

- city Ithaca, NY, 1990); ———, Z. Zhao, "Text Linking and Retrieval: Experiments for Textbook Components," *Technical Report 90-1125* (Computer Science Department, Cornell University, 1990).
44. D. C. Crouch, in *Proceedings of the Ninth Annual International Conference on Research and Development in Information Retrieval* (Association for Computing Machinery, New York, 1986), pp. 58–67; T. E. Doczkocs, *ibid.*, pp. 49–57; P. Ingwersen and I. Wormell, *ibid.*, pp. 68–76.
45. R. N. Oddy, *J. Doc.* 33, 1 (1977); S. E. Pollitt, *ASLIB Proc.* 36, 229 (1984); G. Guida and C. Tasso, *Automatica* 19, 759 (1983).
46. J. J. Rocchio, in *The Smart System—Experiments in Automatic Document Processing*, G. Salton, Ed. (Prentice-Hall, Englewood Cliffs, NJ, 1971), pp. 313–323; G. Salton, *ibid.*, pp. 324–336; E. Ide, *ibid.*, pp. 337–354; S. K. M. Wong and Y. Y. Yao, *J. Am. Soc. Inf. Sci.* 41, 334 (1990).
47. G. Salton and C. Buckley, *J. Am. Soc. Inf. Sci.* 41, 288 (1990).
48. T. C. Brauen, in *The Smart Retrieval System—Experiments in Automatic Document Processing*, G. Salton, Ed. (Prentice-Hall, Englewood Cliffs, NJ, 1971), pp. 456–484.
49. Y. Bar-Hillel, Ed., *Language and Information—Selected Essays on Their Theory and Application* (Addison-Wesley, Reading, MA, 1964), pp. 330–364.
50. This study was supported in part by NSF grant IRI 89-15847.

Animal Choice Behavior and the Evolution of Cognitive Architecture

LESLIE A. REAL

Animals process sensory information according to specific computational rules and, subsequently, form representations of their environments that form the basis for decisions and choices. The specific computational rules used by organisms will often be evolutionarily adaptive by generating higher probabilities of survival, reproduction, and resource acquisition. Experiments with enclosed colonies of bumblebees constrained to foraging on artificial flowers suggest that the bumblebee's cognitive architecture is designed to efficiently exploit floral resources from spatially structured environments given limits on memory and the neuronal processing of information. A non-linear relationship between the biomechanics of nectar extraction and rates of net energetic gain by individual bees may account for sensitivities to both the arithmetic mean and variance in reward distributions in flowers. Heuristic rules that lead to efficient resource exploitation may also lead to subjective misperception of likelihoods. Subjective probability formation may then be viewed as a problem in pattern recognition subject to specific sampling schemes and memory constraints.

THE EMERGING FIELD OF COGNITIVE SCIENCE ATTEMPTS TO explain the nature of thought and the appearance of intelligence. Cognitive analyses have mostly been applied to language capabilities and the acquisition of skills in humans (1), but have been expanded to include problem-solving and communication in animals (2–5). The cognitivist view suggests that the processing of information (by either animals or humans) involves three stages. First, sensory data are translated and encoded into a form that can be manipulated through mental operations. Second, encoded information is acted upon by specific computational rules. And third, these rules produce alternative "representational" states that depend on the informational input. The concept of "representation" remains controversial, especially for animals (5). However, these three stages may be viewed, in a less controversial manner, as three components

of a single dynamical system mechanistically tied to the organism's nervous system. The encoding of information would then correspond to initial inputs, computational rules correspond to transient dynamics, and representations would correspond to the equilibrium configurations resulting from the transient dynamics. The animal reaches a representation of the environment through the operation of specific computational rules applied to a particular pattern of incoming sensory information.

The computational rules used by organisms can be symbol-processing programs, as in most artificial intelligence models (6), or can be models of nervous systems, as in neural networks (7). My thesis is that these computational rules are evolutionarily adaptive. Different computational schemes may generate behaviors or representation of the environment that lead to different efficiencies in the use of resources, acquisition of mates, or acquisition of skills necessary for survival. Differential efficiencies may then confer different evolutionary advantages. The design features of information-processing ("cognitive architecture") may be subject to natural selection in a manner analogous to any other aspect of the organism's phenotype.

The link between mental process, cognition, and evolution originates in Darwin's writings (8) and has found continuous support from many investigators since the Darwinian revolution (9). Many more recent studies have explicitly examined the adaptive nature of specific mental processes in animals and have argued for varying degrees of adaptive specialization in mental function to accommodate specific ecological requirements (10). Few studies, however, have explicitly examined specific computational rules in the evolutionary ecology of organisms, though the adaptive nature of computational rules has been proposed (3). In this article, I summarize research on floral choice behavior in bumblebees (*Bombus* spp.) and argue for an evolutionary basis for the computational rules employed by bees as they acquire floral resources in their natural environment.

Bumblebee as a Model System

The choice of bumblebees as model experimental organisms was not arbitrary. Bumblebees have many features which make them ideal for examining the evolution of decision-making processes. Individual worker bumblebees are almost exclusively engaged in a

The author is in the Department of Biology, University of North Carolina, Chapel Hill, NC 27599–3280.

single task—collecting nectar and pollen for the colony. Worker bumblebees are sterile and, therefore, not concerned with acquiring mates or reproductive decisions. They are largely free from predation and, unlike honeybees, do not communicate with each other about resources. The energy reserves of the colony are tied to the colony's reproductive success, so individual foraging behavior has consequences for genetic representation in the future. Consequently, the bumblebee is an autonomous agent engaged solely in one important activity of obvious evolutionary significance.

Bumblebees are also easy to investigate under controlled conditions. Their foraging can be restricted to sets of artificial flowers whose energetic characteristics can be experimentally manipulated. In the experiments I describe, colonies of wild bumblebees were transferred into a large enclosure consisting of a frame surrounded with netting. Into this enclosure I placed a 1.2 m by 1.2 m by 6 mm sheet of clear plastic with 0.5-cm wells drilled 2.5 cm apart. This generated a field of 2016 possible flower positions. Artificial cardboard flowers of different colors were centered under specified wells generating, for example, a field of 100 blue and 100 yellow flowers randomly dispersed over the 1.44-m² foraging field (11). Known quantities of artificial nectar (diluted honey) were then dispensed into the artificial flowers to generate a specific distribution of rewards associated with each of the floral types. The nectar distributions, flower colors, and spatial distribution varied from experiment to experiment. The basic technique, however, remained the same.

Economics of Choice Under Uncertainty

Cognitive psychologists traditionally divide mental operations into lower level and higher level cognitive tasks. Lower level cognition would include some forms of learning (for example, habituation), pattern recognition, and perception; higher level cognitive tasks include language, communication, complex learning and memory, and problem-solving and task-acquisition (1). I will focus almost entirely on problem-solving and task-acquisition in bumblebees and more specifically on certain forms of “economic” decision-making applied to the efficient exploitation of limiting resources.

Since the 1940s the dominant model of human economic decision-making has been the expected utility model first articulated by Bernoulli (12) in 1793 and axiomatized by von Neumann and Morganstern (13) in 1945 as part of their development of game theory. Bernoulli was interested in explaining the following phenomenon. Imagine a coin-toss gamble. If the coin comes up heads, you get \$100; if it comes up tails you pay \$100. Will you play the game? Because you stand an equal chance of winning or losing, you should be indifferent to playing or not, that is, 50% of the people should play when offered the opportunity. However, the majority of people will not play the game. Also, as the range of the gamble increases (say, \$1000 won or lost), more people will choose not to play. People seem to avoid the risk of losses, and such behavior has been termed “risk-aversion” (14).

Bernoulli resolved this apparent paradox by postulating the existence of a “utility function” (U) that translates the absolute value of an amount of money (X) into the perceived value or utility of that money [$U(X)$]. The important feature of choice is not the absolute amount of money (X) lost or gained, but the utility that is lost or gained. Bernoulli suggested that the “utility of money” showed diminishing returns (a positive but decelerating function), and specifically suggested the logarithm as the utility function. Under such a characterization, the utility lost is greater than the utility gained. Because the gamble is fair, the expected utility of the game $EU(X)$ implies an expected loss in utility and, therefore, individuals should avoid the game. There are many alternative explanations for

this particular example. For my purposes, it is sufficient to note that modern decision theories start with Bernoulli's suggestion that value is best represented by some form of nonlinear translation (U) of a concrete random variable (X) and that this nonlinear translation implies a sensitivity to risk and variability in choice situations.

Choices will also be influenced by the expected or arithmetic mean payoff from a given distribution. If two distributions show the same risk, then the preference will be for whichever choice has the highest arithmetic mean payoff. Any theory of choice must contain both these aspects of preference; high expected gains will be preferred to low expected gains, low risks will be preferred to high risks. The expected utility model embraces these features of choice and explicit models attempt to partition the effects of central tendency and uncertainty as they contribute to decision-making. Formally (14, 15), the expected utility generating from choice over some random variable X will be a function of both the mean and variance in X , that is,

$$EU(X) = G(\mu_x, \sigma_x^2) \quad (1)$$

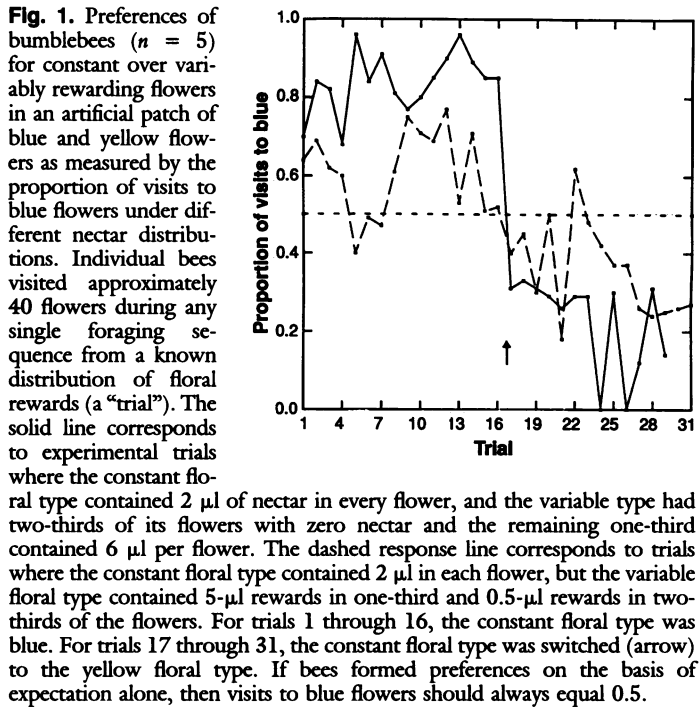
Can we use this model to describe animal choice behavior? Are the animal's choices sensitive to the mean and variance in resources? What are the relative contributions of mean and variance in determining patterns of choice? And, what is the biological and evolutionary basis of “utility” and how can it be related to the information-processing and cognitive capabilities of the individual organism?

Risk-Aversion in Bumblebees

My first experiments ascertained the individual bee's response to variability in floral reward in different floral types (11). Flowers, 100 blue and 100 yellow, were randomly spaced and interspersed throughout the artificial patch. Each blue flower contained 2 μ l of nectar, whereas the yellow flowers contained 6 μ l in one-third of the flowers and no nectar in the remaining two-thirds. Thus, blue and yellow flowers contained the same average (2 μ l), but only yellow was risky. Single bees were introduced into the enclosure and allowed to visit approximately 40 flowers (a “foraging bout”). In these experiments I worked with five marked individual bees (*Bombus pennsylvanicus*). The proportion of flowers of each color visited by the individual bee and the sequence of visits by each bee were monitored for each trial, where a trial consisted of a known distribution of rewards inserted into the enclosure followed by a single bee's sequence of visits. After each trial, the artificial patch was cleaned, and a new or replicate distribution was established. A new bee was then allowed to forage from the patch.

Individual bees favored the constant blue over the variable yellow floral type and visited an average of 84 \pm 1% blue (Fig. 1, solid line, trials 1 to 16). When the constant flower color is switched to yellow and blue is made variable (Fig. 1, trial 17), the bees switch their preferences to the now constant yellow floral type and visited an average of 23 \pm 2% blue (Fig. 1, solid line, trials 17 to 29). There appears to be little, if any, historical influence on the bumblebees' choices evidenced by the rapid switch from blue to yellow preference at trial 17 and rendering irrelevant the order of color presentation. A comparison of the sequence of first ten visits at the beginning of a foraging bout and the last ten visits shows that the bees are random foragers at the onset of each trip and decide on which flower type to visit within each foraging sequence (11).

Avoiding risk is not simply a product of aversion to nonrewarding (0 μ l) flowers. When blue was held constant (2 μ l of nectar per flower) and yellow was variable, with 5 μ l in one-third of the flowers and 0.5 μ l in the remaining two-thirds, the bees still preferred the constant over the variable floral type and visited an average of 62 \pm 1% blue (Fig. 1, dashed line, trials 1 to 16) even though all the flowers



contained some nectar. Preference for constant over variable floral type was preserved when the color for constant was switched to yellow and visited an average of $37 \pm 1\%$ blue (Fig. 1, dashed line trials 17 to 31). Risk-aversion was substantially lower in the second set of experiments (dashed line). This reduction was expected because the variance in reward was smaller for the second set of experiments.

In general, in models of risk-taking, such as Eq. 1, it is assumed that increasing uncertainty can be compensated for by increasing expectation. In the experiments described above, the expected reward per flower was constant and only variance was manipulated. Clearly, bees are risk-averse in fair gambles, but can the variability in reward be offset by increasing the average reward of the variable type? To answer this question, we kept one flower type constant (yellow, $0.5 \mu\text{l}$ of nectar per flower) and constructed the preference set for a range of means and variance in the blue floral type (16). For a given variance in reward in blue, we adjusted the mean reward in blue until bees were indifferent between color types; we continued increasing the mean reward in blue until the bees showed a preference for the variable over the constant type. The data presented are the pooled responses of eight foragers. No single individual bee was subjected to the entire range of means and variances, and no doubt pooling the data obscures some of the natural variation in individual response. However, for the analysis I am using, the variable of interest is the direction of preference, not the magnitude of preference, that is, the analysis is ordinal rather than cardinal, and individual bees show the same directional preferences. When one bee preferred yellow over blue for a given distribution of nectar, all the bees preferred yellow over blue. Consequently, interindividual variation in the magnitude of preference, although important in resolving other kinds of questions, does not enter into this study.

A plot of mean-variance combinations that generate indifferent foraging (50% constant, 50% variable) shows a significant positive and linear relation (Fig. 2, triangles). Consequently, increasing variability and uncertainty in floral reward can be compensated for by increasing the expected reward in a simple linear fashion. This experiment effectively partitions the differential contributions of mean and variance in shaping floral choice for this particular environment by measuring the trade-off between these two variables

along the indifference curve. Such a trade-off between mean and variance follows from several well-known economic models of choice—for example, the Markowitz-Tobin portfolio selection model where Eq. 1 follows:

$$EU(X) = \mu_x - a \sigma_x^2 \quad (2)$$

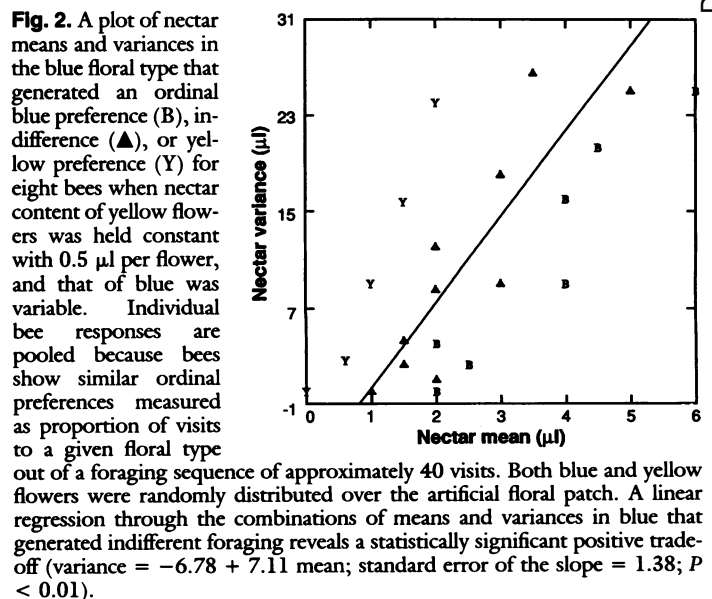
where a is the coefficient of risk-aversion and corresponds to the slope of the indifference curve (15).

Risk-sensitive foraging responses have now been suggested for a large number of animal species, including juncos (17), sparrows (18), bananaquits (19), shrews (20), wasps (11), warblers (21), rats (22), pigeons (23), as well as other species of bumblebees (24). In some cases, the trade-off between mean and variance in reward has also been established (25). Experiments on risk-sensitive responses in animals, however, have often been confounded by uncontrolled variables making critical demonstrations difficult. For example, in the frequently cited study by Caraco *et al.* (17), variance in reward is confounded with time and rate of delivery. Recent reviews examine much of the evidence for risk-sensitive responses in animals (26). Detailed studies on the biological basis of utility and the biomechanics of foraging may help resolve some of the difficulties in assessing the mechanisms that generate observed responses to resource variability.

Biomechanics of Utility in Bumblebees

The behaviors described so far were clearly consistent with the expected utility hypothesis. But unlike economics, where utility is constructed or revealed on the basis of observed behavior (27), in this biological system we can generate a functional account of utility based on biomechanical principles and derived independently from observed choice behavior.

The observed risk-aversion in bumblebees implies that the utility of reward size to a forager is a positive decelerating function of the volume of nectar obtained per flower (11, 14, 16, 26). If individual bees maximize their expected rate of net energy gain, then a nonlinear relation between nectar volume and rate of net energy uptake would account for risk-averse responses to variability in reward size. The maximization of the expected rate of net energy gain at the individual level is related to evolutionary fitness because the reproductive success of colonies is determined by the rate of



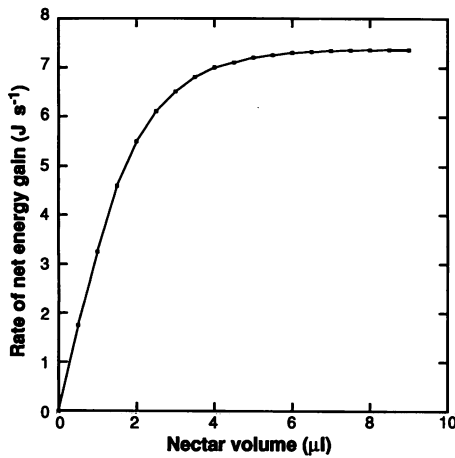


Fig. 3. A plot of the model “utility” function (Eq. 3) based on the empirically established biomechanics of nectar extraction that relates rate of net energy gain in individual bumblebees as a function of nectar volume in visited flowers. When bees maximize the expected rate of net energy gain, the nonlinear relation accounts for the bees’ sensitivity to variance in nectar rewards per flower.

energy entering the colony.

On the basis of biomechanical principles, Harder and Real (28) showed that the rate of net energy uptake (E) per flower for an individual foraging bee is represented by

$$E = \frac{e\rho SV - W(K_p(T_a + V/I) + K_f T_f)}{T_f + T_a + (V/I)} \quad (3)$$

where e is the energy content of 1 mg of sucrose (15.48 J); ρ is the nectar density (mg/ μ l); S is the nectar concentration (percentage of sucrose equivalents); V is the nectar volume (μ l); W is the bee’s mass (g); T_p is duration of visit (s); T_f is flight time between flowers (s); K_p and K_f are energetic costs of probing and flying (J/g/s); T_a is time required for entering and leaving a flower (s); and I is ingestion time (s).

For the specific physiological parameters of an average *B. pennsylvanicus* and for the nectar and flower characteristics used in the artificial patch experiments, this rate of net energy gain is a positive decelerating function of nectar volume (Fig. 3). Harder and Real (28) further showed that this physiologically based model was a fairly accurate predictor of the bumblebees’ choice behavior over the different nectar distributions used in the earlier experiments. Consequently, we have a first principles account of risk-aversion in bumblebees that does not entail a constructed utility function revealed through choice behavior.

Computation, Perception, and the Architecture of Choice

The biological account of risk-aversion that I have suggested assumes that information about floral rewards is processed in a particular fashion. Suppose a bee visits a sequence of n flowers during a foraging bout. For each flower in the sequence let R_i represent the reward and T_i the time required to get the reward from the i th flower. The Harder-Real model assumes that bees maximize the average of the set of R_i/T_i , that is,

$$\max E_1 = \text{Avg} \left\{ \frac{R_1}{T_1}, \frac{R_2}{T_2}, \dots, \frac{R_n}{T_n} \right\} \quad (4)$$

where E_1 represents the expected net energy gain from calculations based on attributes of single flowers. Such an averaging scheme amounts to calculating the expectation of the ratio of two random variables, $E(R/T)$. In a simple stochastic rate process, or technically, a reward-renewal process, such as foraging from flowers, the maximization of $E(R/T)$ leads to the maximization of short-term energy gain (29).

Several authors have argued, however, that organisms should maximize long-term gain rather than short-term gain (30). Maximizing long-term gain amounts to calculating the ratio of expectations of the random variables, that is $E(R)/E(T)$. Actually there is a range of possible computational rules that lie between completely short-term and completely long-term gain corresponding to different ways of processing information from a string of flower visits. For example, rewards and times could be pooled across successive pairs of flowers in the string, that is,

$$\max E_2 = \text{Avg} \left\{ \frac{R_1 + R_2}{T_1 + T_2}, \frac{R_3 + R_4}{T_3 + T_4}, \dots, \frac{R_{n-1} + R_n}{T_{n-1} + T_n} \right\} \quad (5)$$

or across triplets of flowers (E_3), quadruplets (E_4) of flowers, and so on. In the theoretical limit $n \rightarrow \infty$ for which $E_\infty = E(R)/E(T)$ which corresponds to the long-term average rate of energetic gain (31).

Using the physiologically based expressions for R and T from the Harder-Real model (Eq. 3), for each distribution of floral rewards used in our trade-off experiments (Fig. 2), the average rate of net energy gain E_k was calculated as a function of the frame size k over which samples are pooled. The frame size determines the fine structure required to process a particular string of information.

In the experiments, a foraging bumblebee visited approximately 35 to 40 flowers on average. Consequently, strings of 36 visits that could be partitioned according to at least nine different framing rules were modeled: single flowers, doubles, triples, quadruples, sets of 6, 9, 12, 18, and 36 flowers, respectively. Thus there are nine different ways of processing a string of 36 visits from these nine different frame sizes. For each reward distribution used in the trade-off experiments (Fig. 2), preference rank was determined. The observed floral preferences were converted to scores of 1, 0, or -1 for yellow, indifferent, or blue preferences, respectively. Then 100 simulations were run of individual foragers each sampling 36 flowers from each of the distributions for which there were preference rankings. From each set of simulations and for each of the reward distributions the nine different E_k ’s corresponding to each of the nine different computational algorithms were calculated. Then the rank correlation (Kendall’s τ) between the observed preference score (1, 0, or -1) for a given distribution of rewards and the corresponding average was calculated. The correlations were determined over all nine frame lengths. In all the experiments, the correlation between expected rate of net energetic gain and observed preference was highest for short-term optimization (that is, frame length $k = 1$ or 2) and lowest for long-term optimization (frame length $k = 36$) (31).

These results suggest that bumblebees frame their decisions on the basis of individual flowers or on pairs of flowers, which results in short-term energy maximization. This analysis does not suggest that bees remember only their last few flower visits, but that they assemble information and track floral quality on the basis of individual flowers rather than pooling information across flowers. Several investigators, however, have argued that bumblebees have only short-term memory, and recent experiments suggest that foraging decisions in bumblebees are influenced by at least the last three flower visits, but perhaps no more than that (32, 33). Are there scenarios under which calculations based on small frame lengths and short-term memory windows will prove evolutionarily advantageous?

Optimal Computational Rules

There are at least three scenarios under which calculations based on small frame lengths will prove advantageous.

First, if bees are limited in their memory capacity and therefore constrained to decisions over small sample sizes, then calculations based on $E(R/T)$ may be more accurate than calculations based on

$E(R)/E(T)$. Short-term calculations are a more robust estimator of performance when sampling is truncated. More accurate estimation will especially occur when the reward (R) and the time to acquire the reward (T) are correlated. A positive correlation between R and T is guaranteed in floral systems simply because it takes more time to ingest greater quantities of nectar. There is considerable experimental and theoretical evidence suggesting that bees are constrained to only short-term memory windows (32, 33) and severe memory constraints have been invoked to account for the evolution of floral constancy in bumblebees (34). Memory constraints may simply reflect lower neural complexity, and increased memory capacity may be correlated with the development of more complex neural systems (35).

A second scenario suggests that short-term optimization is adaptive given the particular spatial structure of floral rewards in natural systems. In most field situations there is intense local competition among pollinators for floral resources. When a pollinator finds a local patch of flowers with abundant nectar, it generally restricts its foraging to neighboring plants, ultimately exhausting the resources in that local patch (36). Such a foraging pattern will generate a high degree of spatial autocorrelation in nectar rewards, and fields are often characterized by "hot" and "cold" spots (37). If information about individual flowers is pooled, then the spatial structure of reward distributions will be lost and foraging over the entire field will be less efficient. In spatially autocorrelated environments it pays not to pool information.

The implications of correlations in habitat structure on the framing of decisions has been explored by Cuthill *et al.* (38). In a series of experiments, they demonstrated that starlings, *Sturnus vulgaris*, shape their foraging decisions (number of prey eaten in a patch) on the basis of only their most recent experience of travel times between patches. The starlings do not appear to estimate the long-term average in the environment. Cuthill *et al.* argue that when foraging parameters in the environment (for instance, travel time between patches) are highly autocorrelated then short-term memory is advantageous. The high degree of temporal autocorrelation leading to adaptive short-term memory in these birds is analogous to the high degree of spatial autocorrelation among rewards in fields that may lead to short-term memory advantages for bees.

A third explanation for why bees might assess the quality of floral types on the basis of individual flowers combines both memory constraints and the hierarchical processing of information. Since Miller's (39) publication on human memory, psychologists have recognized limits to the recall of specific information. If seven objects of a specific class (say, pieces of fruit) are presented to subjects, recall of the set of objects is fairly accurate. If many more objects are present, recall deteriorates. We obviously can recall more than seven objects, so how do we store large amounts of information in a manner that will allow for recall? Miller suggested that information is hierarchically nested, so that classes of information are pooled together. For example, if subjects were presented with seven pieces of fruit and seven hand tools, recall would be fairly efficient because fruit and tools can be used as class identifiers.

In simple nervous systems, like the bumblebee's, information on floral qualities can be hierarchically nested when rewards and times at individual flowers are pooled into simple quotients. More information may be processed and recalled if the strings are hierarchically structured, and pairing by flowers seems a reasonable hierarchical pattern for nesting the information.

The short-term memory (either as constraint or adaptation) and hierarchical processing schemes depend on some truncation of sampling or limits to recall in the nervous system. Ancillary evidence on how bumblebees form subjective estimates of probabilities support the view that bees engage in truncated sampling leading to bias in subjective probability formation.

Computation and Subjective Probability Bias

All decisions made under uncertainty demand that decision-makers estimate the probability of particular outcomes. Perhaps the most exciting and challenging ideas in contemporary decision theory pertain to the misperception of objective probabilities. Many psychologists have argued (with some experimental evidence) that human beings (and some laboratory rats) consistently and systematically misperceive probabilities (40). Low probability events seem to be overrepresented; high probability events are underrepresented. Humans tend to homogenize probabilities and weight events toward equal likelihood of occurrence. Do other animals show similar subjective probability biases, and how are biases related to information processing schema?

Eliciting probabilities from animals is difficult. Subjective probabilities must be constructed from patterns of choice over different reward options. The economic theories of "acceptance sets" and "state-preference" (41) provide a framework for this kind of analysis.

Imagine that the organism is faced with a choice in which it receives a reward X_1 if state S_1 occurs and X_2 if state S_2 occurs. We wish to determine the subjective probabilities the organism attaches to the occurrence of S_1 and S_2 . For example, nectar rewards within the class of yellow flowers might occur in two states corresponding to two reward levels with fixed probabilities p and $1 - p$. The expected utility from such a set of alternatives is simply

$$EU(X_1, X_2) = p U(X_1) + (1 - p)U(X_2) \quad (6)$$


For a fixed probability pair, we can determine the indifference set of reward states X_1 and X_2 that generate the same expected utility, that is, those values of X_1 and X_2 for which $EU(X_1, X_2) = c$, a constant. We can construct a graph whose axes are the reward levels or states X_1 and X_2 . The 45° line on this graph represents all alternatives for which $X_1 = X_2$; that is, alternative states for which the outcome of choice is certain.

Through some simple mathematical operations on Eq. 6, state-preference theory demonstrates that the slope of the indifference set at the 45° point equals the negative of the ratio of the probabilities, that is, $-p/(1 - p)$. If we experimentally construct an indifference set over reward states with fixed probabilities, then we can ascertain whether the subjective probabilities derived from the animal's choice behavior corresponds to the objective probabilities established by the investigator.

Using an enclosed colony of *Bombus bimaculatus*, I constructed such an indifference set (42). Fifty artificial yellow and 50 artificial white flowers were randomly distributed over the foraging patch. In the yellow flowers, rewards were distributed such that four-fifths of the flowers had one reward level whereas the remaining fifth had another reward level. Thus, the reward states in yellow flowers occurred with objective probabilities of 0.8 and 0.2. White flowers always contained 4 μ l of nectar. While keeping the proportion of yellow flowers with a given reward fixed, I manipulated the amounts of nectar in the two alternative states of the yellow flowers. The indifference set was constructed by titrating the reward states in yellow until bees were indifferent between foraging on the variable yellow or the constant white floral type. The data represent the pooled responses of five foragers from the colony with the majority of responses coming from two individuals. Once again, the pooled responses of individuals were used because directional preferences were the same across individuals, and the analysis requires only ordinal measures of preference.

A plot of the preferences in reward-state space reveals the indifference set (Fig. 4, triangles). Each point in the reward-state space corresponds to a combination of reward levels X_1 and X_2 with fixed probabilities 0.2 and 0.8 in yellow flowers. Combinations marked

Fig. 4. A state-preference plot of combinations of reward states in variable yellow flowers that generated either ordinal preference for yellow (Y), indifference (\blacktriangle), or the constant white floral type (W) in five individual bees. Individual bees visited approximately 40 flowers during each of the foraging sequences and showed similar ordinal preferences when foraging on identical reward distributions.



Reward states in the yellow flowers occurred with fixed probabilities [$p(X_1) = 0.2$ and $p(X_2) = 0.8$]. A linear regression through the indifference points has a slope of -8.5 ($X_2 = 16.88 - 8.5X_1$; standard error of the slope, 2.30; $P < 0.01$). The expected slope based on the objective probabilities is -4.0 . Consequently, bees appear to overestimate the likelihood of common events and underestimate the likelihood of rare events.

with a solid triangle generated indifferent foraging between yellow and white flowers. The slope of the regression through the indifference points operationally measure the subjective probabilities.

If the bees' subjective probabilities corresponded to the objective probabilities, then the slope of the regression should equal $-0.8/0.2 = -4$. In actuality, the slope of the regression is equal to -8.5 (standard error of slope 2.30, $P < 0.01$), indicating that the bees treated a probability of 0.8 more like a probability of 0.9. Similarly, a probability of 0.2 was treated more like 0.1. When an event is common they overrepresent it; when rare, it is underrepresented. This makes some intuitive sense and may simply reflect that the bees pay attention to common phenomena and ignore rare events.

The bees' probability bias indicated in our experiments tends in the opposite direction from the bias proposed for human subjects. In part, this difference may be due to differences in information-processing schemes and memory constraints. We would expect overrepresentation of common events if bees form probabilities after only a few samples from the reward distribution. In the experiments, the rewards in the yellow flowers were highly skewed—events with probabilities 0.8 and 0.2. A truncated sample from such a distribution is most likely to generate overestimation of probabilities if bees form estimates on the bases of encounter frequency. For example, if bees truncate sampling after a single flower, then there is an 80% chance that the 0.8 probability event is actually perceived as equal to 1.0. As the sampling progresses, the subjective probability will converge to the objective probability, but short-term calculations and extremely truncated sampling will lead to high bias in probability estimates.

Additional experiments on bumblebee "switching" behavior also suggest that bumblebees underperceive rare events and overperceive common events (43). When bees were restricted to foraging over an artificial patch of equal numbers of blue and yellow cardboard flowers (100 blue, 100 yellow) with constant and identical rewards (2 μl in every flower), the bees were indifferent to foraging on blue or yellow and visited them in equal proportions. However, when yellow flowers were made relatively rare (still with 2 μl in every flower), the bees disproportionately overvisited blue flowers. When blue was made relatively rare, the bees disproportionately overvisited yellow. The bees switched their preferences as a function of relative abundance of the two flower types.

Disproportionate visitation to the common floral type can be explained by simple nonlinearities in the perceptual field. The total

weight of sensory input from blue flowers may be disproportionately great when blue is relatively more abundant. Excitation from yellow in the perceptual field may be reduced when input from blue is relatively higher. Predatory switching as a function of relative abundance has been observed often, and similar arguments for the role of information processing and signal detection have been articulated (44). The perception of frequency and probability may then be recognized as a special problem in pattern recognition.

Conclusion

Bees may be viewed as short-term energy maximizers that employ computational rules that efficiently exploit the resources of their environments subject to possible neural or memory constraints, or both. Bees appear to form probabilities on the bases of frequency of encounter of different types of reward states and begin with no prior estimation of likelihoods. The observed probability bias results from short-term calculations, truncated sampling, or both. Short-term computational rules will prove adaptive when animals are subject to (i) memory or perceptual constraints, (ii) attempting to efficiently exploit resources in a spatially autocorrelated environment, or (iii) the hierarchical classification or nesting of information.

Within the bumblebees environment, misperception may actually be adaptive. The computational algorithm that seems to lead to the most efficient exploitation of resources generates, as a by-product, a bias in probability estimation. Tversky and Kahneman (45) argue that perceptual biases may result from organisms generally adopting "heuristic" rules for decision-making. However, we must recognize that what constitutes an adaptive heuristic device in one species or in one environment may be mal-adaptive for other species in other environments.

Throughout this article, I have been using economic models of decision-making developed to account for patterns of choice in humans, and it seems natural to compare the performance of bees with the performance of humans when subjected to the same kinds of tasks that employ the same kinds of methods. However, evolutionary comparisons between bees and humans is exceedingly simplistic. Humans and bees do not share a recent evolutionary history or similar ecological requirements. What I have demonstrated here is how the computational rules employed by bees adapt them for efficient exploitation of the resources in their environments. Single species arguments for adaptive specialization in mental processes must be augmented by comparisons among closely related species and species performing similar ecological tasks. Comparative studies on the adaptive nature of cognitive function in vertebrates are under way (3, 5, 10) and invertebrate models, including bees, are starting to be explored (46). Nonetheless, an evolutionary theory of cognition is still in its infancy. A truly predictive biological theory of decision-making, problem-solving, and choice behavior in animals, including humans, will explore the evolutionary and functional significance of the computational and information processing rules employed by individual organisms in their struggle to survive, acquire mates, and reproduce.

REFERENCES AND NOTES

1. H. A. Simon and C. A. Kaplan, in *Foundations of Cognitive Science*, M. I. Posner, Ed. (MIT Press, Cambridge, MA, 1989), pp. 1–48.
2. J. M. Pearce, *An Introduction to Animal Cognition* (Erlbaum, London, 1987).
3. C. R. Gallistel, *Annu. Rev. Psychol.* **40**, 155 (1989); C. R. Gallistel, *The Organization of Learning* (MIT Press, Cambridge, MA, 1990).
4. D. Cheney, R. Seyfarth, B. Smuts, *Science* **234**, 1361 (1986); D. L. Cheney and R. M. Seyfarth, *Cognition* **37**, 167 (1990).
5. H. L. Roitblat, *Behav. Brain Sci.* **5**, 353 (1982); *Introduction to Comparative Cognition* (Freeman, New York, 1987).
6. A. Newell, P. S. Rosenbloom, J. E. Laird, in *Foundations of Cognitive Science*, M. I. Posner, Ed. (MIT Press, Cambridge, MA, 1989), pp. 93–131.

7. D. E. Rummelhart and J. L. McClelland, Eds., *Parallel Distributed Processing: Explorations in the Microstructure of Cognition* (MIT Press, Cambridge, MA, 1986); D. E. Rummelhart, in *Foundations of Cognitive Science*, M. I. Posner, Ed. (MIT Press, Cambridge, MA, 1989), pp. 133-159.
8. R. J. Richards, *Darwin and the Emergence of Evolutionary Theories of Mind and Behavior* (Univ. of Chicago Press, Chicago, 1987).
9. J. E. R. Staddon, *Adaptive Behavior and Learning* (Cambridge Univ. Press, Cambridge, 1983); E. M. Macphail, *Behav. Brain Sci.* 10, 645 (1987); J. R. Anderson, *ibid.*, p. 467; F. M. Wuketits, *Evolutionary Epistemology* (SUNY Press, Albany, NY, 1990).
10. D. F. Sherry and D. L. Schacter, *Psychol. Rev.* 94, 439 (1987); B. Wilson, N. J. Mackintosh, R. A. Boakes, *Q. J. Exp. Psychol. B* 37, 313 (1985); S. J. Shettleworth, in *Advances in Analysis of Behavior*, M. Zeiler and P. Harzem, Eds. (Wiley, New York, 1983), pp. 1-39; J. R. Krebs, S. D. Healy, S. J. Shettleworth, *Anim. Behav.* 39, 1127 (1990); R. P. Balda and A. C. Kamil, *ibid.* 38, 486 (1989); P. H. Harvey and J. R. Krebs, *Science* 249, 140 (1990).
11. L. A. Real, *Ecology* 62, 20 (1981).
12. D. Bernoulli, *Commentar. Acad. Sci. Imp. Petropolitanae* 5, 175 (1738) [for an English translation, see L. Sommer, *Econometrica* 22, 23 (1954)].
13. J. von Neumann and O. Morganstern, *Theory of Games and Economic Behavior* (Princeton Univ. Press, Princeton, NJ, 1944).
14. S. French, *Decision Theory* (Halstead Press, New York, 1986); R. L. Keeney and H. Raiffa, *Decisions with Multiple Objectives* (Wiley, New York, 1976).
15. H. Markowitz, *Portfolio Selection* (Wiley, New York, 1959); J. Tobin, *Rev. Econ. Stud.* 25, 65 (1958).
16. L. A. Real, J. R. Ott, E. Siverfine, *Ecology* 63, 1617 (1982).
17. T. Caraco, S. Martindale, T. Whittam, *Anim. Behav.* 28, 820 (1980); T. Caraco, *Behav. Ecol. Sociobiol.* 8, 213 (1981).
18. T. Caraco, *Behav. Ecol. Sociobiol.* 12, 63 (1983).
19. J. M. Wunderle and T. G. O'Brien, *ibid.* 17, 371 (1986).
20. C. J. Barnard and C. A. J. Brown, *ibid.* 16, 161 (1985).
21. F. R. Moore and P. A. Simm, *Experientia* 42, 1054 (1986).
22. R. C. Battalio, J. H. Kagel, D. N. McDonald, *Am. Econ. Rev.* 75, 596 (1985).
23. R. J. Herrnstein, *J. Exp. Anal. Behav.* 7, 179 (1964); M. Davison, *ibid.* 12, 247 (1969).
24. R. V. Carter and L. M. Dill, *Behav. Ecol. Sociobiol.* 26, 121 (1990); K. D. Waddington, T. Allen, B. Heinrich, *Anim. Behav.* 29, 779 (1981).
25. T. Caraco and S. L. Lima, *Quant. Anal. Behav.* 6, 1 (1986).
26. L. A. Real and T. Caraco, *Annu. Rev. Ecol. Syst.* 17, 371 (1986); D. W. Stephens and J. R. Krebs, *Foraging Theory* (Princeton Univ. Press, Princeton, NJ, 1986); J. R. Krebs and A. Kacelnik, in *Behavioral Ecology*, J. R. Krebs and N. B. Davies, Eds. (Blackwell, Oxford, 1991).
27. J. M. Henderson and R. E. Quandt, *Microeconomic Theory: A Mathematical Approach* (McGraw-Hill, New York, 1980).
28. L. D. Harder and L. A. Real, *Ecology* 68, 1104 (1987).
29. M. J. Turelli, J. H. Gillespie, T. W. Schoener, *Am. Nat.* 119, 879 (1982).
30. J. F. Gilliam, R. F. Green, N. E. Pearson, *ibid.*, p. 875; H. P. Possingham, A. I. Houston, J. M. McNamara, *Ecology* 71, 1622 (1990).
31. L. A. Real, S. Ellner, L. D. Harder, *Ecology* 71, 1625 (1990).
32. R. Menzel, *Nature* 281, 368 (1979); R. Menzel, in *Experimental Behavioral Ecology*, B. Holdobler and M. Lindauer, Eds. (Sinauer, Sunderland, MA, 1985), p. 55; B. Heinrich, in *The Biology of Learning*, P. Marler and H. Terrace, Eds. (Springer-Verlag, Heidelberg, 1984), p. 135.
33. R. Dukas and L. A. Real, unpublished manuscript.
34. N. M. Waser, *Am. Nat.* 127, 593 (1986); R. Dukas and L. Real, unpublished manuscript.
35. W. F. Angermeier, *The Evolution of Operant Learning and Memory* (Karger, Basel, 1984); J. M. Pearce, *Introduction to Animal Cognition* (Erlbaum, London, 1987).
36. G. H. Pyke, in *The Ecology of Animal Movement*, I. R. Swingland and P. J. Greenwood, Eds. (Clarendon Press, Oxford, 1983), p. 7; D. E. Loria and F. R. Moore, *Behav. Ecol.* 1, 24 (1990).
37. P. Feinsinger, *Ecol. Monogr.* 48, 269 (1978); J. M. Pleasants and M. Zimmerman, *Oecologia* 41, 283 (1979).
38. I. C. Cuthill *et al.*, *Anim. Behav.* 40, 625 (1990).
39. G. A. Miller, *Psychol. Rev.* 63, 81 (1956).
40. D. Kahneman and A. Tversky, *Econometrica* 47, 263 (1979); U. Karmarkar, *Organ. Behav. Human Perform.* 21, 61 (1979); M. Machina, *Econometrica* 50, 277 (1982); R. C. Battalio, J. H. Kagel, D. N. MacDonald, *Am. Econ. Rev.* 75, 597 (1985); J. H. Kagel, D. N. MacDonald, R. C. Battalio, *ibid.* 80, 912 (1990); R. C. Battalio, J. H. Kagel, K. Jiranyakul, *J. Risk Uncertainty* 3, 25 (1990).
41. M. Yaari, *Q. Journ. Econ.* 79, 278 (1965); R. N. Rosssett, *Rev. Econ. Stud.* 38, 481 (1971); L. A. Real, *Am. Nat.* 130, 399 (1987).
42. L. A. Real, R. Dukas, W. Zielinski, unpublished data.
43. L. A. Real, in *Behavioral Mechanisms in Diet Selection*, R. H. Hughes, Ed. (Springer-Verlag, Heidelberg, 1989), p. 1.
44. T. Getty, *Am. Nat.* 125, 239 (1985); T. Getty and J. R. Krebs, *ibid.*, p. 39.
45. A. Tversky and D. Kahneman, *Science* 185, 1124 (1974); D. Kahneman, P. Slovic, A. Tversky, Eds., *Judgement Under Uncertainty: Heuristics and Biases* (Cambridge Univ. Press, Cambridge, UK, 1982).
46. R. Dukas and L. A. Real, *Anim. Behav.*, in press.
47. I wish to thank L. Antony and J. Levine for their many helpful discussions on aspects of animal cognition; E. Werner, D. Real, S. Ellner, and three anonymous reviewers made many helpful suggestions on the manuscript. Supported by NSF grants BSR 8500203, BNS 8719292, and BNS 9096209.

Mid-Ocean Ridges: Discontinuities, Segments and Giant Cracks

KEN C. MACDONALD, DANIEL S. SCHEIRER, SUZANNE M. CARBOTTE

Geological observations reveal that mid-ocean ridges are segmented by numerous rigid and nonrigid discontinuities. A hierarchy of segmentation, ranging from large, long-lived segments to others that are small, migratory, and transient, determines the pattern and timing of creation of new ocean floor. To the extent that spreading segments behave like giant cracks in a plate, the crack

propagation force at segment tips increases with segment length, which may explain why long segments tend to lengthen and prevail over shorter neighboring segments. Partial melting caused by decompression of the upper mantle due to plate separation and changes in the direction of spreading result in the spawning of new short segments so that a balance of long and short segments is maintained.

THE MID-OCEAN RIDGE IS THE LARGEST MOUNTAIN CHAIN and the most active system of volcanoes in the solar system. In plate tectonic theory, the ridge is located between plates of the earth's rigid outer shell that are separating at speeds of 10 to 170

mm/year. The ascent of molten rock from deep in the earth (~30 to 60 km) to fill the void between the plates creates new sea floor and a volcanically active ridge. This ridge system wraps around the globe like the seam of a baseball and is approximately 70,000 km long. However, the ridge itself is only ~5 to 30 km wide, very small compared to the plates, which can be thousands of kilometers across (Fig. 1).

The authors are with the Department of Geological Sciences and Marine Science Institute, University of California, Santa Barbara, CA 93106.