

A neural network model of foraging decisions made under predation risk

SCOTT L. COLEMAN

University of North Texas Health Science Center, Fort Worth, Texas

VINCENT R. BROWN

Hofstra University, Hempstead, New York

and

DANIEL S. LEVINE and ROGER L. MELLGREN

University of Texas, Arlington, Texas

This article develops the *cognitive-emotional forager* (CEF) model, a novel application of a neural network to dynamical processes in foraging behavior. The CEF is based on a neural network known as the *gated dipole*, introduced by Grossberg, which is capable of representing short-term affective reactions in a manner similar to Solomon and Corbit's (1974) opponent process theory. The model incorporates a trade-off between approach toward food and avoidance of predation under varying levels of motivation induced by hunger. The results of simulations in a simple patch selection paradigm, using a lifetime fitness criterion for comparison, indicate that the CEF model is capable of nearly optimal foraging and outperforms a run-of-luck rule-of-thumb model. Models such as the one presented here can illuminate the underlying cognitive and motivational components of animal decision making.

Foraging behavior has been characterized as an inherently dynamic process (Mangel & Clark, 1988; McFarland, 1977; McNamara & Houston, 1986; Stephens & Charnov, 1982). In order to survive, actively feeding animals must make a number of important foraging decisions on any given day. For example, an animal must decide where to search for food, what foods to consume, and how long to forage before moving to a new location. Assessing the costs and benefits of foragers' decisions has led foraging modelers to borrow theoretical concepts from economics and decision theory, such as utility, risk sensitivity, and average rate maximization (see Stephens & Krebs, 1986, for a review). In addition, experimental psychologists have shown that decision making under risk is heavily influenced by short-term emotional reactions (Townsend & Busemeyer, 1989; Tversky & Kahneman, 1974, 1981).

Animal foraging theory, like human decision-making theory, has moved back and forth between normative and descriptive models. The normative models are usually some variant of optimal foraging theory (OFT), which assumes that the principles of evolutionary biology apply to animal decision making in a foraging context (Krebs, 1978; McFarland, 1977). Animals that forage efficiently

are favored by natural selection and thus should have evolved as optimal (or nearly optimal) decision makers. As a normative model, OFT prescribes how an animal should behave in order to maximize a particular currency (e.g., net energy intake or probability of survival), which is then interpreted in terms of a forager's lifetime fitness. Although OFT does not explain the proximate mechanisms that underlie actual behavior, it gives insights into some decision-making processes that might be used by animals.

Yet an increasing number of scientists recognize that both animal and human decision making is heavily influenced by short-term contexts and may therefore deviate from maximization of lifetime fitness. For example, Reboreda and Kacelnik (1991) and Marsh and Kacelnik (2002) have noted that foraging starlings are sensitive to risk and not just to expected value in choices among both amounts and delays of food. Like Tversky and Kahneman's (1981) human subjects deciding between monetary gambles, these starlings tend to be risk seeking for choices among losses (and, sometimes, risk averse for gains). This suggests that a foraging theory is required that incorporates short-term emotional influences and yet predicts adequate, if not optimal, fitness.

How can such emotional influences on decision making be understood quantitatively? There have been a few models of foraging decisions using neural networks (Mangel, 1990; Niv, Joel, Meilijson, & Ruppin, 2002a, 2002b; Real, 1992). However, none of those models incorporates biologically realistic principles for dynamic affective evaluation in real time. One such principle is opponent

We thank Stuart Humphries, Graeme Ruxton, Tom Langen, Raymond Jackson, Martha Mann, and Donald Hantula for their valuable comments on earlier versions of our manuscript. Correspondence should be addressed to V. R. Brown, Department of Psychology, 135 Hofstra University, Hempstead, NY 11549 (e-mail: psyvrb@hofstra.edu).

processing, based on comparison of current stimuli with ongoing expectations (Solomon & Corbit, 1974). Opponent processing has played a major role in neural models of cognitive–emotional interactions in human decision making (Grossberg & Gutowski, 1987; Leven & Levine, 1996). Here we introduce a neural network model based on opponent processing that models animal foraging decisions made under risk of predation.

Some Data on Real-Time Foraging Decisions

The choices made by foragers often depend on the motivational and physiological state of the animal at the time of the decision, and a decision made in the present will therefore affect future decision making (McFarland, 1982). For example, Caraco, Martindale, and Whittam (1980) documented the effect of energy level on risk-taking behavior in foraging birds. Hungry birds having a negative energy budget (i.e., energy levels below a survival threshold) preferred a variable (risk-prone) return to the constant (risk-averse) return preferred by birds with a positive energy budget (Caraco & Lima, 1987). Thus, in risk-sensitive foraging, the current level of hunger affected the forager's feeding strategy.

Hunger is known to influence feeding behavior in many ways. Some effects associated with increased hunger are reduced selectiveness in diet (Schoener, 1971), increased impulsiveness (Snyderman, 1983), risk proneness (Caraco et al., 1980), reduced neophobia and attenuated defensive reactions (Shevelkin, 1994), enhanced conditioned food preference (Rozin & Schulkin, 1990), increased activity, and decreased postmeal satiety (Le Magnen, 1992). Hunger may also influence subjective values of food tastes (Cabanac, 1985).

Foraging under risk may refer to either of two quite different situations. One type of risk is energy risk in a situation in which resources are not evenly spread through a habitat. Under risk of starvation, the forager must choose between (1) some patches that have a large payoff but low probability of payoff and (2) other patches that have a small payoff but high probability of payoff. The other type of risk is the risk of the forager itself becoming prey. Foraging decisions made under risk of predation may differ from decisions based on energetic considerations alone (see Lima & Dill, 1990, for a review); for example, predation risk can influence the choice of habitat or patch (Caldwell, 1986; Dill, 1987; Gilliam & Fraser, 1987). When a forager is faced with a condition in which it is not possible to freely pursue a desired resource because of the risk of predation, then foraging becomes a compromise between food search behavior and predator avoidance behavior. Thus, the trade-off between foraging and predation risk precludes a pure energy maximization strategy (Goss-Custard, 1977; Milinski & Heller, 1978). Lima, Valone, and Caraco (1985) showed that exposure to potential predation resulted in cautious food-seeking behavior. Rather than maximizing energy gains, predation pressure forced squirrels to return to cover after brief bouts of foraging.

In general, to maximize survival, animals tend to forage in safe environments until resources are depleted (Cowlshaw, 1997). On the other hand, foragers tolerate a high risk of predation in patches with high gain. However, optimal foraging models predict that increases in energy states should decrease the time allocation and vigilance of risky patch use (Brown, 1999; Houston, McNamara, & Hutchinson, 1993). Contrary to this prediction, Gilliam and Fraser (1987) demonstrated that minnows (*Semotilus atromacalatus*) tended to accept risk of predation when feeding in a high-quality microhabitat, thus behaving less than optimally if foraging were determined solely by a gain maximization rule.

The neural network model we introduce incorporates both starvation and predation risks. It includes short-term trade-offs between risk and desirability of food.

The Patch Selection Paradigm

The simulations implemented here use the simple patch selection model under both risk of predation and risk of starvation (reviewed by Mangel & Clark, 1988). It is assumed that the animal is in the patch utilization phase of foraging—that is, the properties of the patches are known to the animal. The forager is confronted with different options (patches) that have different expected gains in terms of energy and different expected losses in terms of predation risk. We simulated a forager that encounters two patches that vary along two dimensions. A patch may yield either a low-valued or a high-valued food item and pose either a low or a high danger of predation for the forager.

For each decision the forager makes, it must decide how to avoid both starvation and predation. Each patch is characterized by a set of parameters (see Table 1): probability of death by predation (β_i), probability of finding food (λ_i), and the energetic value of food (Y_i). The change in the state variable $X(t)$ (i.e., energy reserves) depends on the outcome of a patch visit. Over the time span of the simulation, the patch parameters are assumed to remain constant.

In this work, we compared our model in foraging simulations with several other specific models (optimal, rule of thumb, and random forager). For purposes of comparison, the parameter values of the less risky patch (Patch 1) were held constant, and those of the more risky patch (Patch 2) were systematically varied. A foraging simulation run proceeds as follows. The model forager's energy reserves are set to an initial value (50 in the simulations reported here). An initial patch is selected on the basis of the initial energy level using the decision rule for the model being simulated. A random process then determines whether the forager survives predation with probability $1 - \beta_i$. If the forager does not succumb to predation, another random process determines whether food is obtained in the current patch during the current time cycle with probability λ_i . If food is obtained, the forager's energy reserves are increased by the value of the food in the current patch, Y_i . A forager's energy reserves

Table 1
Parameters and Initial Values Used in the Patch Selection Simulations

Symbol	Value	Description
All Models		
T	20	Number of time cycles
t		Current time period $t < T$
Y_1	10	Value of energy gain in Patch 1
Y_2	18	Value of energy gain in Patch 2
α	10	Energy expenditure per cycle
λ_1	.4	Probability of finding food in Patch 1
λ_2	.1–.9	Probability of finding food in Patch 2
β_1	.005	Probability of predation in Patch 1
β_2	.005–.07	Probability of predation in Patch 2
E_{\min}	0	Minimum energy reserve (starvation occurs)
E_{\max}	100	Maximum energy reserve
E_{start}	50	Initial value of energy reserve
Optimal and Energy Threshold Models		
$F(x, t, T)$		Function that maximizes lifetime fitness, at time t , given current energy reserve x
$X(t)$		State variable X at time t (equal to E_{start} at $t = 0$)
$I^*(x, t)$		Patch in which $F(x, t, T)$ occurs
CEF Model		
I	Equation 4	Tonic input assumed to be related to hunger level or inversely to energy reserves
J_1^+	4.0	Affectively meaningful input for anticipated gains in Patch i
J_2^+	5.0–15.0	
J_1^-	0.5	Affectively meaningful input for anticipated losses in Patch i
J_2^-	1.0–5.0	
ψ	98	Value of life (single-attribute model)
ψ_β	70	Value of life loss (multiattribute model)
$\psi_{(1-\beta)}$	10	Value of life gain (multiattribute model)
Δ	Equation 1	Decision criterion; see Equation 1
ROL Model		
m	1–8	Number of successive successful time periods
n	1–5	Number of successive unsuccessful time periods

Note—Number of patches in the basic patch selection paradigm was two, with Patch 1 parameters being fixed and Patch 2 parameters varying. Parameters with two values represent Patch 1 and Patch 2, respectively.

are capacity limited (i.e., they are bounded by an upper limit; 100 in the present simulations). If food is not obtained, the forager's energy reserves are decreased by subtracting α , the cost of foraging. If the forager's energy reserves fall to zero, the forager dies. If the forager survives the time cycle, the forager now decides whether to remain in the current patch or to switch to the other one. How the decision is made is determined by the model being simulated. The length of a foraging cycle could vary according to environmental conditions or life history of the forager being simulated; for example, the simulated foraging cycle could correspond to a single day in the case of a small bird in winter (McNamara & Houston, 1986). The details of each model's decision-making process are given next.

Stochastic Dynamic Programming (SDP)

Although the normative and descriptive models of foraging described in this study utilize different quantitative techniques, direct comparisons between them (and other models) are possible using the common currency of lifetime fitness. Such comparisons may lead to a more complete understanding of the nature of foraging behavior.

Sibly and McFarland (1974) introduced the state-space approach to the study of motivation in animal behavior.

In this approach, a state variable (standing for the motivational and physiological state of an animal) is represented in a vector space (see McFarland & Houston, 1981, for a review). The advantage of the state-space system is that it provides a framework for the investigation of physiological and psychological factors that underlie behavior and emphasizes the ways the state is changed by behavior over time. Theoretically, decisions proceed by a form of cost/benefit analysis based on the utility of various actions available to the forager. Costs and benefits are affected by internal factors, such as the state of food deprivation, and external factors, such as the risk of predation. Knowledge of both internal and external factors is necessary for determining the optimal trajectory that minimizes costs and maximizes benefits over time (i.e., the fitness of a behavioral sequence; Sibly & McFarland, 1976).

In the 1980s, behavioral ecologists adopted stochastic dynamic programming (SDP) as a technique for the study of trade-offs between various courses of action within an animal's repertoire and used it to model foraging decisions based on the motivational state (energy reserves) of the forager over a fixed time period (Houston, Clark, McNamara, & Mangel, 1988). The technique finds the optimal decision trajectory through a recursive process

that maximizes a survival function. Computer simulation of optimal foraging (OF) decisions uses a state variable $X(t)$ to characterize the state of a forager at time t . The “state” of the forager is usually defined as its energy reserves at that time, but it is also assumed to be a motivational and physiological state, such as hunger.

The success of a forager is measured by the lifetime fitness function $F(x, t, T)$, which represents the maximum probability of a forager with energy reserve X surviving from time t until time T . That is, $F(x, t, T)$ is the probability that the forager survives from time t until time T , assuming that it makes the optimal patch selection $i^*(x, t)$ at each time interval between t and T . The factors determining the probability of survival over a time interval $F(x, t, T)$ are the probability of not being preyed upon ($1 - \beta_i$) and the probability of not starving [λ_i , when $X(t) \leq \alpha$, where α is the energy expenditure per cycle]. The optimal patch trajectory is found by first choosing the patch that maximizes the probability of surviving until a terminal time T , assuming that the forager has already survived until time $T - 1$ and currently has energy reserves X . Next, the patch that maximizes survival from $T - 2$ until T is chosen, and so on. This backward time optimization algorithm is computationally simpler than searching n^T possible decision paths in a forward procedure (Bellman, 1957). The backward time algorithm produces a table listing the best choice patch $i^*(x, t)$ at time t given energy reserves $X(t)$. Because the choice is time dependent, the same energy reserves X can lead to different patch selections at different times. Generally, when energy reserves are sufficiently larger than the cost α (i.e., the forager is not currently in danger of starvation), the optimal choice becomes more conservative as t approaches T ; therefore, the optimal model tends to choose a less risky but less nutritious patch over a more risky but more nutritious patch near the end of a simulated foraging cycle.

Energy Threshold Model

The optimal model makes three important assumptions about the knowledge and decision-making strategies of the forager: (1) The forager has and uses knowledge of the patch parameters (probability of death by predation, probability of obtaining food, and quality of food in a given patch), (2) the forager’s decisions depend on the current energy reserves (hunger level) of the forager, and (3) the forager’s decisions are time dependent (i.e., the forager is assumed to have knowledge of the foraging time cycle). The SDP technique, albeit capable of determining the optimal decision trajectory, is unlikely to accurately represent actual proximate mechanisms involved in foraging decisions. Nevertheless, because the optimal strategy represents the upper bound on performance, it can serve as a standard by which to compare other decision-making algorithms.

Assumption 3, that the forager knows near the beginning of its life or foraging cycle when its life or foraging cycle will end, may not always be valid. By eliminating Assumption 3, we obtain an energy threshold (ET) model

under which the forager attempts to maximize lifetime fitness as a function of current energy reserves, but the energy level at which the forager chooses one patch over another does not change with time. That is, the forager is assumed to adopt a constant ET above which it chooses a safer but lower return patch and below which it chooses a more dangerous but higher return patch.

The Gated Dipole Neural Network

Neural networks can provide a theoretical methodology for studying how animals make decisions under realistic constraints and under the influence of both short-term and long-term motivations. Hence, we will employ a neural network model to represent the cognitive and motivational processes that may represent an animal’s actual foraging decisions. We will show that such a model, which utilizes a time-independent energy threshold, can be made to perform at a level that does not quite match the performance of an optimal forager but comes close.

Our neural network model of foraging, called the *cognitive-emotional forager* (CEF) model, is adapted from the gated dipole, a neural network originally developed to model emotional antagonistic rebounds (Grossberg, 1972a, 1972b). That is, the reduction of an existing stimulus or the absence of an expected stimulus that has a particular affective value leads to an experience of the opposite emotion; for example, cessation of a painful stimulus, such as electrical shock, is experienced as relief (positive affect) and absence of an expected pleasurable stimulus, such as food, is experienced as frustration (negative affect). A typical example of such a rebound is shown in Figure 1, in which the “on-channel activity” is the primary or initial affective response and the “off-channel activity” is the rebound. These effects are similar to the mechanisms assumed by the theory of motivation known as *opponent processing* (Solomon & Corbit, 1974). We first describe the gated dipole in general, then its adaptation to human decision making under risk by Grossberg and Gutowski (1987) and our own adaptation of the Grossberg-Gutowski model to the CEF model of foraging decisions.

The neural network mechanism that Grossberg proposed involved pairs of antagonistic channels with habituating chemical transmitters that bias (“gate”) responses to channel inputs, as shown in Figure 2. In Figure 2, the left side (x_1 to x_3 to x_5) is called the “on” channel and responds to an affectively significant input J (e.g., electric shock). The right side of the network (x_2 to x_4 to x_6) is called the “off” channel and represents affect of the opposite valence. Both on and off channels also receive a constant tonic input, I , that is interpreted as nonspecific arousal.

The synapses between x_1 and x_3 and between x_2 and x_4 , marked with completely or partially filled squares, have a chemical transmitter that tends to be depleted with activity of the node at the previous stage (x_1 or x_2) and to recover when that node is inactive. They are called *gates* because they let in a fraction of the activity of the

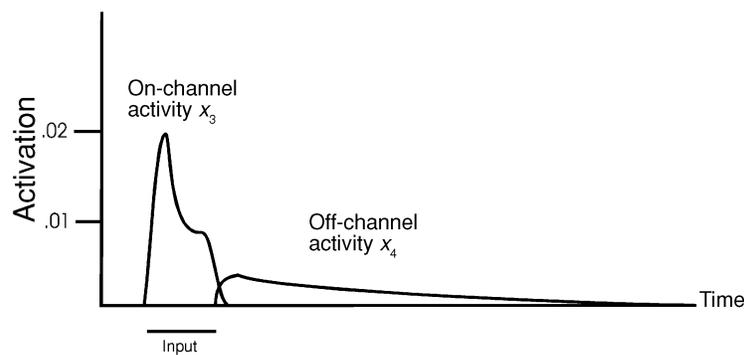


Figure 1. The time course of on-channel and off-channel activity in the gated dipole neural network. From “Modeling Some Effects of Frontal Lobe Damage: Novelty and Perseveration,” by D. S. Levine and P. S. Prueitt, 1989, *Neural Networks*, 2, p. 109. Copyright 1989 by Elsevier Science. Adapted with permission.

previous node equal to the amount of undepleted transmitter. These gates achieve the off-channel rebound of Figure 1 in the following manner. While J is on, the “on” channel receives more input than the “off” channel; hence, transmitter is more depleted on the “on” side. But the greater input overcomes the more depleted transmitter, so on-channel activity x_3 exceeds off-channel activity x_4 . This leads, by feedforward competition (i.e., inhibition) between channels, to net positive activity from the on-channel output node x_5 . For a short time after shock is removed, both channels receive equal inputs I , but the right channel is less depleted of transmitter than is the left channel. Hence, off-channel activity x_4 now exceeds x_3 , until the depleted transmitter recovers. Now competition between channels leads to net positive activity from the right-channel output node x_6 . If the significant input J is affectively negative (e.g., shock), off-channel activity is interpreted as relief and can positively reinforce a motor response that leads to such activity. If J is affectively positive (e.g., food), off-channel activity is interpreted as frustration and can suppress motor responses.

Affective Balance and the Cognitive–Emotional Forager (CEF) Model

Grossberg and Gutowski (1987) developed *affective balance theory*, a theory that slightly modifies the gated dipole to include simultaneous inputs of positive and negative affective valence to both channels. These inputs can represent amounts and/or probabilities of monetary gain or losses, which underlie decisions between risky alternatives. Specifically, affective balance theory accounts for the dependence of human preferences on expectations and on the framing of choices (e.g., Tversky & Kahneman, 1974, 1981).

Figure 3 shows the CEF model’s gated dipole, which is based on a version of Grossberg and Gutowski’s. The gated dipole is activated by two types of input. First, a nonspecific arousal input I activates both channels tonically. In the version of this model that we develop for foraging behavior, this tonic input is proportional to the forager’s

hunger level—that is, inversely related to the forager’s energy reserves. Second, there are specific affectively meaningful inputs to each channel (J^+ and J^-) that form the basis of decision making. Furthermore, it is assumed that the net affective response at each channel is the product of the short-term memory of the stimulus event and the long-term memory trace of the event’s expected outcome. The expected long-term event outcome is encoded through transmitter gating at connections (“synapses”) of the neural network.

In the case of foraging, for our first version of the model, J^+ encodes the presence and desirability of food, whereas J^- denotes the danger and likelihood of being the object of predation. Thus, the on-channel response is sensitive to expected gains in the form of energy reserves, and the off-channel response is sensitive to expected losses in the form of injury or death encountered in a choice situation. In a more elaborate version, there are two separate gated dipoles denoting food and predation, and the patch decision is made by summing the outputs of both dipoles in response to each patch.

In the CEF model, the energy threshold that determines patch selection is the result of a psychophysiological process that compares potential losses and gains (J^- and J^+ , respectively). This mechanism is compatible with several previous theories that have suggested a relationship among motivation, emotion, and cognitive processes that control the basic survival systems of feeding and avoidance behavior (e.g., Buck, 1985; Gray, 1977, 1991; Panksepp, 1982). Motivation, affect, and cognition are considered to be part of an integrated, adaptive, goal-directed appetitive/avoidance mechanism. Motivation leads to arousal and the general activation of behavior; in addition, it serves to maintain internal homeostasis. Thus, motivation can be considered to be an internal source of activation underlying foraging behavior (this is represented by the tonic input I in the model). Affect refers to a motivational state elicited by specific stimuli and not necessarily to an emotional state or a feeling (the affective states elicited by the different foraging patches are represented by the J^+ and J^- in the model); Berridge and

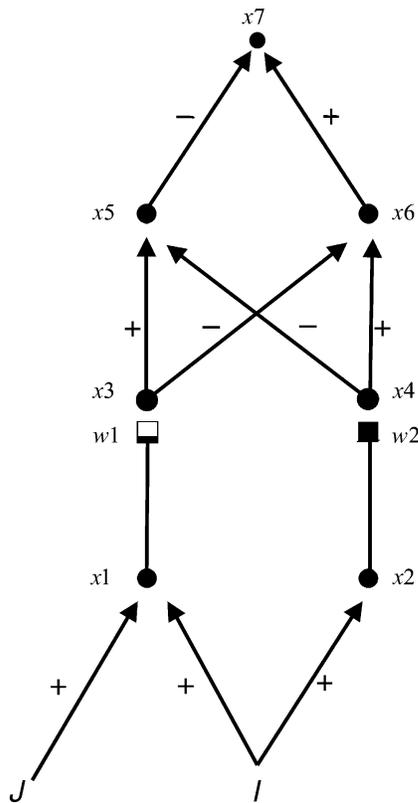


Figure 2. Schematic gated dipole network. J is a significant input (in the example of Grossberg, 1972b, electric shock), whereas I is nonspecific arousal. Synapses w_1 and w_2 can undergo depletion (as w_1 has in this diagram), as indicated by partial lightening of boxes. After J is shut off, $w_1 < w_2$ (transiently), so $x_3 < x_4$. By competition, x_6 is activated, enhancing a motor output suppressed by J . From *Introduction to Neural and Cognitive Modeling* (2nd ed., p. 63), by D. S. Levine, 2000, Mahwah, NJ: Erlbaum. Copyright 2000 by Lawrence Erlbaum Associates. Adapted with permission.

Robinson (1998) made a similar distinction between wanting (incentive motivation toward an object) and liking (pleasure derived from the object). Panksepp (1982) proposed the existence of several neurochemical systems within the limbic system—such as a foraging expectancy system and a fear avoidance system—that promote neural processing of biologically important stimuli. Thus, affect may be considered as the activation of separate neural systems that regulate basic survival responses.

The decision mechanism in the CEF model operates by means of an algorithm that Grossberg and Gutowski (1987) called *affective balance*, which combines the J^+ and J^- elicited by each patch and the internal motivational state represented by the parameter I . The decision maker, in our case the forager, chooses between two alternatives by attending to both positive and negative aspects of a potential food source, or “patch,” and then forming a net affective response (r_1) to positive and negative aspects of a patch, under the motivation level I . Then, the decision maker shifts attention to positive and

negative aspects of the second patch and forms a net affective response (r_2) to that alternative; r_2 depends nonlinearly not only on I and the J^+ and J^- of the second patch but also on the the J^+ and J^- of the first patch with which it is being compared. Letting $\Delta = r_1 - r_2$, the first alternative is chosen if $\Delta > 0$, and the second alternative is chosen if $\Delta < 0$. The actual formula for Δ is derived in Grossberg and Gutowski (1987) on the basis of steady-state values for a dynamical system of nonlinear differential equations arising from the comparison process. If the subscripts 1 and 2 attached to J^+ and J^- denote the positive and negative affective inputs associated with Patch 1 and Patch 2, respectively, this formula is

$$\Delta = G \cdot \left[(J_1^+ - J_1^-) - (J_2^+ - J_2^-) \right] + C \cdot \left[(J_1^+ \cdot J_2^-) - (J_1^- \cdot J_2^+) \right], \quad (1)$$

where

$$G = (A + I) \cdot C \quad (2)$$

and

$$C = \frac{(A \cdot B)}{(A + I + J_1^+) \cdot (A + I + J_1^-)}, \quad (3)$$

and A and B are positive constants.¹

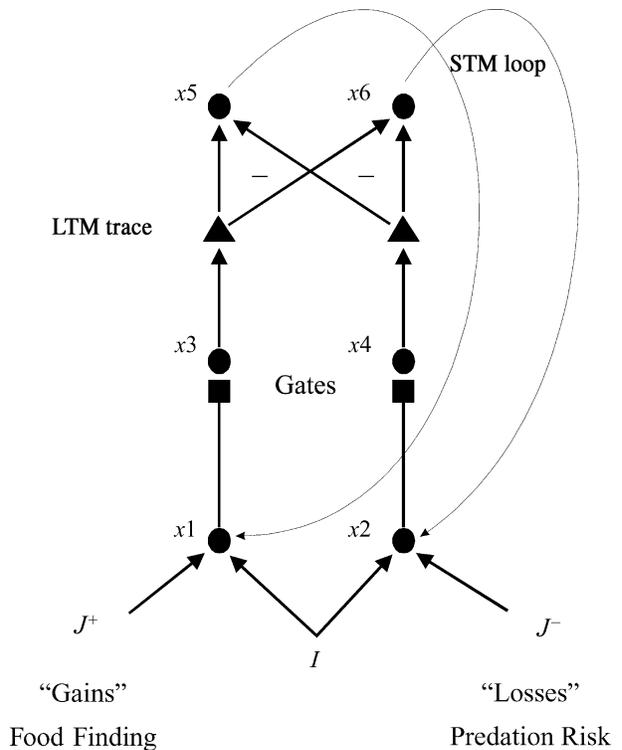


Figure 3. Schematic of the cognitive-emotional forager model. Cognitive and emotional interactions of the gated dipole are activated by affect-specific drive inputs (J^+ and J^-) and nonspecific arousal input (I). The net affective activity is the product of the activation of the short-term memory for an event gated by the affective value of the event in long-term memory.

We chose the gated dipole model as the basis of our foraging decision model for several reasons. First, the dipole has successfully modeled aspects of risky and contextually influenced decision making in humans, including framing effects (Grossberg & Gutowski, 1987) and multiattribute consumer preferences (Leven & Levine, 1996). This is because the affective value of an alternative depends nonlinearly both on the alternative chosen and on the context (other alternative or current state) to which it is compared. Second, several decision researchers (e.g., Shafir, 2000; Weber, Shafir, & Blais, 2004) have proposed that the magnitude of affective value of reward (sometimes called the *utility of the reward*) can be a Weber law function (analogous to a just noticeable difference in perception) of either amount or probability of reward, whereby the affective utility of an increase or decrease in reward depends on the initial amount of reward and not just on the amount of change. Such Weber law functions can arise from the dipole Equations 1–3. Finally, the dipole accounts for these decision data by means of a network that includes plausible analogs of brain processes, as discussed in the next section.

The expected return of a patch (J_i^+), representing potential gains, is the product of the probability of finding food (λ_i) and the value of the food item (Y_i). The product of the probability of predation (β_i) and a parameter called the *value of life* (ψ) represents the anticipated losses (J_i^-) of a patch. The tonic arousal level I is assumed to depend on the forager's current energy level $X(t)$. For the purposes of simulation, we need to specify the exact relationship between X and I . I is assumed to increase as X decreases from its maximum value. Thus, I is a function of $E_{\max} - X$, where E_{\max} is the upper bound on the forager's energy reserves. More generally, we assume that

$$I = \left[\frac{(E_{\max} - X)}{L} \right]^K, \quad (4)$$

where K and L are positive constants that are optimized as described in the Comparison of Models section.

One distinction between the CEF model and OF is the CEF model's inclusion of a parameter for the value of life (ψ), corresponding to the possible losses encountered during foraging as a result of a trade-off between reproductive success and predation risk. In the case of the CEF model, the parameter for the value of life (i.e., anticipated losses) is patch independent, although it is distinct from the generic cost of foraging α , which is a loss of energy reserves subtracted at the time of patch choice. In our model, the value of life is considered to be in part a measure of the forager's future reproductive success; for example, animals are known to avoid dangerous places to forage and take fewer risks if they are either young or reproducing (see Kavaliers & Choleris, 2001, for a review). Thus, it may be representative of individual differences in risk proneness between animals, which is not taken into account by OFT.

Figure 4 shows the decision criterion Δ plotted as a function of tonic arousal level (I) for three different sets of values of J^+ and J^- in a two-patch simulation. Patch selection is determined by the sign of Δ . Positive values

result in the selection of Patch 1, and negative values result in the selection of Patch 2. A preference reversal between patches as a function of I occurs when Δ changes from positive to negative, or vice versa. In Figure 4, as I is increased (i.e., increased arousal and higher motivation), the forager shifts from Patch 1 to Patch 2, represented by the solid line. If values for Δ remain either entirely positive or entirely negative in a simulation (dashed lines in Figure 4), then patch choice is constant for that set of patch parameters.

Table 2 depicts the choice between patches, designated as either Patch 1 or Patch 2 for a series of simulations of the CEF model. By Equation 3, the value of I at which $\Delta = 0$ is

$$I_0 = -A + \frac{(J_1^+ \cdot J_2^-) - (J_1^- \cdot J_2^+)}{(J_1^- - J_1^+) - (J_2^- - J_2^+)}, \quad (5)$$

where A is a positive constant ($A = 1.0$ in the present simulations). Equation 5 represents the point of switching from Patch 1 to Patch 2 (i.e., a preference reversal) during foraging. In Table 2, Patch 1's affectively meaningful inputs (J_1^+ , J_1^-) are held constant, and Patch 2's affectively meaningful inputs (J_2^+ , J_2^-) are systematically increased. When predation in Patch 2 (J_2^-) is low, Patch 2 is preferred (i.e., Δ is negative) when the expected food gain J_2^+ is greater than the expected food gain J_1^+ . However, the CEF model tends to select Patch 1 over Patch 2 (i.e., Δ is positive) when the anticipated predation J_2^- is greater than J_1^- .²

Figure 5 shows the changes in the value of the input I over a fixed time period in a CEF simulation. The simulation involves patch choice between a profitable patch, with relatively higher predation risk designated as Patch 2 and a less profitable but safe patch designated as Patch 1. In this particular run, the forager successfully avoided predators and starvation to survive until the end of the simulation ($T = 20$). The value of input I_0 is the point that demarcates a preference reversal between the alternative patches and is a constant threshold for patch selection over time. Changes in input I occur as a consequence of the probabilistic outcomes encountered while foraging and form a decision trajectory over the time period as values of I cross I_0 . The rate of change in I over time is greater for food discoveries in the more profitable Patch 2 than in Patch 1, which has a less valuable return. Failure to find food in either patch has the same effect on I because the cost of foraging α is a constant and not patch specific in this simulation.

The type of neural network methodology we employ is based on analyzing psychological processes into component subprocesses. The nodes or modules in the network that represent subprocesses typically are assumed to correspond not to single neurons but to a large number of neurons or to brain regions or to both (see Levine, 2000, for a review).

At this stage of our modeling, we do not venture to assign specific brain areas to our network structures. How-

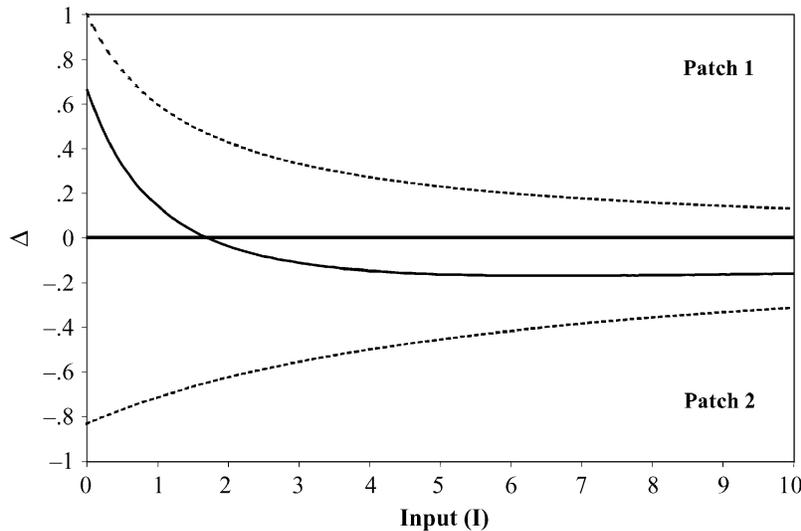


Figure 4. The decision criterion (Δ) plotted against increasing levels of I for three different choice situations. The solid line represents a preference reversal occurring as a function of I . The dashed lines represent constant patch choices. Parameter values for the switch are $J_1^+ = 5.0, J_1^- = 0.5; J_2^+ = 11.0, J_2^- = 3.0$ (i.e., Patch 2 has a higher return but is more dangerous than Patch 1); for the constant Patch 1 choice, $J_1^+ = 5.0, J_1^- = 0.5; J_2^+ = 5.0, J_2^- = 2.0$ (i.e., Patch 2 is more dangerous than Patch 1 with the same return); for the constant Patch 2 choice, $J_1^+ = 5.0, J_1^- = 0.5; J_2^+ = 10.0, J_2^- = 0.5$ (i.e., Patch 2 has a higher return than Patch 1 with the same danger level). Patch 1 is chosen when $\Delta > 0$, and Patch 2 is chosen when $\Delta < 0$.

ever, the neural structures found in the limbic system have been identified as important components in integrating and regulating neural processing of emotion, learning, memory, and cognition (Gray, 1991). The history of neural network modeling has included numerous cases of network models being originally suggested mainly on psychological grounds and then being partially mapped at a later time to specific brain regions as confirmatory physiological data emerge. An example is Grossberg and Merrill's (1996) assignments of various subparts of pre-

vious networks modeling eyeblink conditioning to regions of the hippocampus and cerebellum.

Rule-of-Thumb (ROT) Models

Janetos and Cole (1981) suggested that animals need not decide optimally to survive; rather, animals may employ simple tactics when confronted with complex decision making. Generally, good *rule-of-thumb* (ROT) models tend to approximate optimal strategies without requiring unrealistic cognitive ability or complete knowledge of the environment by the forager. Rules of thumb do not usually depend on the animal knowing detailed information about patch parameters, nor do they depend on the current motivational state of the animal. For example, a forager may efficiently select prey items on the basis of size alone (Werner & Hall, 1974) or may leave a patch after a number, n , of prey items are consumed (Iwasa, Higashi, & Yamamura, 1981). Although not strictly optimal, ROT models may be efficient under some circumstances and may actually have a selective advantage over mechanisms that yield optimal results when one includes cognitive effort as a consideration (Levine & Elsberry, 1997). ROT models are predicted to perform close to the optimal strategy when predation risk is unknown or overestimated; however, the relative advantage of an ROT is diminished when learning and energetic costs have low or negligible impact on mortality (Bouskila & Blumstein, 1992; Welton, McNamara, & Houston, 2003).

The ROT implemented here is an adaptation of the *run-of-luck* (ROL) rule, which has been shown to perform under some circumstances in a close to optimal fashion

Table 2
CEF Model Patch Selection and Input Threshold Values I_0 for Preference Reversals in a Two-Patch Choice Simulation

J_2^+	J_2^-									
	0.5	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
5	—	P1	P1	P1	P1	P1	P1	P1	P1	P1
6	P2	3.0	P1	P1	P1	P1	P1	P1	P1	P1
7	P2	P2	3.0	12.0	P1	P1	P1	P1	P1	P1
8	P2	P2	0.8	3.0	7.5	21.0	P1	P1	P1	P1
9	P2	P2	P2	1.2	3.0	6.0	12.0	30.0	P1	P1
10	P2	P2	P2	0.4	1.5	3.0	5.3	9.0	16.5	39.0
11	P2	P2	P2	P2	0.8	1.7	3.0	4.8	7.5	12.0
12	P2	P2	P2	P2	0.3	1.0	1.9	3.0	4.5	6.6
13	P2	P2	P2	P2	P2	0.6	1.2	2.0	3.0	4.3
14	P2	P2	P2	P2	P2	0.2	0.8	1.4	2.1	3.0
15	P2	P2	P2	P2	P2	P2	0.4	0.9	1.5	2.2

Note—P1 and P2 denote Patch 1 and Patch 2, respectively. The dash indicates where $\Delta = 0$. The Patch 1 parameters are held constant at $J_1^- = 0.5, J_1^+ = 5.0$. The Patch 2 parameters are varied at $J_2^- = 0.5-5.0, J_2^+ = 5.0-15.0$. A P1 or P2 designation indicates an invariant patch choice. For noninvariant choices, Patch 2 is chosen when $I > I_0$.

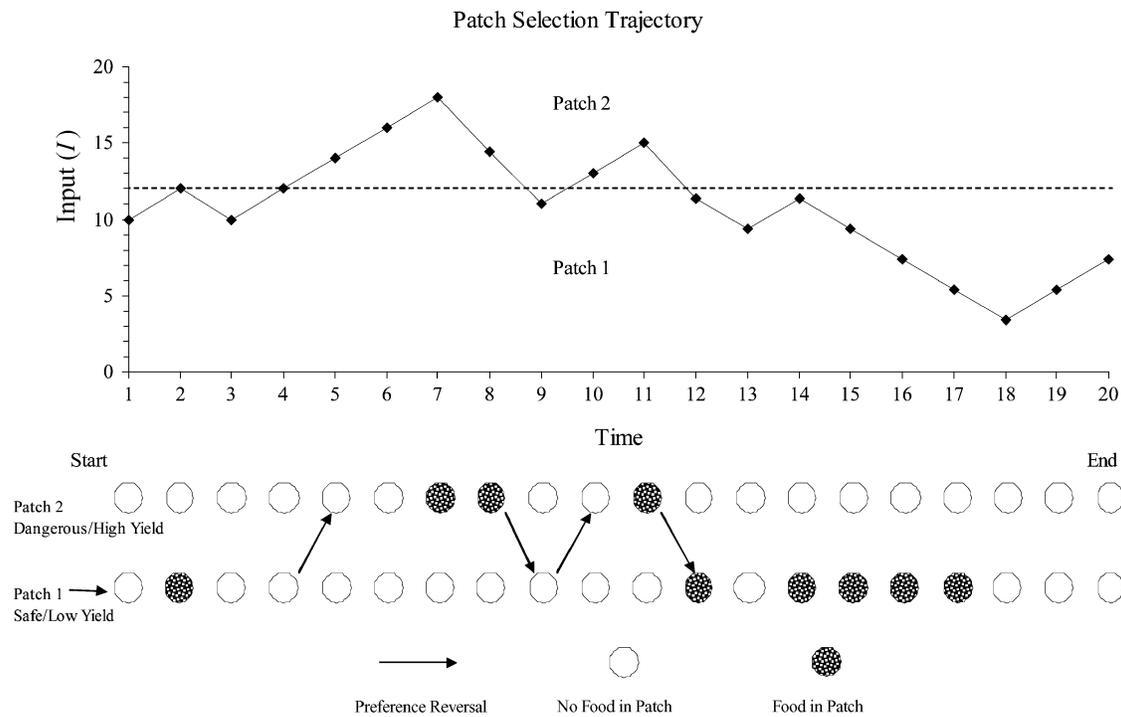


Figure 5. Patch selection trajectory determined by the level of input as a function of patch choice of either a “safe” but low-value food patch ($J_1^+ = 5.0$, $J_1^- = 0.5$) or a “dangerous” but high-value food patch ($J_2^+ = 9.0$, $J_2^- = 3.5$). The figure shows a sequence of 20 patch choices determined by the CEF model. The dashed line represents the input threshold (i.e., value of $I_0 = 12.0$), which indicates the motivational level at the point of a preference reversal (i.e., switching between patches). Changes in input I occur as a consequence of the probabilistic outcomes encountered while foraging (shown below the graph) and form a decision trajectory. I is inversely proportional to the energy reserves of the forager and can be thought of as proportional to the forager’s hunger level.

(Gettys, Reid, & Mellgren, 1988). The forager is assumed to start in the safer but less profitable patch and remain there as long as it continues to find food. Specifically, the forager is assumed to remain n successive trials without finding food (run of bad luck). It then switches to the more dangerous but more lucrative patch. Now, the goal of the forager is to remain in the dangerous patch for as short a time as possible, just long enough to recover sufficient energy reserves without being preyed upon. In this case, the appropriate rule becomes a run of good luck; the forager remains in the dangerous patch until it encounters m successive trials in which food is found. The integers n and m that give the best outcome over the entire foraging cycle does not depend on the current energy reserves of the forager, but n and m are computed as functions of the patch parameters. Although it would be possible to formulate a state-dependent ROL model, the role of the ROL model implemented here is to provide an instance of a state-independent model for comparison purposes.

Random Forager Model and Constant Patch Models

Three simple foraging strategies were introduced as control conditions to provide a basis for comparison of the OF, ET, CEF, and ROL strategies. The random for-

ager (RF) strategy simply flips a coin on each time cycle to determine which patch to visit. One would expect the random forager, who does not utilize patch parameters as the basis for decision making, to fare far worse than the others and thus to provide a useful lower bound on foraging performance. The two constant patch strategies (CP1) and (CP2) assume the animal simply remains in the same patch throughout its lifetime. Remaining solely in the dangerous (albeit high-return) patch should prove particularly disastrous and provides a baseline similar to that provided by the random forager. On the other hand, an animal that remains in the safe (but lower return) patch should fare much better with regard to death by predation but may be susceptible to death by starvation. However, any constant patch strategy should fall short of optimal performance. The constant patch strategy is an important control, since one would hope that any complicated, cognitively demanding strategy would do better than a simple stay-put strategy, so it needs to be demonstrated that the extra effort required by a more complex strategy is indeed worthwhile. Note that neither the random foraging nor the constant patch strategy depends on the forager’s energy reserves or the patch parameters.

Figure 6 graphically illustrates of the relationships between the models discussed here as a function of the properties of time dependence, state dependence, and

patch parameter dependence. Again, the time-independent ET model provides an upper bound on the performance of the CEF, which in turn should perform better than the state-independent ROL, random forager (RF), and constant patch (CP) models.

Comparison of Models

To compare the various decision rules, we plotted the proportion of deaths occurring in 30,000 runs against varying levels of a single Patch 2 parameter (either probability of finding food λ_2 or probability of predation β_2), while all other variables in Patch 2 and Patch 1 were held constant. The total proportion of deaths is attributable to predation and starvation combined.

In comparison with the optimal model, the CEF model is time independent; thus, no assumption about knowledge of the length of the foraging cycle is necessary. All foraging cycles have an endpoint—the point here is whether the model assumes that the forager is aware of it ahead of time. In the energy threshold model, the energy threshold that determines foraging decisions is constant over the foraging cycle T . This is similar to the mechanism employed by the CEF model, in which the value of tonic input I_0 , which produces a preference reversal, is constant over time and for a given set of patch parameters. Since the CEF model relies only on the product of λ_i and Y_i (i.e., the expected value), and not their specific values as in the optimal model, the energy threshold version of the optimal model actually represents the upper bound on performance by the CEF model. The parameters L , K , and ψ of the CEF model were chosen such that the tonic arousal level I_0 at which the model reverses its patch preference corresponds as closely as possible (in terms of least squared differences) to the patch reversal energy level of the energy threshold model. Presumably, the appropriate parameter values are learned through the animal’s experience with its environment, perhaps during an exploratory phase.

Comparisons among the various decision rules described previously involve computer simulations in which

a forager chooses between two patches designated as Patch 1 and Patch 2. The proportion of deaths as a function of increasing probability of predation risk ($\beta_2 = .005$ to $.070$) in Patch 2 is shown in Figure 7. In this set of simulations, there is no difference between the probabilities of finding food in the two patches ($\lambda_1 = \lambda_2 = .4$); however, the value of the food item in a patch is different ($Y_1 = 10, Y_2 = 18$). The range of predation probabilities plotted is assumed to represent realistic levels of risk encountered during foraging and has been used in other models (e.g., Bouskila & Blumstein, 1992). When the probabilities of predation are equal ($\beta_1 = \beta_2 = .005$), the optimal, ET, and CEF models always choose the larger payoff available in Patch 2. In general, there is a range of β_2 (.02 to $.045$) where the CEF model performs at near optimal levels. The CEF decision rule diverges from the OF and ET models in the upper range of β_2 values shown in Figure 7. The CEF model ultimately converges with the Patch 1 and the ROL decision rules at higher values of predation.

The proportion of patch choices for the CEF and ET models (i.e., patch residency) as a function of probability of predation β_2 in Patch 2 is shown in Figure 8. It can be seen that the proportion of Patch 1 residency increases as the probability of predation in Patch 2 increases, when using the same patch food item values as in Figure 7. The pattern of patch residency over the range of β_2 is similar for the CEF and ET models, except that the ET decision rule tends to utilize both patches more often in a simulation. Thus, the shift from Patch 2 to Patch 1 as a function of patch parameters is less of an all-or-none phenomenon in the ET model than in the CEF model.

Figure 9 shows the results from the simulations in which the probability of finding food in Patch 2 was varied from $\lambda_2 = .1$ to $.9$. In this set of simulations, the probabilities of predation in Patch 1 and Patch 2 were held constant at $\beta_1 = .005$ and $\beta_2 = .01$, respectively. The value of the food items in each patch was $Y_1 = 10$ and $Y_2 = 18$. The food-finding probabilities plotted are assumed to cover the range encountered during foraging.

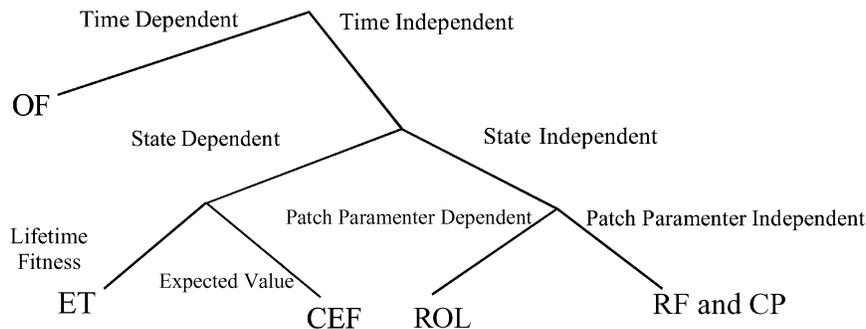


Figure 6. Tree diagram of the relationships between the models compared here as a function of the properties of time dependence, energy dependence, and patch parameter dependence. OF, optimal forager; ET, energy threshold; CEF, cognitive–emotional forager; ROL, run of luck; RF, random forager; CP, constant patch.

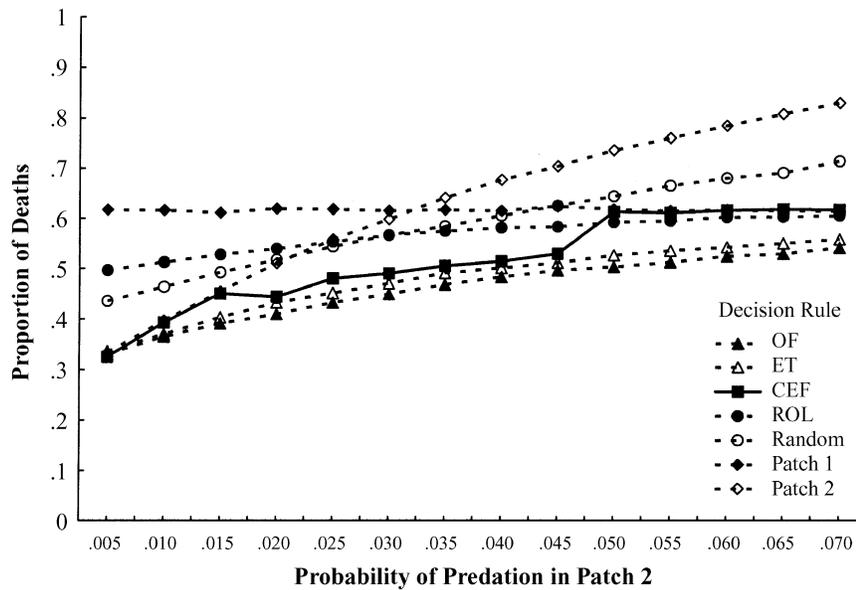


Figure 7. Forager mortality for each model in a two-patch choice simulation. The proportion of deaths plotted as a function of increasing probability of predation (β_2) in Patch 2. Patch 1 parameter values ($\lambda_1 = .4, \beta_1 = .005, Y_1 = 10.0$) and other Patch 2 parameter values ($\lambda_2 = .4, Y_2 = 18.0$) were held constant.

In general, the results of the computer simulations show that there is a range of λ_2 in which the CEF model performs at near optimum levels. The CEF decision model, like the optimal and ET models, predominately prefers Patch 1 at the lower λ_2 range. Evidently, the small chance of finding food in conjunction with the increased risk of

predation in Patch 2 make this a less acceptable choice to an optimal forager. As the chance of finding food in Patch 2 increases, the results show that the OF, ET, and CEF models diverge from Patch 1 and begin to prefer Patch 2. The CEF decision rule ultimately diverges from optimality as the values of λ_2 fall in the upper range for

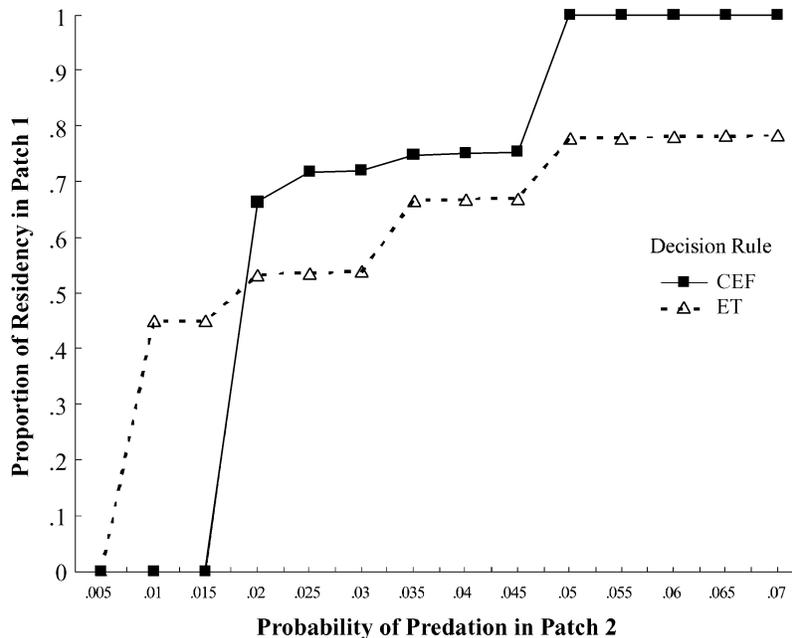


Figure 8. Patch 1 residency for the energy threshold (ET) model and the cognitive-emotional forager (CEF) model in a two-patch choice simulation.

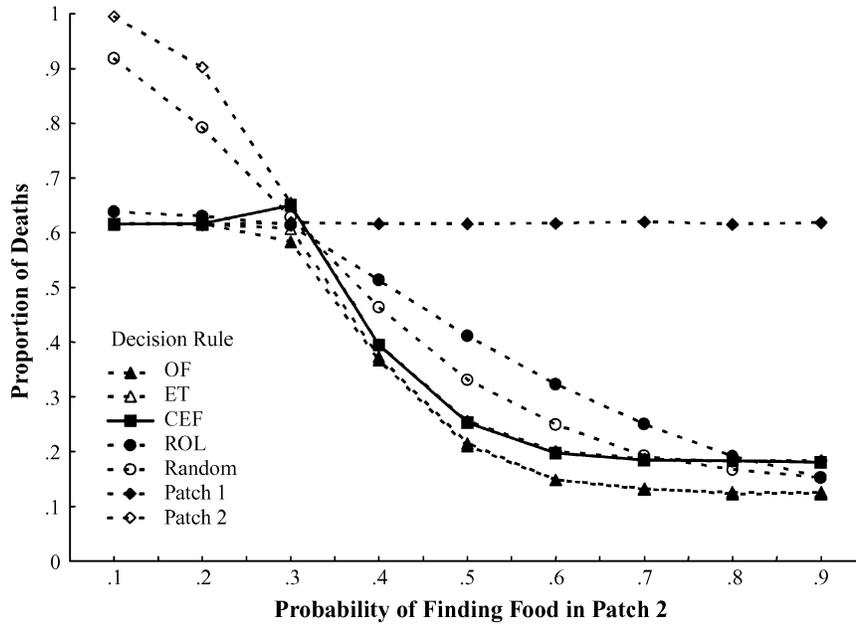


Figure 9. Forager mortality for the decision rules in a two-patch choice simulation. The proportion of deaths plotted as a function of increasing probability of finding food (λ_2) in Patch 2. Patch 1 parameter values ($\lambda_1 = .4$, $\beta_1 = .005$, $Y_1 = 10.0$) and other Patch 2 parameter values ($\beta_2 = .01$, $Y_2 = 18.0$) were held constant.

this set of simulations. Apparently, the CEF models a forager that ignores the higher probability of predation to take advantage of the higher expected gain afforded in Patch 2, as indicated by its similarity to the CP2 decision rule's performance at high values of λ_2 . The CEF model is comparatively better than the ROL decision rule at values of λ_2 greater than .3. The number of bad luck cycles (n) diminishes as a function of increasing λ_2 ; consequently, the ROL model tends to decrease the potential risk of starvation, albeit at the expense of a greater predation risk.

The ROL model is by definition a simple heuristic. It reflects decision making without the assumption of knowledge of the animal's internal state. Thus, the CEF model's enhanced performance relative to the ROL model may be attributed to the forager having knowledge about internal states. In general, foraging models that incorporate feedback from changes in the forager's internal state over time often have a quantitative advantage over equivalent time-independent models (Houston & McNamara, 1982, 1985). Thus, the computer simulations indicate that the CEF model is capable of near optimal foraging relative to the random and constant patch controls and is generally better than the ROL model.

Multiattribute CEF Model

Although the single-dipole version of the CEF model does well relative to the optimal model, it may not provide an entirely accurate representation of the trade-offs encountered by the forager. The opponent process idea generally involves the notion of opposites. Death by pre-

datation, although the strongest negative aspect of the situation, is not the opposite of successfully locating and consuming food. The two pairs of outcomes might be more properly characterized as (1) success versus failure in finding and consuming food and (2) success versus failure in avoiding predation. These two components, food and predation, can be represented by two separate dipoles for each patch the forager may encounter (see Figure 10). The output of the dipoles is assumed to be combined by the forager and then compared across patches. Therefore, the total anticipated gain is

$$\lambda_i Y_i + (1 - \beta_i) \psi_{1-\beta}$$

and the total anticipated loss is

$$(1 - \lambda_i) \alpha + \beta_i \psi_\beta$$

For simplicity, we assume the forager adds the positive outputs from the two dipoles, adds the negative outputs from the two dipoles, and takes the difference. A more general procedure would be to weight the output of each dipole differently, representing the differential importance of the two attributes. We assume here that the parameters representing the values of the attributes (i.e., the value of food and the value of life) serve to weight the attributes. We selected parameters of the multiattribute CEF that would improve performance, to some degree, over the single-dipole version but that were not systematically optimized, as they were for the single-dipole version.

In Figures 11 and 12, the performance of the multiattribute CEF (M) is plotted relative to the performance of

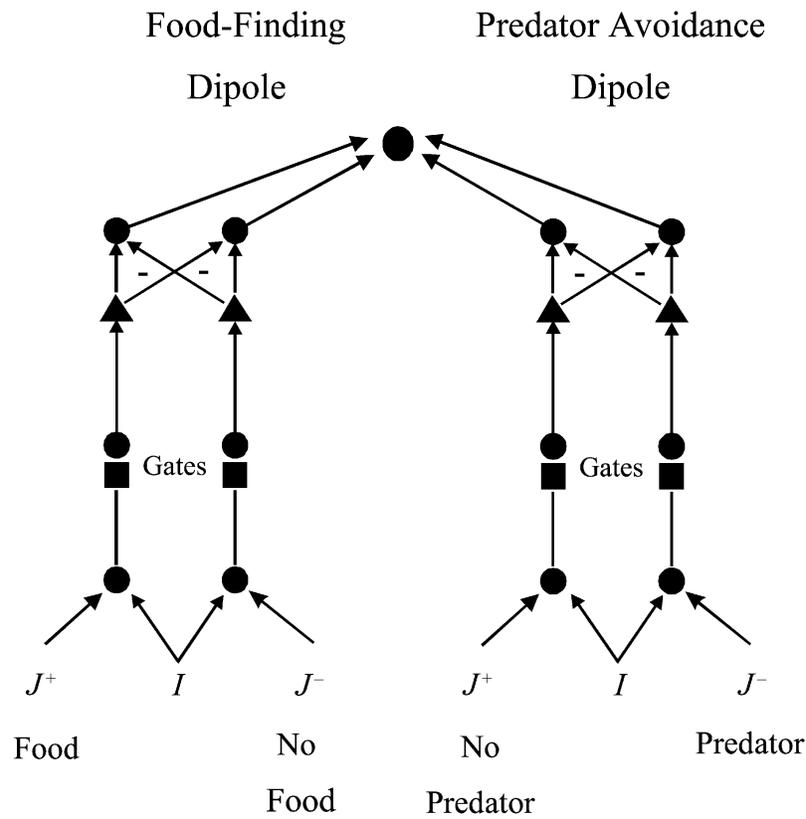


Figure 10. Two gated dipoles, as in Figure 3, are combined to create a multiattribute CEF model in which the outputs of pairs of opposite affective components are summed. The two channels in one dipole, labeled the “Food-Finding Dipole,” represent finding and not finding food. The two channels in the other dipole, labeled the “Predator Avoidance Dipole,” represent avoiding and being subject to predation.

the single-dipole CEF (S) and the ET model. In the multiattribute version of CEF, the near optimal performance is extended over a slightly wider range when survival is plotted against the probability of predation (β_2) and the probability of finding food (λ_2) (Figures 11 and 12, respectively). In general, the overall performance obtained by the multiattribute CEF model is consistent with the single-dipole version and appears to be slightly superior to that of the single-dipole version. The additional architecture of the multiattribute version of the CEF model may be a more realistic neural network and thus theoretically more appealing for modeling the psychophysiological mechanisms of foraging under risk of predation. For example, the food-finding dipole could be implemented as part of Panksepp’s (1982) foraging expectancy system and the predator avoidance dipole could be part of the fear–anxiety system.

The multiattribute CEF model enables foraging decisions to be based on both gain and loss components of fitness related to the cost of predation, as discussed in Olsson, Brown, and Smith (2002). Olsson et al. define cost of predation as encompassing not only the probability of predation itself (i.e., risk) but also the quality of the habitat and the expected reproductive fitness of the forager. That is, a

fitter animal with high expectation of future foraging success has more to lose by being killed than one that is less fit or less likely to find food later. Similar to the multiattribute model’s value of life, the cost of predation includes the expected gain in fitness when surviving (i.e., a cost of being killed). In this way, foraging decisions are a reflection of fitness when modeled by state-dependent changes in energy gain and survival. Thus, the patch choice model using the multiattribute version of CEF is the sum of costs and gains from independent components of fitness. When the value of life is low (i.e., low expected fitness), the multiattribute CEF will tend to select the more dangerous patch, whereas when the value of life is high, the CEF will tend to select the safer patch. This is because the final patch choice is gated by the combination of expected loss due to predation and expected energy intake.

Discussion

Comparison with other neural network models of foraging. In a previous application of a neural network in evolutionary optimization, Mangel (1990) compared a neural network model with the SDP solution of a primitive organism (invertebrate) foraging under predation risk. The simulation results indicated that the neural net-

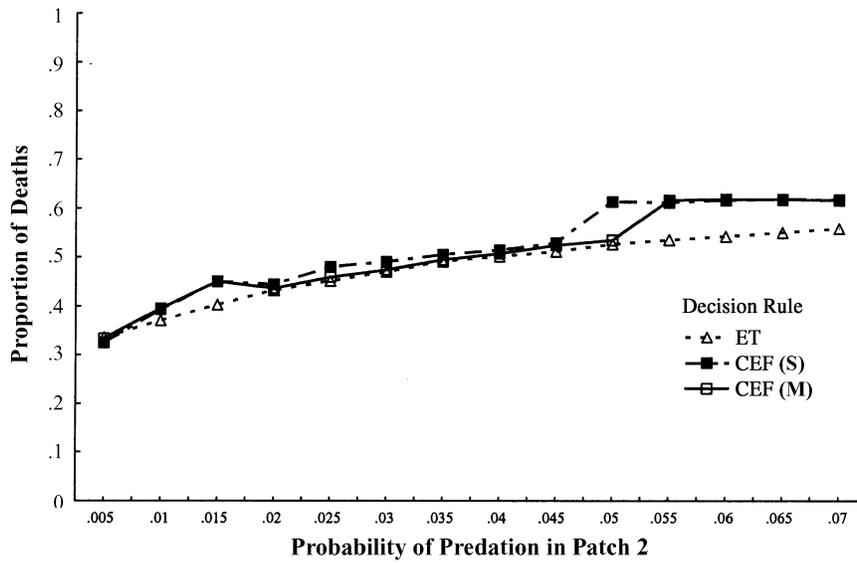


Figure 11. Forager mortality for the multiattribute (M) and the single-dipole (S) versions of the CEF model in a two-patch choice simulation. The proportion of deaths is plotted as a function of increasing probability of predation (β_2) in Patch 2 using Figure 7 parameter values.

work model was capable of obtaining between 85% and 99% of the optimal fitness as measured by individual survival and reproduction. This finding supports the view that neural network models can achieve fitness at near optimal levels, which lends credibility to insights about real biological systems derived from neural network mechanisms. While Mangel's model is based on the minimum number of neural groups necessary to sim-

ulate foraging behavior (e.g., motor, sensory, and energy reserve nodes), our CEF model utilizes a more general (less species-specific) neural network capable of modeling foraging decisions based on motivational, affective, and cognitive processes. Thus, the theory of decision making of the CEF model is intended to explain foraging decisions made under predation risk in a wider range of organisms than the model developed by Mangel.

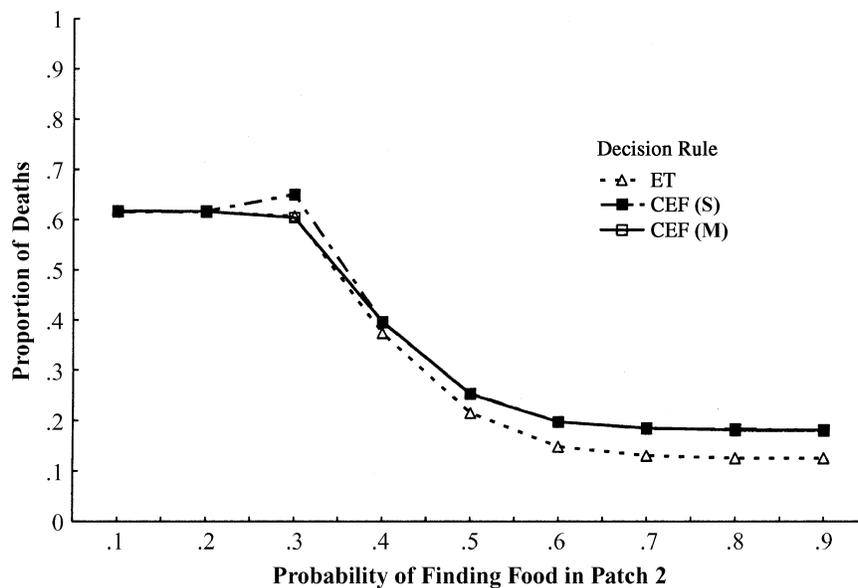


Figure 12. Forager mortality of the multiattribute (M) and the single-dipole (S) versions of the CEF model in a two-patch simulation. The proportion of deaths is plotted as a function of increasing probability of finding food (λ_2) in Patch 2 using Figure 9 parameter values.

Real (1992) anticipated many of the themes of this article in his sophisticated review of data on context-sensitive decision making, for which he proposed a series of inter-related quantitative models, including one neural network model. Real's neural network was based on a variant of the back propagation network of Rumelhart, Hinton, and Williams (1986), which is widely utilized in applications since theorems have shown it can learn an arbitrary mapping from n -dimensional to m -dimensional vectors. The back propagation network, however, is not based on biologically realistic network architectures or functional psychological principles (see Levine, 2000, chap. 6, for review); hence, Real's network does not include mechanisms for actual foraging decisions in real time.

Niv et al. (2002a, 2002b) utilized a neural network to simulate evolutionary development of reinforcement learning in foraging bees. They showed that risk-averse behavior emerged parsimoniously from their model without the need of a nonlinear utility function of the amount of reward. Unlike our model, their model considered only the risk of starvation and not predation. The model of Niv and her colleagues is based on a variant of the temporal difference (TD) paradigm of Sutton and Barto (1998). This paradigm, like the gated dipole utilized in our network, is sensitive to temporal contrasts in amount of reinforcement. Yet the TD model lacks the gated dipole's sensitivity to motivational level (nonspecific arousal), which plays a major role in our CEF model.

To our knowledge, the CEF is the first neural network foraging model that incorporates in a systematic way the affective and contextual influences that form the basis for most animal, as well as human, decisions under risk. The modeling framework in this article is based on general neural opponent processing principles: We hope in future articles to specify particular quantitative instantiations of these principles that can lead to simulations of real-time data on foraging decisions under risk of either starvation (Marsh & Kacelnik, 2002; Reboresda & Kacelnik, 1991) or predation (Olsson, Brown, & Smith, 2001, 2002).

Implications and predictions of the CEF model.

Of particular interest to those who study foraging is the interaction between external and internal causal factors that underlie behavior. Historically, this has been a rather difficult problem when considering the consequences of several activities over the life of an animal (e.g., foraging, predator avoidance, territorial defense, etc.). Fitness in terms of future reproductive success is determined by the cost and benefits of engaging in different activities. Thus, fitness provides a "common currency" of forming decisions (Houston et al., 1988). OF theory has utilized stochastic dynamic programming to investigate the trade-offs between activities (i.e., foraging under predation risk) in the form of this common currency. Although beneficial in the investigation of animal decision making, the SDP approach is limited because it fails to specify the proximate mechanisms that produce optimal foraging, and the analytical solution does not realistically depict

actual cognitive abilities possessed by animals. Nevertheless, SDP does provide the basis for comparison between models, as it was used here to compare our CEF model, via Monte Carlo simulations, with an OFT model and with an ROL ROT model.

Simulations of the CEF model can at times yield the same kind of impulsive decision making that occurs in real animals under conditions that combine high risk and high gain (e.g., Gilliam & Fraser, 1987). Under such conditions a CEF forager tends to choose a high-valued food item, whereas the optimal model chooses the alternative patch under the same conditions of predation risk and energy reserves (see Figure 9). High values of expected gains or losses (i.e., predation risk) may have counterintuitive effects on behavior.

Olsson et al. (2001, 2002) have shown that risk taking in European starlings (*Sturnus vulgaris*) is best characterized by more than just the probability of encountering a predator. The expected fitness after surviving predation is also an important consideration. Starlings accustomed to low habitat quality showed decreased avoidance responses to predation risk and increased approach responses to gains, possibly due to low perceived fitness. Thus, the prospect of high-quality food may alter behavioral antipredator mechanisms underlying patch selection via manipulation of cost of predation.

In fact, Olsson et al. (2001, 2002) define the cost of predation as the product of the probability of predation and the fitness level of the organism after surviving a predation attempt. This is captured in the CEF model as the product $J^- = \beta\psi$ (the expected loss due to predation—which is the input to the off channel—is the product of the probability of predation and the value of life). As a consequence, the high expected value of a food item in a dangerous patch can be the incentive for making risky decisions. The incentive motivation incorporated in the gated dipole serves to function as a psychological mechanism capable of dealing with long- and short-term state-dependent foraging effects evident in studies of foraging under predation risk. Some evidence suggests that incentive motivation increases nonspecific arousal and is known to alter preferences on concurrent schedules of reinforcement (Killeen, 1982a, 1982b, 1985).

The gated dipole model allows for both utility-maximizing and affective tendencies to influence decision making. Thus, it seems likely that, for at least some species of foragers, the CEF model simulates the data better than does the optimal model, in that CEF demonstrates some degree of the risk taking in which real animals engage. In CEF, the amount of risk taking can be varied by parametric changes in our value of life parameter: The higher the value of life, the greater the risk aversion. Perhaps, in addition to differences based on age and reproductive success, there are significant genetic differences within a foraging species in the trait of risk aversion, as has been demonstrated in humans (e.g., Zuckerman, 1990); this would be a fruitful avenue for further empirical research.

A central tenet of the CEF model is the concept of opponent processes of motivation. Clearly, foraging behavior is mediated by appetitive (i.e., food finding) neural mechanisms that elicit approach responses. Foraging under the risk of predation will lead to an interaction between appetitive systems and aversive systems that mediate predator avoidance responses. In the simple version of our model, the appetitive drive of food finding is directly linked to the predator avoidance drive, forming a single system. The activation of internal representations of emotionally charged appetitive and aversive events encountered during foraging activates a "central" motivational system that mediates the behavioral response. Thus, the opponent process "translates" the emotionally charged events into a common currency of affect, and the decision maker attempts to maximize the subjective affective response.

In the multiattribute version of CEF, appetitive and aversive stimuli are processed by separate gated dipoles combined in a central neural connection. This network represents opponent processing more accurately than does the single-attribute network, because, in each gated dipole, the two channels of the network represent actual opposites (food and its absence; predation and its absence). It is perhaps for this reason that the multiattribute model is closer to optimal than is that of the single-attribute model, particularly when food-finding probability is low or predation risk is high (see Figures 11 and 12).

In our model, the source of motivational change is the nonspecific tonic arousal input of the gated dipole. Several implications may result from linking arousal with behavioral change. The term *arousal* itself may have many connotations in relation to motivation. For example, some meanings of arousal involve changes in activation, whereas others refer to changes in drive (see Andrew, 1974, for a review). Grossberg (1982) considered I to be the constant amount of activation received by the gated dipole. I is assumed to have an additive relation to specific drive inputs; therefore, arousal and drive are not the same concept in Grossberg's theory. Grossberg observed several effects of sudden or unexpected changes in I that relate to motivation and learning. For example, a burst of I , as the result of a surprising or unexpected event, has effects similar to that of the offset of specific drive input J —namely, eliciting the opposite response. Grossberg and Gutowski (1987) predicted a change in risk sensitivity in underaroused individuals (p. 306). Thus far, there have been few experimental studies that examine the effects of arousal on risk-sensitive decision making. Our model further suggests that this could be a fruitful line of empirical research.

In sum, we argue, as have others, that foraging under predation risk is a dynamic process in which decision making involves the ability to trade off desirable goals against potential negative outcomes. Animals are assumed to possess psychophysiological mechanisms that mediate decision making. Future investigation should

determine how such a network could learn about the patch parameters. The CEF model presented here demonstrates performance comparable to that of normative OF models and includes plausible proximate mechanisms that mediate (optimal or suboptimal) behavior. Thus, neural networks may provide a theoretical basis for the development of realistic models of foraging behavior through an interdisciplinary approach based on psychological and behavioral ecology. Such biologically based foraging models can also influence the development of models for human consumer decision making, including effects of relevant economic variables, such as price and delay (Rajala & Hantula, 2000; Smith & Hantula, 2003). Hence, similar opponent processing principles appear to explain the influence of short-term affective reaction in decisions by a wide range of species, from insects to humans.

REFERENCES

- ANDREW, R. J. (1974). Arousal and the causation of behaviour. *Behaviour*, **51**, 135-164.
- BELLMAN, R. (1957). *Dynamic programming*. Princeton, NJ: Princeton University Press.
- BERRIDGE, K. C., & ROBINSON, T. E. (1998). What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews*, **28**, 309-369.
- BOUSKILA, A., & BLUMSTEIN, D. T. (1992). Rules of thumb for predation hazard assessment: Predictions from a dynamic model. *American Naturalist*, **139**, 161-176.
- BROWN, J. S. (1999). Vigilance, patch use, and habitat selection: Foraging under predation risk. *Evolutionary Ecology Research*, **1**, 49-71.
- BUCK, R. (1985). Prime theory: An integrated view of motivation and emotion. *Psychological Review*, **92**, 389-413.
- CABANAC, M. (1985). Influence of food and water deprivation on the behavior of the white rat foraging in a hostile environment. *Physiology & Behavior*, **35**, 701-709.
- CALDWELL, G. S. (1986). Predation as a selective force on foraging herons: Effects of plumage color and flocking. *Auk*, **103**, 494-505.
- CARACO, T., & LIMA, S. L. (1987). Survival, energy budgets, and foraging risk. In M. L. Commons, A. Kacelnik, & S. J. Shettleworth (Eds.), *Foraging: Quantitative analyses of behaviour* (Vol. 6, pp. 1-21). Hillsdale, NJ: Erlbaum.
- CARACO, T., MARTINDALE, S., & WHITTAM, T. S. (1980). An empirical demonstration of risk-sensitive foraging preferences. *Animal Behaviour*, **28**, 820-830.
- COWLISHAW, G. (1997). Trade-offs between foraging and predation risk determine habitat use in a desert baboon population. *Animal Behaviour*, **53**, 667-686.
- DILL, L. M. (1987). Animal decision making and its ecological consequences: The future of aquatic ecology and behaviour. *Canadian Journal of Zoology*, **65**, 65-71.
- GETTYS, C. F., REID, P. J., & MELLGREN, R. L. (1988, November). *Can we distinguish between optimal foraging decisions and rules of thumb?* Paper presented at the annual meeting of the Psychonomic Society, Chicago.
- GILLIAM, J. F., & FRASER, D. F. (1987). Habitat selection under predation hazard: Test of a model with foraging minnows. *Ecology*, **68**, 1856-1862.
- GOSS-CUSTARD, J. D. (1977). Optimal foraging and the size selection of worms by redshank, *Tringa totanus*, in the field. *Animal Behaviour*, **25**, 10-29.
- GRAY, J. A. (1977). Drug effects on fear and frustration: Possible limbic sites of action of minor tranquilizers. In L. L. Iverson, S. D. Iverson,

- & S. H. Snyder (Eds.), *Handbook of psychopharmacology* (Vol. 8, pp. 433-529). New York: Plenum.
- GRAY, J. A. (1991). Neural systems, emotion and personality. In J. Madden IV (Ed.), *Neurobiology of learning, emotion, and affect* (pp. 273-306). New York: Raven.
- GROSSBERG, S. (1972a). A neural theory of punishment and avoidance: I. Qualitative theory. *Mathematical Biosciences*, **15**, 39-67.
- GROSSBERG, S. (1972b). A neural theory of punishment and avoidance: II. Quantitative theory. *Mathematical Biosciences*, **15**, 253-285.
- GROSSBERG, S. (1982). Processing of expected and unexpected events during conditioning and attention: A psychophysiological theory. *Psychological Review*, **89**, 529-572.
- GROSSBERG, S., & GUTOWSKI, W. E. (1987). Neural dynamics of decision making under risk: Affective balance theory and cognitive-emotional interactions. *Psychological Review*, **94**, 300-318.
- GROSSBERG, S., & MERRILL, J. W. L. (1996). The hippocampus and cerebellum in adaptively timed learning, recognition, and movement. *Journal of Cognitive Neuroscience*, **8**, 257-277.
- HOUSTON, A. I., CLARK, C., MCNAMARA, J. M., & MANGEL, M. (1988). Dynamic models in behavioural and evolutionary ecology. *Nature*, **332**, 29-34.
- HOUSTON, A. I., & MCNAMARA, J. M. (1982). A sequential approach to risk-taking. *Animal Behaviour*, **30**, 1260-1261.
- HOUSTON, A., & MCNAMARA, J. M. (1985). The variability of behaviour and constrained optimization. *Journal of Theoretical Biology*, **112**, 265-273.
- HOUSTON, A. I., MCNAMARA, J. M., & HUTCHINSON, J. M. C. (1993). General results concerning the trade-off between gaining energy and avoiding predation. *Philosophical Transactions of the Royal Society of London: Series B*, **341**, 375-397.
- IWASA, Y., HIGASHI, M., & YAMAMURA, N. (1981). Prey distribution as a factor determining the choice of optimal foraging strategy. *American Naturalist*, **117**, 710-723.
- JANETOS, A., & COLE, B. (1981). Imperfectly optimal animals. *Behavioral Ecology & Sociobiology*, **9**, 203-210.
- KAVALIERS, M., & CHOLERIS, E. (2001). Antipredator responses and defensive behavior: Ecological and ethological approaches for the neurosciences. *Neuroscience & Biobehavioral Reviews*, **25**, 577-586.
- KILLEEN, P. R. (1982a). Incentive theory. In D. J. Bernstein (Ed.), *Nebraska Symposium on Motivation: Response structure and organization* (Vol. 9, pp. 169-216). Lincoln: University of Nebraska Press.
- KILLEEN, P. R. (1982b). Incentive theory: II. Models for choice. *Journal of the Experimental Analysis of Behavior*, **38**, 217-232.
- KILLEEN, P. R. (1985). Incentive theory: IV. Magnitude of reward. *Journal of the Experimental Analysis of Behavior*, **43**, 407-417.
- KREBS, J. R. (1978). Optimal foraging: Decision rules for predators. In J. R. Krebs & N. B. Davies (Eds.), *Behavioural ecology* (pp. 23-62). Oxford: Blackwell.
- LE MAGNEN, J. (1992). *Neurobiology of feeding and nutrition*. San Diego: Academic Press.
- LEVEN, S. J., & LEVINE, D. S. (1996). Multiattribute decision making in context: A dynamic neural network methodology. *Cognitive Science*, **20**, 271-299.
- LEVINE, D. S. (2000). *Introduction to neural and cognitive modeling* (2nd ed.). Mahwah, NJ: Erlbaum.
- LEVINE, D. S., & ELSBERRY, W. R. (Eds.) (1997). *Optimality in biological and artificial networks?* Mahwah, NJ: Erlbaum.
- LEVINE, D. S., & PRUEITT, P. S. (1989). Modeling some effects of frontal lobe damage: Novelty and perseveration. *Neural Networks*, **2**, 103-116.
- LIMA, S. L., & DILL, L. M. (1990). Behavioural decisions made under the risk of predation: A review and prospectus. *Canadian Journal of Zoology*, **68**, 619-640.
- LIMA, S. L., VALONE, T. J., & CARACO, T. (1985). Foraging efficiency-predation risk trade-off in the grey squirrel. *Animal Behaviour*, **33**, 155-165.
- MANGEL, M. (1990). Evolutionary optimization and neural network models of behavior. *Journal of Mathematical Biology*, **28**, 237-257.
- MANGEL, M., & CLARK, C. (1988). *Dynamic modeling behavior in ecology*. Princeton, NJ: Princeton University Press.
- MARSH, B., & KACELNIK, A. (2002). Framing effects and risky decisions in starlings. *Proceedings of the National Academy of Sciences*, **99**, 3352-3355.
- McFARLAND, D. J. (1977). Decision making in animals. *Nature*, **269**, 15-21.
- McFARLAND, D. J. (1982). Introduction to functional analysis of behaviour. In D. J. McFarland (Ed.), *Functional ontogeny* (pp. 3-22). London: Pitman.
- McFARLAND, D. J., & HOUSTON, A. (1981). *Quantitative ethology: The state-space approach*. London: Pitman.
- MCNAMARA, J. M., & HOUSTON, A. (1986). The common currency for behavioral decisions. *American Naturalist*, **127**, 358-378.
- MILINSKI, M., & HELLER, R. (1978). Influence of a predator on the optimal foraging behaviour of sticklebacks (*Gasterosteus aculeatus*). *Nature*, **275**, 642-644.
- NIV, Y., JOEL, D., MEILJSON, I., & RUPPIN, E. (2002a). Evolution of reinforcement learning in foraging bees: A simple explanation for risk-averse behavior. *Neurocomputing*, **44**, 951-956.
- NIV, Y., JOEL, D., MEILJSON, I., & RUPPIN, E. (2002b). Evolution of reinforcement learning in uncertain environments: A simple explanation for complex foraging behaviors. *Adaptive Behavior*, **10**, 5-24.
- OLSSON, O., BROWN, J. S., & SMITH, H. G. (2001). Gain curves in depletable food patches: A test of five models with European starlings. *Evolutionary Ecology Research*, **3**, 285-310.
- OLSSON, O., BROWN, J. S., & SMITH, H. G. (2002). Long- and short-term state-dependent foraging under predation risk: An indication of habitat quality. *Animal Behaviour*, **63**, 981-989.
- PANKSEPP, J. (1982). Toward a general psychobiological theory of emotions. *Behavioral & Brain Sciences*, **5**, 407-467.
- RAJALA, A. K., & HANTULA, D. A. (2000). Towards a behavioral ecology of consumption: Delay-reduction effects on foraging in a simulated Internet mall. *Managerial & Decision Economics*, **21**, 145-158.
- REAL, L. A. (1992). Information processing and the evolutionary ecology of cognitive architecture. *American Naturalist*, **140**(Suppl.), S108-S145.
- REBORDA, J. C., & KACELNIK, A. (1991). Risk sensitivity in starlings: Variability in food amount and food delay. *Behavioral Ecology*, **2**, 301-308.
- ROZIN, P., & SCHULKIN, J. (1990). Food selection. In E. M. Stricker (Ed.), *Handbook of behavioral neurobiology: Neurobiology of food and fluid intake* (Vol. 10, pp. 297-328). New York: Plenum.
- RUMELHART, D. E., HINTON, G. E., & WILLIAMS, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1. Foundations* (pp. 318-362). Cambridge, MA: MIT Press.
- SCHOENER, T. W. (1971). Theory of feeding strategy. *Annual Review of Ecological Systems*, **2**, 369-404.
- SHAFIR, S. (2000). Risk-sensitive foraging: The effect of relative variability. *Oikos*, **88**, 663-669.
- SHEVELKIN, A. V. (1994). Facilitation of defense reactions during the consumption of food in snails: The participation of glucose and gastrin/cholecystokinin-like peptide. *Neuroscience & Behavioral Physiology*, **24**, 115-124.
- SIBLY, R., & McFARLAND, D. J. (1974). A state-space approach to motivation. In D. J. McFarland (Ed.), *Motivational control systems analysis* (pp. 213-250). London: Academic Press.
- SIBLY, R., & McFARLAND, D. J. (1976). On the fitness of behavioral sequences. *American Naturalist*, **10**, 601-617.
- SMITH, C. L., & HANTULA, D. A. (2003). Pricing effects on foraging in a simulated Internet shopping mall. *Journal of Economic Psychology*, **24**, 653-674.
- SNYDERMAN, M. (1983). Optimal prey selection: The effects of food deprivation. *Behavior Analysis Letters*, **3**, 359-369.
- SOLOMON, R. L., & CORBIT, J. (1974). An opponent-process theory of motivation: I. Temporal dynamics of affect. *Psychological Review*, **81**, 119-145.
- STEPHENS, D. W., & CHARNOV, E. L. (1982). Optimal foraging: Some simple stochastic models. *Behavioral Ecology & Sociobiology*, **10**, 251-263.

- STEPHENS, D. W., & KREBS, J. R. (1986). *Foraging theory*. Princeton, NJ: Princeton University Press.
- SUTTON, R. S., & BARTO, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge, MA: MIT Press.
- TOWNSEND, J. T., & BUSEMEYER, J. R. (1989). Approach-avoidance: Return to dynamical decision behaviour. In C. Izawa (Ed.), *Current issues in cognitive processes* (pp. 107-133). Hillsdale, NJ: Erlbaum.
- TVERSKY, A., & KAHNEMAN, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, **185**, 1124-1131.
- TVERSKY, A., & KAHNEMAN, D. (1981). The framing of decisions and the psychology of choice. *Science*, **211**, 453-458.
- WEBER, E. U., SHAFIR, S., & BLAIS, A.-R. (2004). Predicting risk sensitivity in humans and lower animals: Risk as variance or coefficient of variation. *Psychological Review*, **111**, 430-445.
- WELTON, N. J., MCNAMARA, J. M., & HOUSTON, A. I. (2003). Assessing predation risk: Optimal behaviour and rules of thumb. *Theoretical Population Biology*, **64**, 417-430.
- WERNER, E., & HALL, D. (1974). Optimal foraging and the size selection of prey by the bluegill sunfish (*Lepomis macrochirus*). *Ecology*, **55**, 393-425.
- ZUCKERMAN, M. (1990). The psychophysiology of sensation seeking. *Journal of Personality*, **58**, 313-345.

NOTES

1. The CEF model may be expanded to accommodate more than two patches. In a CEF model of three or more patches, the forager compares the current or resident patch with all others in a pairwise fashion and then selects the patch that produces the largest positive output of the gated dipole. Preliminary results indicate that with three patches, the model's performance is similar to that of the two-patch case.

2. Under some parameter values, the CEF model adopts the invariant patch selection strategy. This might be reasonable when one of the patches provides an extremely low return or a dangerously high risk of predation. However, if more flexibility is desired, a future modification of the CEF model could be to use the drive input (J) rather than the tonic arousal input (I) as the source of motivation produced by hunger and fear of predation. Here, weighted drive-specific inputs (J^+ , J^-) could be used as an alternate method to model decision making. Furthermore, the current multiattribute version of CEF inputs equal levels of I into each gated dipole. This could be modified, say, to include separate nonspecific arousal levels relating to hunger and to fear.

(Manuscript received October 22, 2004;
revision accepted for publication June 22, 2005.)