_1)V_{21}] + (1 - P)[(1 - a_2)V_{12} + a_2V_{22}] \tag{3}$$

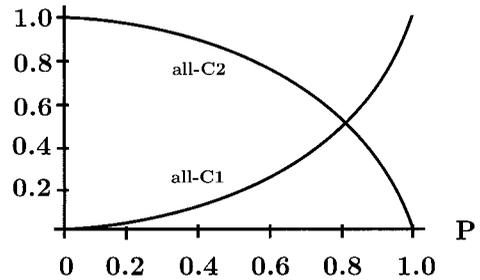
This simple model includes an assumption that some persons may see as biasing the case in favor of plasticity. We assume that plasticity comes for “free”; we are not imposing additional costs for setting up and maintaining the mechanisms that make plasticity possible.

Before we ask about the conditions under which Flex is the best strategy, we can determine first when one inflexible strategy is better than the other:

$$P(V_{11} - V_{21}) > (1 - P)(V_{22} - V_{12}) \tag{4}$$

In signal detection theory (discussed later), the quantity  $(V_{11} - V_{21})$  is known as the *importance* of  $S_1$ , and  $(V_{22} - V_{12})$  is the importance of  $S_2$ . The importance of a state of the world is the size

$a_1 (=a_2)$



**Figure 18.1**  
Flexible strategy. Minimal acceptable level of  $a_1$  and  $a_2$  for various values of  $P$ .

of the difference between the payoffs associated with suitable and unsuitable responses to that state. Sober (1994) introduces a useful term here. The expected importance of a state of the world is its importance multiplied by its probability. So according to formula (4), All-1 is better than All-2 if  $S_1$  has a higher expected importance than  $S_2$ .

So if you are going to always do the same thing, you should do the thing suited to the state of the world with the higher expected importance. When will the flexible strategy be better than both inflexible options? Flex is better if and only if:

$$\begin{aligned} a_2/(1 - a_1) > (1 - P)(V_{22} - V_{12})/P(V_{11} - V_{21}) \\ > a_1/(1 - a_2) \end{aligned} \tag{5}$$

An example is given in figure 18.1. For simplicity, I assume that  $a_1 = a_2$ . The payoffs used are as follows:  $V_{11} = 3$ ,  $V_{12} = 1$ ,  $V_{21} = 1$ ,  $V_{22} = 10$ . The figure graphs the minimum acceptable level of  $a_1$  and  $a_2$  for various values of  $P$ . The area above both lines is the only area where flexibility is favored over both inflexible options.

In this model, we aim only to describe the behavioral profile of a well-adapted individual. The model does not describe the internal psychological structure of the agent. However, Christopher

Stephens (in press) has extended the model in a psychologicistic way. Stephens accepts that the behavioral profile of a well-adapted organism is described by the model given here, but adds that an actual organism can also be psychologically described using a matrix of subjective probabilities and subjective utilities. These are the internal causes of the organism's behavior. Clearly, if an organism's subjective probabilities and subjective utilities exactly match the objective probabilities and the real fitnesses, the organism will do well. But the consequences of inaccurate subjective probabilities and pathological utilities can be complicated because combinations of "bad" subjective probabilities and utilities can interact and sometimes result in an adaptive behavioral profile.

Stephens' extension of the model is interesting because it connects the biological modeling of the evolution of adaptive behavior with a well-studied model of individual psychological processing, the Bayesian model. In other work (Godfrey-Smith 1996) I have connected the simple model of plasticity with another modeling tradition, signal detection theory.

Signal detection theory (Green and Swets 1966) augments the model described here by giving a finer-grained description of how organisms use cues to achieve adaptive flexibility. The model described here looks at relationships between flexible and inflexible ways of dealing with environments while assuming that the cue used by a flexible strategy is a "given," a fixed constraint within which the organism optimizes. Signal detection theory enables us to ask, given the general nature of an organism's physical connections to the world, what is the best cue available for guiding its behavior with respect to a particular problem?

Suppose that the organism must fashion its "cue" by monitoring the concentration of some environmental chemical, which varies in a continuous way. Two likelihood functions describe the relation between the concentration of the chemical and the alternative states of the world.

These functions,  $F(X | S_1)$  and  $F(X | S_2)$ , contain information about how likely a particular value of  $X$  is, given a particular state of the world. Let us suppose that both these functions are normal distributions and that the mean of  $F(X | S_2)$  is higher than that of  $F(X | S_1)$ . The likelihood ratio function,  $lr(X)$ , is the ratio between the values of the functions. That is,  $lr(X) = F(X | S_2)/F(X | S_1)$ . In the case described here,  $lr(X)$  will be a continuously increasing function. The likelihood ratio of an observation gives us a way of describing the evidential quality of that observation.

Let us suppose that because of the payoff matrix and the value of  $P$ , the best inflexible strategy is to always produce  $C_1$ . Which observed values of  $X$  are sufficiently good evidence for  $S_2$ , that it is worth producing  $C_2$  instead?

Signal detection theory can determine a threshold value of  $X$ , such that if this value or a higher one is observed, then this observation is worth acting on (the expected payoff from producing  $C_2$  given this observation is higher than the expected payoff from producing  $C_1$ ). That threshold in turn gives us the optimal way for an agent to treat levels of  $X$  as a cue directing behavior. The threshold is described initially in terms of the minimal acceptable likelihood ratio of an observation. This can be used to give us in turn the threshold value of the observable variable  $X$ . The best threshold is the value of  $X$  so that:

$$\begin{aligned} & F(X | S_2)/F(X | S_1) \\ & > P(V_{11} - V_{21})/(1 - P)(V_{22} - V_{12}) \end{aligned} \quad (6)$$

The kind of assessment being done using the signal detection model is very similar to that seen in the simple Moran–Sober model of plasticity. Again the crucial relationship is that between the reliability properties of what is observed and the expected importances of the alternative states of the world. These in turn depend on the costs of different kinds of errors and the benefits of different kinds of correct decisions. However, signal

detection theory gives us a finer-grained way of describing the use of environmental cues than the other model does because it does not simply take a cue with its reliability properties as a given, but instead describes how an organism might use a continuously varying environmental variable as a cue.

We can also add to a model of this kind various ways of representing the costs of plasticity. I supposed that the likelihood functions  $F(X|S_1)$  and  $F(X|S_2)$  were normal distributions with different means. Clearly, the likelihood ratio (which measures evidential quality) will depend not just on the means but also on the variances of these distributions. The smaller the overlap of the distributions, the more that can be learned from observing  $X$ . We might suppose that the organism makes its observation by a sampling process of some kind and that the likelihood functions  $F(X|S_1)$  and  $F(X|S_2)$  apply to a sample of a given size. If the organism were to observe a larger sample, it could in effect observe a signal with associated likelihood functions that have smaller variances and hence less overlap. But such sampling might have costs. In that case, the organism with optimal cognitive mechanisms will strike a balance between the costs of sampling and the benefits of observing a more informative signal, just as a scientist does when determining how large a sample is needed to answer a statistical question.

These optimality models are obviously extremely simple and in many ways unrealistic. They are no substitute for full evolutionary models or for empirical descriptions of how animals deal with their environments via psychological mechanisms. However, the simplicity of the models gives them the ability to illustrate some themes in a clear and stark way. They give us a ground-floor understanding of what it means for an organism to make use of an environmental cue in controlling its behavior, and how the costs of various kinds of flexibility must be traded off against the benefits. As such, they

are small pieces of a future theory of mind that stresses the role of cognition and protocognition in controlling behavior to deal with environmental complexity.

---

## References

- Allen, C., Bekoff, M., and Lauder, G. (eds.) (1998). *Nature's Purposes: Analyses of Function and Design in Biology*. Cambridge Mass.: MIT Press.
- Byrne, R. W. and Whiten, A. (eds.) (1988). *Machiavellian Intelligence: Social Expertise and the Evolution of Intellect in Monkeys, Apes and Humans*. Oxford: Clarendon Press.
- Godfrey-Smith, P. (1996). *Complexity and the Function of Mind in Nature*. Cambridge: Cambridge University Press.
- Godfrey-Smith, P. (2002). Environmental complexity and the evolution of cognition. In *The Evolution of Intelligence*, R. Sternberg and J. Kaufman, eds., pp. 223–249. Hillsdale, N.J.: Lawrence Erlbaum Associates.
- Green, D. M. and Swets, J. A. (1966). *Signal Detection and Psychophysics*. New York: Wiley.
- Humphrey, N. (1976). The social function of intellect. In *Growing Points in Ethology*, P. P. G. Bateson and R. A. Hinde, eds., pp. 303–317. Cambridge: Cambridge University Press. Reprinted in Byrne and Whiten (1988).
- Laland, K., Odling-Smee, J., and Feldman, M. (2000). Niche construction, biological evolution and cultural change. *Behavioral and Brain Sciences* 23: 131–175.
- Levins, R. (1968). *Evolution in Changing Environments*. Princeton, N.J.: Princeton University Press.
- Lewontin, R. C. (1985). The organism as the subject and object of evolution. In *The Dialectical Biologist*, R. Levins and R. Lewontin, eds., pp. 85–106. Cambridge, Mass.: Harvard University Press.
- Lively, C. (1986). Canalization versus developmental conversion in a spatially variable environment. *American Naturalist* 128: 561–572.
- Moran, N. (1992). The evolutionary maintenance of alternative phenotypes. *American Naturalist* 139: 971–989.
- Roberts, W. A. (1998). *Principles of Animal Cognition*. Chapter 7. Boston: McGraw-Hill.

Schlichting, C. and M. Pigliucci (1998). *Phenotypic Evolution: A Reaction Norm Perspective*. Sunderland, Mass.: Sinauer Associates.

Silvertown, J. and D. Gordon (1989). A Framework for plant behavior. *Annual Review of Ecology and Systematics* 20: 349–366.

Sober, E. (1994). The adaptive advantage of learning and a priori prejudice. In *From a Biological Point of View*, E. Sober, ed., pp. 50–69. Cambridge: Cambridge University Press.

Stephens, C. (in press). When is it selectively advantageous to have true beliefs? Sandwiching the better safe than sorry argument. *Philosophical Studies*.

Stephens, D. (1989). Variance and the value of information. *American Naturalist* 134: 128–140.

Stephens, D. (1991). Change, regularity, and value in the evolution of animal learning. *Behavioral Ecology* 2: 77–89.

Sterelny, K. (2001). *The Evolution of Agency, and Other Essays*. Cambridge: Cambridge University Press.

