Bulletin of Mathematical Biology Vol. 54, No. 2/3, pp. 355–378, 1992.
Printed in Great Britain.

0092–8240/92$5.00 + 0.00
Pergamon Press plc
Society for Mathematical Biology

RISK-SENSITIVE FORAGING: A REVIEW OF THE THEORY

■ JOHN M. MCNAMARA
School of Mathematics,
University of Bristol,
University Walk,
Bristol BS8 1TW, U.K.
(*E.mail: majmm@UK.ac.bristol.npla*)

■ ALASDAIR I. HOUSTON
NERC Unit of Behavioral Ecology,
Department of Zoology,
University of Oxford,
South Parks Road,
Oxford OX1 3PS, U.K.

Introduction. The energy gained by an animal during a foraging bout typically increases the animal's expected future reproductive success (EFRS). The greater the energy gain the greater the increase, but the relationship between gain and EFRS need not be linear. To be specific, consider a foraging bout of fixed length. Let $\mathrm{f}(y)$ denote the EFRS of an animal that has gained y units of energy by the end of the bout. Then f will typically be an increasing function but may for example be concave, i.e. $\mathrm{f}''(y) \leqslant 0$ for all y; or may be convex, i.e. $\mathrm{f}''(y) \geqslant 0$ for all y; or may be concave for part of its range and convex on another.

Suppose that an animal has a choice between a number of foraging options during the bout. The energy gained during the interval, Y, is a random variable whose distribution is determined by the choice of option. The EFRS of an animal at the beginning of the bout is $s\ \mathbb{E}\{\mathrm{f}(Y)\}$ where s is the probability of surviving the bout. If we assume that s is the same for all available options then an animal maximizes its fitness by choosing any alternative which maximizes $\mathbb{E}\{\mathrm{f}(Y)\}$. When f is linear we have $\mathbb{E}\{\mathrm{f}(Y)\} = \mathrm{f}(\mathbb{E}\{Y\})$ so that fitness is maximized by maximizing the mean gain over the interval, $\mathbb{E}\{Y\}$. When f is nonlinear this need no longer be true.

In this paper we are primarily interested in a range of feeding options which yield the same mean gain, $\mathbb{E}\{Y\}$, but differ in their variability. In particular, suppose an animal can choose between an option which yields $\mu = \mathbb{E}\{Y\}$ with certainty and an alternative which yields a stochastic amount of food with the same mean μ. When f is concave we have $\mathbb{E}\{f(Y)\} \leqslant f(\mathbb{E}\{Y\})$ by Jensen's inequality, and fitness is maximized by choosing the deterministic option. Such a preference for an option over an alternative yielding the same mean gain but having higher variance in reward magnitude is called risk averse behaviour. When f is convex Jensen's inequality implies that $\mathbb{E}\{f(Y)\} \geqslant f(\mathbb{E}\{Y\})$ and fitness is maximized by choosing the variable option. Preference for a variable option over a less variable alternative with the same mean is called risk prone behaviour. Risk sensitive behaviour has been analysed in a number of theoretical models. Here we review functional models which predict such behaviour. Risk sensitive behaviour has also been sought and found empirically. Observational evidence is reviewed by Real and Caraco (1986) and Ellner and Real (1989).

Whether it is optional to be risk averse or risk prone depends on the shape of f, but an animal's EFRS and its dependence on the energy gain over an interval depend on circumstance. Our review of the risk-sensitive foraging theory concentrates on models that attempt to derive the shape of f from the animal's biology. As we demonstrate below, the shape of f depends on:

- (i) the energy reserves of the animal and time of day;
- (ii) the future foraging environment. In particular:
    - (a) the quality of food,
    - (b) whether there is fixed time such as dusk at which foraging must stop,
    - (c) whether foraging is likely to be disturbed by interruptions such as snow cover;
- (iii) the biological significance of the energy. For example:
    - (a) whether energy is used to avoid starvation,
    - (b) whether energy is put directly into reproduction,
    - (c) whether energy is put directly into growth.

One of the main messages of this review is that there is no universal model of risk-sensitive behaviour. There are, however, modelling methods and principles which can be used to predict risk-sensitive behaviour if the biological circumstances are identified.

The next section presents two ways in which the detailed process of food intake can be formulated. Most of the risk-sensitive models we present use one of these formulations. Although the formulations are very different, it is reassuring that predictions of risk-sensitive behaviour do not seem to depend on which is used in a model.

The models of risk-sensitive behaviour presented fall into three categories. We begin by analysing a range of models in which food is used solely to avoid starvation. Predictions of risk sensitivity then depend on the assumptions of the model and are summarized in Table 1.

To show that predictions can also depend crucially on the biological significance of energy obtained we modify a starvation model so that an animal is using energy both for survival and reproduction. Table 2 compares the predictions of this modified model with the original pure starvation model.

In our third category of model we explicitly consider models in which there is variability in the delay until an item is found rather than variability in the amount obtained. We demonstrate that changing the type of variability leads to very different predictions in some models, while giving similar predictions in others.

Finally we consider the strength of selection pressure on risk-sensitive decision making.

Models of Food Intake. Sutherland and Anderson (1987) discuss various ways in which differences in variance in energetic gain can arise. One possibility that has been given considerable attention is that an animal can reduce the variance in its intake rate by foraging in a group rather than foraging alone (e.g. Caraco, 1981; Pulliam and Millikan, 1982; Clark and Mangel, 1984, 1986; Clark, 1987; Ekman and Rosander, 1987).

We focus on two rather different models of food intake. The first model is highly schematic in that both time and reserves are modelled as taking integer values. The model captures the essence of the risk sensitive element in foraging and is often amenable to easy computation and analytic analysis. The second model is based on the standard prey choice paradigm (e.g. Charnov, 1976; Stephens and Krebs, 1986), and while being a more realistic representation of an animal's actual foraging process is usually more difficult to analyse.

Discrete time and state model. At each of the times $t=0, 1, 2, \ldots$ an animal must choose one of two options labelled Option 1 and Option 2. If it chooses Option i at time t it finds 0, 1 or 2 food items between t and $t+1$ with probabilities q_i, $1-(q_i+p_i)$ and p_i, respectively. All items have energetic content of one unit. During each time one unit of energy is used in metabolic expenditure. Thus the transition law for the animal's energy reserves $X(t)$ at time t is given by:

$$\begin{aligned} &P(X(t+1)=x-1 \mid X(t)=x, \text{ Option } i)=q_i \\ &P(X(t+1)=x \quad\;\; \mid X(t)=x, \text{ Option } i)=1-(p_i+q_i) \\ &P(X(t+1)=x+1 \mid X(t)=x, \text{ Option } i)=p_i. \end{aligned} \tag{1}$$

The mean net rate of energy gain and variance in gain per unit time under Option i are given by:

$$\mu_i = p_i - q_i \tag{2}$$

and

$$\sigma_i^2 = p_i + q_i - \mu_i^2. \tag{3}$$

Two prey model. Foraging takes place in continuous time. Food items are of two types. Type i items have handling time h_i, energetic value e_i and are found as a Poisson process of rate λ_i while searching. On finding an item an animal can either accept the item, in which case energy is gained but searching time is lost, or reject the item and immediately continue searching with no loss of time. One can think of there being three options available to an animal at a given time. Option 1: eat only Type 1 items; Option 2: eat both types of item; and Option 3: eat only Type 2 items. We assume that $e_1/h_1 > e_2 h_2$ so that Type 1 items are the more profitable. Under this assumption it turns out that it is essentially never optimal to use Option 3 and the choice at any time is between Options 1 and 2. We assume that metabolic expenditure is d per unit time. The mean net rates of gain under Options 1 and 2 are then:

$$\mu_1 = \frac{\lambda_1 e_1}{1 + \lambda_1 h_1} - d$$

and

$$\mu_2 = \frac{\lambda_1 e_1 + \lambda_2 e_2}{1 + \lambda_1 h_1 + \lambda_2 h_2} - d$$

respectively. The variance in gain per unit time under Options 1 and 2 are, respectively:

$$\sigma_1^2 = \frac{\lambda_1 e_1^2}{(1 + \lambda_1 h_1)^3}$$

and

$$\sigma_2^2 = \frac{\beta^2}{\alpha}\left[\lambda_1\left[\frac{e_1}{\beta} - \frac{h_1}{\alpha}\right]^2 + \lambda_2\left[\frac{e_2}{\beta} - \frac{h_2}{\alpha}\right]^2\right]$$

where $\alpha = 1 + \lambda_1 h_1 + \lambda_2 h_2$ and $\beta = \lambda_1 e_1 + \lambda_2 e_2$ (McNamara, unpublished manuscript).

A remark on randomized policies. The options considered in the above models of an animal's intake are not randomized. We could also have allowed an animal to choose an option with some probability p. It is, however, easy to show that an animal cannot improve its performance by inclusion of such randomized actions in its repertoire, and such actions are ignored.

Pure Survival Models

Overnight survival. Here we look at a single daytime period which ends at dusk (time T). We ignore the possibility that an animal could starve during the day and assume that it survives the following night if its reserves at dusk exceed a critical level x_c. An optimal policy maximizes overnight survival and thus maximizes $P(X(T) > x_c)$.

We first consider a constrained problem analysed by Stephens (1981). In Stephens' z-scores model an animal is constrained to choose one of two available options at the beginning of the foraging period (time 0) and cannot later switch to the other option. (Stephens does not state his assumptions in this manner but Houston and McNamara, 1982, show that these are his implicit assumptions.) The gain per unit time under Option i has mean μ_i and variance σ_i^2. Let x_0 be the reserves at time 0. Using a normal approximation the probability that an animal choosing Option i survives the night can be written as:

$$P(X(T) > x_c) = \Phi\left[\frac{x_0 + \mu_i T - x_c}{\sigma_i \sqrt{T}}\right] \tag{4}$$

where Φ is the normal distribution function, i.e.:

$$\Phi(z) = \int_{-\infty}^{z} \frac{1}{\sqrt{2\pi}} e^{-x^2/2}\,\mathrm{d}x \qquad -\infty < z < \infty.$$

The best choice of option maximizes $(x_0 + \mu_i T - x_c)/\sigma_i$. In the case $\mu_1 = \mu_2 = \mu$ and $\sigma_1 < \sigma_2$ it is thus optimal to choose Option 1 if and only if $x_0 + \mu T > x_c$. This result is known as the daily energy budget rule and can be "rephrased as "if the mean net gain is sufficient for overnight survival choose the low variance option, if the mean net gain is not sufficient choose the high variance option".

The unconstrained problem, in which an animal can switch freely between the options, can be analysed by modelling the energy reserves as a diffusion process. In technical tems this assumes that the change in energy reserves is given by the stochastic differential equation:

$$\mathrm{d}X(t) = \mu(u(X(t), t))\,\mathrm{d}t + \sigma(u(X(t), t))\,\mathrm{d}B(t),$$

where $B(t)$ is Brownian motion. Here a policy $u(x, t)$ specifies the option chosen in state x at time t. If the policy specifies that Option i is chosen then $\mu(u(x, t)) = \mu_i$ and $\sigma(u(x, t)) = \sigma_i$. McNamara (1983) derives the optimal policy when $\mu_1 = \mu_2 = \mu$ and $\sigma_1 < \sigma_2$. This policy can be given in terms of the switching line:

$$x + \mu(T - t) = x_c$$

as follows. While above the line; i.e. for reserves x satisfying $x > x_c - \mu(T - t)$, it is optimal to choose the low variance option, Option 1. While below the line it is optimal to choose the high variance option, Option 2. Above the line overnight survival seems probable and it pays the animal to play safe. Below the line overnight survival seems improbable and it pays the animal to take risks.

It can be seen that this optimal sequential rule has some formal similarities to the daily energy budget rule. In particular, under both rules Option 1 is chosen at time 0 if and only if $x_0 + \mu T > x_c$. But of course the rules are fundamentally different in that the constrained rule specifies one decision at time 0, whereas the optimal sequential rule specifies which option to choose as a function of how well the animal is doing during the foraging period.

One can analyse the advantage of sequential rules over constrained rules by comparing payoffs at time 0. Suppose $x_0 + \mu T = x_c$. Then the probability of survival under the optimal constrained rule is 0.5. McNamara (1983) shows that the probability of survival under the optimal sequential rule is $\sigma_2/(\sigma_1 + \sigma_2)$. This is greater than 0.5 and will be close to 1 if σ_2/σ_1 is large.

McNamara (1984) analyses the general case when $\mu_1 \neq \mu_2$ and $\sigma_1 < \sigma_2$ by modelling energy reserves as a diffusion process. In this case it is optimal for an animal with reserves x to choose Option 1 at time t if and only if:

$$x > x_c + \left[\frac{\mu_2\sigma_1 - \mu_1\sigma_2}{\sigma_2 - \sigma_1}\right](T - t). \tag{5}$$

Thus it is again optimal to choose Option 1 when reserves are high and Option 2 when reserves are low, but the switching line depends on all four parameters μ_1, σ_1, μ_2 and σ_2.

Clear analytic results seem possible only when reserves are modelled as a diffusion process. Nevertheless, the exact diffusion results appear to give good approximations to the optimal policy in other circumstances. Houston and McNamara (1985) present a model in which the food supply is described by the Two Prey Model. As before a fixed amount of energy x_c is required to survive the night. This model also includes the possibility of death from starvation during the foraging peiod, but including this feature has an insignificant effect on the form of the optimal policy when reserves are high and dusk is approaching. Houston and McNamara find the policy which maximizes survival probability by numerical computations.

When $e_1/h_1 > e_2/h_2$ the classical prey choice model based on rate maximization predicts that Type 2 prey items will not be included in the diet if $\lambda_1 e_1/(1+\lambda_1 h_1) > e_2 h_2$ (Stephens and Krebs, 1986). But Houston and McNamara (1985) show in their model that when this condition is met, it can still be optimal to include Type 2 items in the diet. Their computations reveal a switching line. When reserves are above this line it is optimal to choose both prey types (i.e. Option 2). When reserves are below this line it is optimal to take only the more profitable Type 1 prey items (i.e. Option 1). This result can be qualitatively understood in terms of risk sensitive behaviour as follows. Above the switching line overnight survival seems probable and it pays the animal to play safe and reduce the variance in food intake by taking all items, even though this means reducing the mean rate. Below the switching line the animal is in danger of starvation and must take only Type 1 items in order to maximize both the mean and variance in food gain. The position of the switching line can be understood quantitatively by approximating reserves by a diffusion process. Let the mean gains under Options 1 and 2 be μ_1 and μ_2, respectively, and let the corresponding variances in gain per unit time be σ_1^2 and σ_2^2, respectively. Then the switching line found numerically by Houston and McNamara (1985) is in close agreement with that given by equation (5).

McNamara (unpublished) analyses overnight survival using the Discrete Time and State Model of the food supply. If the variance in reward magnitude is large (i.e. $p_i + q_i$ close to 1 for some option) it is possible to obtain solutions which manifest periodic effects even when the time to go till final time T is large. These are artifacts due to the discrete grid on which computations are performed and disappear when the variance is reduced. One then obtains a switching line above which the lowest variance option is used and below which the highest variance option is used. Again the switching line agrees closely with the predictions of equation (5). The special case when there are two options giving the same mean gain is investigated by McNamara and Houston (1986). They consider the case where the terminal reward at dusk is not a step function in reserves but a smoothed step function. The switching line then has slope equal to the common mean. Thus equation (5) correctly predicts the optimal policy even when the terminal reward is not a simple step function. A modified version of the Discrete Time and State Model is investigated numerically by Houston and McNamara (1986).

Lower lethal boundary. The previous model ignored death from starvation during foraging and concentrated on death overnight. Here we reverse the emphasis and assume that an animal can forage continually; death now occurs if the animal's reserves ever fall to zero during foraging. In this section we are interested in the policy which best allows an animal to escape from the lower lethal boundary. To be more precise we assume that the animal's energy reserves are unbounded and concentrate on the case where all foraging options yield the same mean net rate of gain μ. If $\mu \leqslant 0$ then an animal is sure to starve. We thus assume $\mu > 0$. Then an animal will either starve or its reserves will tend to infinity as time tends to infinity. We find the policy which minimizes the probability of starvation. This model thus looks at a one-off escape from the lower lethal boundary. Such a model would be appropriate for an animal with low reserves, which nevertheless has a large capacity to carry reserves, and is effectively safe from starvation if it can build up reserves to its full capacity. When an animal's capacity is small the criterion of a single escape is no longer appropriate and one must introduce a different optimality criterion. This is done in the section on long term survival.

Uninterrupted foraging. We here assume that food sources are always available, although the amount of food obtained from these sources is stochastic. We begin by looking at the Discrete State and Time Model. In this model if Option i is chosen at time t, reserves one time unit later have decreased by one unit with probability q_i and increased by one unit with probability p_i. If this option is used at every level of energy reserves the probability of ever starving given initial reserves of x is:

$$\phi_i(x) = \begin{cases} (q_i/p_i)^x & \text{if } q_i < p_i \\ 1 & \text{if } q_i \geqslant p_i. \end{cases}$$

The best fixed (i.e. state independent) policy is thus to choose the option minimizing q_i/p_i. Merad (1991) shows that this fixed policy is also the optimal state dependent policy. If all options have the same mean net gain $\mu = p_i - q_i$, then minimizing q_i/p_i is equivalent to minimizing the variance in the option chosen. Thus risk-averse behaviour is optimal at all levels of reserves.

When the food suply is described by the Two Prey Model, Houston and McNamara (1985) show that it is optimal to take every item for which the net energy gain $e - dh$ is positive.

The rate at which food can be gained is unimportant in both the models described here, all that is important is that reserves increase rather than decrease. As we now describe, this conclusion need not hold if the foraging animal is uncertain to gain access to a food source.

Interrupted foraging. Events such as bad weather or the presence of a predator can mean that an animal is unable to forage for some time. To incorporate this feature in the Discrete Time and State Model we assume that periods of uninterrupted foraging have a geometric length and are interspersed by periods when the animal cannot forage and loses one unit of energy per unit time. Merad (1991) has analysed this modified model. He considers the case where interruptions also last a geometric number of time intervals, and shows that it is still optimal to choose the option with the lowest variance at all levels of reserves. When interruptions do not have a geometric distribution, however, minimizing variance may not be optimal and the optimal choice of action may depend on reserves.

Barnard *et al.* (1985) analyse a similar model to that considered by Merad. They present an example in which interruptions last a negative binomial number of time units. Two feeding options are available. In one case these options yield the same mean net gain. The optimal policy in this example is to choose the lower variance option when reserves are very low, choose the higher variance option at intermediate reserves, and choose the lower variance option again at high reserves. This policy can be understood as follows. At low reserves the animal is in imminent danger of starvation even if there is no interruption and it pays to play safe. At slightly higher reserves the main danger is that an interruption will occur. The animal thus takes risks in order to get its reserves above the level at which an interruption is liable to kill it. For high reserves the main concern is not to allow reserves to drift down and the lower variance option is optimal.

Long term survival. The model with no limit on energy reserves, in which an animal with sufficiently high reserves is almost sure to survive, is not appropriate for a small overwintering bird. This is because a small bird's maximum level of reserves is only sufficient for a few days survival. To analyse survival strategies of the bird we must include its limited ability to carry reserves in a model, and introduce a criterion for survival over an extended period.

We assume that the foraging process is time homogeneous, so that there is no day/night cycle. An animal starves if its reserves fall to zero. There is a fixed upper limit L on reserves, and food which would have carried reserves above L is lost. The food supply is described by the Discrete State and Time Model and is not interrupted. By taking L to be an integer we can restrict attention to integer levels of energy reserves x in the range $x = 0, 1, \ldots, L$.

We consider two seemingly different ways to look at survival. Although the approaches differ it turns out that both lead to the same optimality criteria for long term survival. Here we present the ideas behind the two approaches using a purely verbal argument. Precise mathematical formulations of the approaches and resulting theorems are deferred to Appendix 1.

Approach 1. Equilibrium mortality rates. A stationary behavioural policy π is a rule specifying which action to choose as a function of energy reserves but not as a function of time. We wish to introduce the idea of an equilibrium rate of mortality for an animal that follows this policy. We motivate this by considering a large group of animals independently following the same policy

π. Over time, numbers in the group will decline as animals die of starvation. At first the numbers starving in each time interval depend strongly on the initial reserves of group members. But eventually an equilibrium will be reached in which the proportion of animals alive at time t which are still alive at time $t+1$ is a constant, λ_π, which depends only on the policy π used. For large t, a particular animal in this group is alive at time t with probability $C_\pi \lambda_\pi^t$, where the constant C_π depends on the policy π and the animal's initial level of reserves. We can compare the animal's probability of survival under two different policies π_1 and π_2, where $\lambda_{\pi_1} < \lambda_{\pi_2}$. It can be seen that if t is sufficiently large:

$$C_{\pi_1}\lambda_{\pi_1}^t < C_{\pi_2}\lambda_{\pi_2}^t.$$

Thus in the long term, policy π_2 will be better, even though it may be disadvantageous in the short term. Motivated by these considerations, we define a policy to be optimal in the long term if it maximizes λ_π, or equivalently minimizes the equilibrium mortality rate $1-\lambda_\pi$.

Approach 2. Survival till a fixed time. The time dependent policy which maximizes the probability of survival until some fixed time T can be found by dynamic programming. This policy specifies the action taken at each level of energy reserves and time. When the time-to-go till final time T is small the best choice of action may depend strongly on time-to-go as well as energy reserves. As the time-to-go increases the best choice of action tends to a limiting value which depends only on energy reserves. Thus in the limit we obtain a stationary policy π^*. We can think of π^* as the policy maximizing "long term" survival and regard it as an optimal policy.

Fortunately these two approaches lead to the same optimality criterion since it can be shown that λ_π is maximized when $\pi = \pi^*$. Furthermore, as is explained in Appendix 1, the policy π^* can be further characterized by an eigenvalue equation for the optimality or dynamic programming operator.

McNamara (1990) discusses these characterizations and proves their equivalence. In his model the food supply is described by the Discrete State and Time Model, but there may be any finite number of feeding options rather than just two. These options may differ from one another in mean gain as well as variance. He gives an analytic proof of a number of qualitative results describing the form of the optimal policy. The main results are:

(i) If one of the feeding options gives a positive mean net gain (i.e. $\mu_i > 0$ for some option i), then an animal's expected future reproductive success (EFRS) is a concave function of reserves at all reserves. Thus it is always optimal to be risk-averse.

(ii) If all options yield a negative mean net gain (i.e. $\mu_i < 0$ for all options i), then the EFRS is concave for high reserves, but may be convex for low reserves. Consequently it is optimal to be risk averse at high reserves, but may be optimal to be risk-prone at low reserves.

These results are discussed and illustrated by McNamara and Houston (1990a). Further illustrations can be found in McNamara and Houston (1990b).

Daytime and night-time starvation. We now consider models which allow both starvation during foraging, and have a day/night cycle, so that starvation can occur overnight. The policy which maximizes overnight survival probability does not maximize daytime survival probability and vice versa. As we shall see, the policy that maximizes overall survival probability is a compromise between these two simple objectives.

Houston and McNamara (1985) consider a single day during which the food supply is described by the Two Prey Model, followed by night during which an animal rests, decreasing its reserves by x_c. The animal dies if reserves reach zero during the day or are below x_c at dusk. Houston and McNamara compute the policy which maximizes the probability of survival till the following dawn. They present an example for which the classical rate maximization model predicts exclusive choice of the more profitable Type 1 prey items. When there is a lot of time left till dusk the policy which maximizes survival probability prescribes that both prey types should be taken at all levels of energy reserves. This can be understood in terms of the model of escaping the lower lethal boundary: when there is a lot of time left an animal can concentrate on not starving now. As dusk approaches the optimal policy depends on reserves. For very low reserves it is always optimal to take any prey item in order to avoid immediate starvation. As reserves increase it becomes optimal to take Type 1 items alone and then becomes optimal to take both types for large reserves. These effects can be understood in terms of surviving the night, and have already been explained.

The above model desscribes survival over a 1 day period. One can also analyse long term survival when there is a day/night cycle and reserves are bounded by an upper limit L. Suppose that foraging starts at dawn (time 0) and ends at dusk (time T) on each day. Overnight a possibly random amount of energy reserves are used up. We now require an optimality criterion which is appropriate for long term survival when there is a day/night cycle. As before we give a purely verbal description, deferring mathematical proofs to an appendix (Appendix 2). We consider survival over an extended period of many days such as winter. This involves using the optimal policy on one day to find the optimal policy on the preceding day; a procedure which Mangel and Clark (1988) refer to as sequential coupling. Working backwards from final time in this way one can find the policy which maximizes survival till the end of winter. On a given day the best choice of action will depend on reserves, time of day and the number of days-to-go till the end of the period. As the number of days-to-go increases, the policy over a day settles down to a limiting policy which depends only on reserves and time of day. We can regard this daily policy as the optimum policy which maximizes long-term survival.

Results from models of long term survival when there is a day/night cycle (e.g. McNamara and Houston, 1986; Houston and McNamara, 1986) show the same general features as were found by Houston and McNamara (1985) when analysing survival over a single night: there is a wedge-shaped region near dusk in which the animal should be risk-prone.

A comparison of the predictions of the various survival models is given in Table 1.

Table 1. A summary of the predictions of the survival models presented

Model	Reserves	Risk-averse	Risk-prone
(a) Overnight survival	Low		√
(z-scores and sequential)	High	√	
(b) Escape from lower lethal boundary ($\mu>0$)			
(i) No interruptions	All reserves	√	
(ii) Interruptions	Very low	√	
	Low		√
	High	√	
(c) Long term survival (no interruptions or day/night)			
(i) $\mu>0$	All reserves	√	
(ii) $\mu<0$	Low		√
	High	√	
(d) Long term survival with day/night ($\mu>0$)			
(i) Near dusk	Very low	√	
	Low		√
	High	√	
(ii) Near dawn	All reserves	√	

Starvation and Reproduction. In the models discussed so far the optimality criterion has been maximization of survival probability. This is an appropriate criterion for, say, a small bird in winter. The bird must survive if it is to reproduce in the future, and food items obtained enhance future reproduction through their effect on survival probability (McNamara and Houston, 1982). Food items may, however, enhance reproduction directly. In such cases we must seek new optimality criteria. In this section we analyse a model in which an animal uses food energy both for immediate survival and reproduction. There is then a conflict between avoiding starvation, and hence increasing future reproductive success, and increasing immediate reproductive success.

Food is described by the Discrete Time and State Model. The foraging process is time homogeneous, so that there is no day/night cycle, and is not interrupted. An animal starves if its energy reserves fall to zero. There is an upper limit L on reserves as in the long term survival model, but we now suppose that food which would have taken reserves above L is converted directly into immediate reproductive output. We introduce a background mortality due to predation, accident or simply death from natural causes which kills an animal with probability $1-\theta$ in each time interval $(0<\theta\leqslant 1)$. Death from the background mortality is independent of the animal's energy reserves and the option chosen. The level of background mortality has no effect on the optimal policy when the optimality criterion is pure survival. As we will demonstrate below it can have a strong effect when an animal is also reproducing. The animal continues to forage and reproduce until it dies either from starvation or the background mortality. We look at the policy which maximizes an animal's lifetime reproductive success.

This problem is an example of an infinite horizon positive programming problem of Markov decision theory (e.g. Whittle, 1983), and has been analysed by Merad (1991). Let $\mu_1=\mu_2=\mu$. Merad shows that when $\mu<0$ it is optimal to be risk-prone at all levels of reserves for all values of θ. This can be compared with the policy which maximizes long term survival, in which it is optimal to be risk-prone at low reserves and risk-averse at high reserves. When $\mu>0$ and there is no background mortality $(\theta=1)$ in the reproduction model, it is optimal to be risk-averse at all levels of reserves. When $\mu>0$ and $\theta<1$ it is optimal to be risk-averse at low reserves and risk-prone at high reserves. In contrast the long term survival model predicts risk-averse behaviour at all levels of reserves. As θ decreases, the range of reserves over which it is optimal to be risk-prone in the reproduction model increases. McNamara *et al.* (1991) give numerical examples comparing the prediction of the long-term survival model and the above reproduction model.

Table 2 compares the predictions of the survival model with those of the reproduction model. Figure 1b shows examples of the relative fitness function for the reproduction model.

Table 2. A comparison of the model based on long term survival with the model based on lifetime reproductive success

	Positive mean		Negative mean	
	Low reserves	High reserves	Low reserves	High reserves
(a) Survival model	Averse	Averse	Prone	Averse
(b) Reproduction model ($\theta=1$)	Averse	Averse	Prone	Prone
(c) Reproduction model ($\theta<1$)	Averse	Prone	Prone	Prone

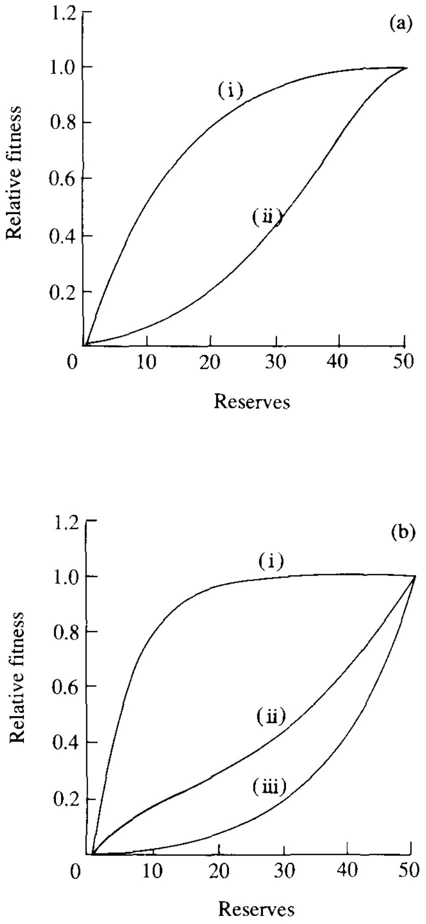

Figure 1. (a) The starvation model. Let $\psi^*(x)$ be the maximum probability an animal with reserves x survives to some given time in the distant future. Then the figure shows the relative fitness $\psi^*(x)/\psi^*(L-1)$ for: (i) $\mu=0.01$; and (ii) $\mu=-0.02$. (b) The reproduction model. Let $\psi^*(x)$ be the maximum expected lifetime reproductive success for an animal with reserves x. Then the figure shows the relative fitness $\psi^*(x)/\psi^*(L-1)$ for: (i) $\mu=0.025, m=0$; (ii) $\mu=0.025, m=0.013$; and (iii) $\mu=-0.015$, $m=0$. $\sigma_1^2=0.3$, $\sigma_2^2=0.4$, $L=51$ throughout.

Variability in Delay. So far, we have only discussed variability in the amount of energy associated with a foraging option. In general, foraging options will also involve delays before food is obtained. McNamara and Houston (1987) present a general model in which both the energetic gain and the associated delay are random variables. The model is based on a single choice between options, after which the animal follows a background foraging process that has a mean net rate of energetic gain γ, with variance σ^2. When there is no variance in delay then (as we have already seen) the optimal action is to maximize the variance in gain if $\partial^2 f/\partial x^2 > 0$ and to minimize the variance in gain if $\partial^2 f/\partial x^2 < 0$. When there is no variance in gain, then the optimal action is to maximize variance in delay if $\partial^2 f/\partial t^2 > 0$ and to minimize the variance in delay if $\partial^2 f/\partial t^2 < 0$. McNamara and Houston illustrate these conditions by considering behaviour over a single day, at the end of which the bird survives the night if its reserves are above the critical level x_c. Risk-sensitivity in terms of reserves is given by the expected daily energy budget rule, i.e. prefer variability in gain if and only if the mean rate of gain is not sufficient to get reserves above x_c at the end of the day. In a space with axes reserves when choice is made and time left until the end of the day, this condition gives us two regions separated by a straight line with slope equal to the mean gain γ. Below the line it is optimal to be risk-prone in terms of reserves, above the line it is optimal to be risk-averse in terms of reserves. Risk-sensitivity in terms of time has a different pattern. There are typically four regions in the space of reserves versus time: in two of these regions it is optimal to be risk-prone in time, in the other two it is optimal to be risk-averse in time. In some regions it is optimal to be risk-prone in time but risk-averse in reserves, while in other regions it is optimal to be risk-averse in time and risk-prone in reserves.

Zabludoff *et al.* (1988) present a model in which one option has a constant delay of magnitude d and the other has a variable delay which has magnitude $d-\delta$ with probability $\frac{1}{2}$ or $d+\delta$ with probability $\frac{1}{2}$. The animal can forage for a total of T time units and must get its reserves above x_c by the end of this time if it is to survive. Let the initial level of reserves be x_0 and all items have unit value. On the assumption that all the time is devoted either to the constant option or to the variable option, Zabludoff *et al.* find the decision that maximizes survival probability. One can analyse their model as follows. Let:

$$(x_c - x_0)/T > 1/d.$$

Then an animal using the constant option is sure to starve and it is optimal to choose the variable option. Let:

$$(x_c - x_0)/T < 1/d.$$

Then an animal using the constant option is sure to survive and it is optimal to use this option.

This is a temporal version of the daily energy budget rule.

It is not clear that this result will hold when both options are variable in time. We have carried out calculations of the optimal sequential policy using dynamic programming. Our results suggest that the rule: "take the low variance option if and only if the mean gain is sufficient for survival", gives a very good approximation to the optimal policy. We have also considered cases in which the amount of food obtained is a random variable that has the same distribution for both options. When the amount obtained is always non-negative, the above sequential policy gain appears to be approximately optimal. When there is the possibility that the net effect of obtaining an item is to decrease reserves then our computations reveal that the optimal policy has a totally different form. At a given time, there are now four regions of reserves. As reserves increase the pattern is "choose high variance, choose low variance, choose high variance, choose low variance". McNamara and Houston (1987) looked at the best single choice given a background foraging process and found a similar decomposition in the space of reserves and time.

There is another, totally different, reason why it can be advantageous to prefer variability in delays. Future food items may not be obtained, either because they are taken by other animals or because the foraging process is interrupted. McNamara and Houston (1987) analyse the consequences of this source of uncertainty. Suppose that items become unobtainable at a rate α per unit time. If there is a delay, h, before an item can be taken, then the item will be available with probability $e^{-\alpha h}$. If the delay H is a random variable, then the probability that the item is available is $\mathbb{E}(e^{-\alpha H})$. We compare an option with a fixed delay h with an option with a variable delay H whose mean is equal to h. By Jensen's inequality:

$$\mathbb{E}(e^{-\alpha H}) \geqslant e^{-\alpha h}.$$

Thus the variable delay gives the greater probability of obtaining the item.

Selection Pressure. McNamara and Houston (1986) introduce the canonical cost as a measure of the cost of deviating from an optimal policy. Let an animal in state x at time t perform action a and then follow the optimal policy. Then $c(x, a, t)$ is the loss in expected future reproductive success that results from performing action a rather than the optimal action. It follows that if a is the optimal action then $c(x, a, t) = 0$. The canonical costs provide a common currency for comparing actions. They also give some indication of the robustness of an optimal policy, in the sense that if canonical costs are small then deviations are not costly. They do not, however, supply us with sufficient information to evaluate the strength of selection on behaviour. The cost of a suboptimal action in a given state at a given time may be very high, but if the animal is very unlikely to be in this state at the relevant time then the selection pressure in favour of the optimal decision will be small. Houston and McNamara (1986) represent this idea by defining the selection pressure $s(x, a, t)$ to be the product of the canonical cost $c(x, a, t)$ and the probability $p(x, t)$ of being in state x at time t. Houston and McNamara consider a model in which there is a day/night cycle and an upper limit to an animal's energy reserves, using maximization of long term survival probability as the optimality criterion. They illustrate $c(x, a, t)$ and $s(x, a, t)$ in the case where all options give the same mean net energy gain while foraging and the mean net gain over a 24 hr period is positive. Within the wedge-shaped region in which it is optimal to prefer the high variance option, the canonical cost of taking the low variance option can be high. It is unlikely, however, that the animal will be in this region because an animal in the region would be very likely to die overnight. Thus $p(x, t)$ is small and the selection pressure is relatively low. In general, the canonical cost of taking the low variance option when it is optimal to take the high variance option can be comparable in magnitude to the canonical cost of taking the high variance option when it is optimal to take the low variance option, whereas the selective pressure in favour of risk-prone behaviour tends to be much less than the selective pressure in favour of risk-averse behaviour.

As well as looking at the selective pressure, Houston and McNamara (1986) present an analysis of the advantage of the optimal policy that is based directly on the resulting mortality. Imagine that the animal can choose between two actions a_1 and a_2. Each has the same mean net energy gain but a_1 has the lower variance. The mortality m^* under the optimal state-dependent policy can be compared with the mortality m_i that results when the animal can only adopt action i. Houston and McNamara find that m_1 is only slightly greater than m^*, but m_2 is much greater than m_1. Thus an animal which uses only action a_1 can slightly reduce its rate of mortality by choosing the higher variance action a_2 when it is in imminent danger of starving during the coming night (i.e. in the wedge shaped region). An animal which uses action a_2 can drastically reduce its rate of mortality by choosing the lower variance action a_1 on all occasions, and can reduce mortality slightly further by choosing a_1 on all occasions except in the wedge shaped region near dusk.

The above results were obtained for an environment that has a day/night cycle. We have seen that when there is no day/night cycle or interruptions to the foraging process, the qualitative form of the optimal policy is different. When all options have the same positive mean net gain μ, it is always optimal to take the option with the lowest variance. In this case there is never a selective advantage to taking the high variance option, and, as Fig. 2a illustrates, doing so may result in a significant increase in mortality. As described above, when the common mean μ is negative it is always optimal to choose the low variance option at high reserves, and it may be optimal to choose the high variance option at low reserves. McNamara and Houston (1990a) give some examples of selection pressure in this case. We now look at this topic in more detail using the Discrete State and Time Model of the food supply.

To start with, we assume that an animal has a range of foraging environments that differ in their variability. Once the animal chooses an environment, it cannot later switch to another environment. To investigate

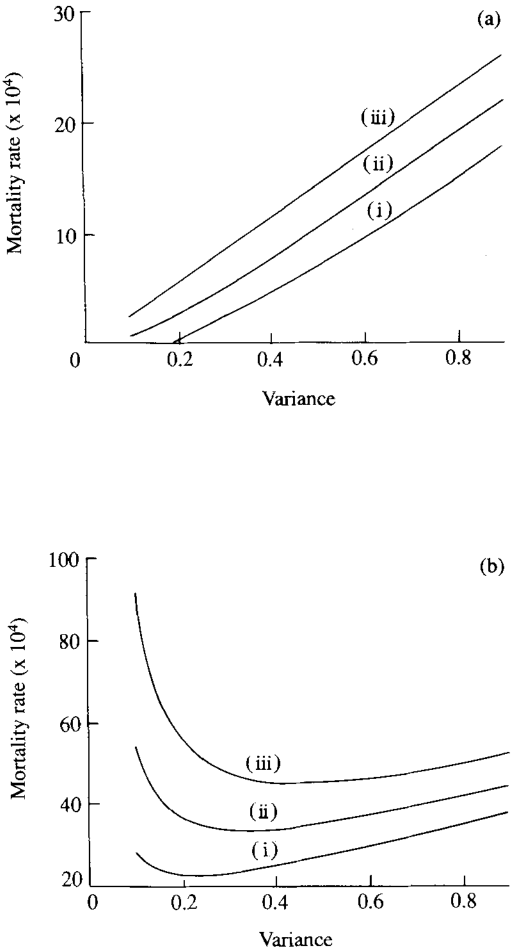

Figure 2. The equilibrium rate of mortality as a function of the variance in the food supply when an animal cannot switch between options. In (a) (i) $\mu=0.02$; (ii) $\mu=0.01$; (iii) $\mu=0.001$. In (b) (i) $\mu=-0.02$; (ii) $\mu=-0.03$; (iii) $\mu=-0.04$; $L=20$ throughout.

selection pressure on choice, we find how the rate of mortality depends on the variance of the environment that is chosen. Figure 2b illustrates this dependence for three values of the mean μ. For each value of μ, mortality is minimized at some intermediate level of variance. This best level of variance increases as μ decreases. It can be seen that in all cases the mortality rate depends strongly on the variance of the environment that is chosen.

To look at the advantage of state-dependent decision, we now assume that the animal can switch freely between two options. Table 3 gives the rate of mortality under an optimal state-dependent policy as a function of the options available. It can be seen that it is always advantageous to have a choice of options rather than a single option, and the advantage of choice increases as the options become more diverse in their variance.

Table 3. Mortality rates under different combinations of foraging options. There is no day/night cycle or interruptions to the foraging processes. Food intake is described by the Discrete Time and State Model with the two feeding options giving the same mean net energy gain

		σ_1^2				
(a)	σ_2^2	0.1	0.3	0.5	0.7	0.9
	0.1	27.9				
	0.3	17.7	23.0			
	0.5	16.0	21.5	27.0		
	0.7	15.2	20.9	26.5	32.1	
	0.9	14.8	20.6	26.2	31.8	37.5

		σ_1^2				
(b)	σ_2^2	0.1	0.3	0.5	0.7	0.9
	0.1	90.9				
	0.3	42.7	47.4			
	0.5	34.5	39.9	45.0		
	0.7	31.2	36.9	42.2	47.5	
	0.9	29.4	35.3	40.8	46.2	51.6

Table entries show equilibrium rates of mortality $\times 10^4$ [i.e. $(1-\lambda_{\pi^*})\times 10^4$] for various combinations of variances in food intake under the two options. In (a) the common mean is $\mu=-0.02$, in (b) $\mu=-0.04$. $L=20$ throughout.

We can summarize the results shown in Fig. 2 and Table 3 as follows. If an animal cannot switch between options, then the selective advantage associated with making the correct choice is strong. In the negative mean case, the animal can get an additional advantage by being able to switch between options that are significantly different.

Discussion. Our discussion of variability in the amount of food that an animal obtains from a given foraging option has been based on the second moment, i.e. the variance. It is also possible to consider the third moment, i.e. the skew. Caraco and Chasin (1984) provide a general account of the effect of skew based on a Taylor series expansion of the fitness function. They also present a model based on a single choice between two distributions with the same mean and variance, but one has positive skew and the other has negative skew. If the animal has to get its reserves to some critical level x_c in order to survive, then it can be optimal to prefer positive skew when reserves are low or high and to prefer negative skew when reserves are intermediate. Houston and McNamara (1986) show that this form of policy is also found in a dynamic model. Houston and McNamara (1986) also give an analytic treatment of the effect of skew when foraging is not interrupted and there is no upper limit on the animal's energy reserves.

Although the model that includes reproduction can capture some aspects of growth (see McNamara *et al.*, 1991, for discussion) the remainder of our models ignore this possible benefit of energy. Houston and McNamara (1990) investigate some effects of environmental variability on the fitness of an organism that grows for a period of time in a given habitat. In contrast to most of the models that we have discussed, there is only one decision. This is made by a mother in choosing the site in which her offspring grows. The quality of the site in terms of growth is characterized by a parameter w that is unknown to the mother, and so is considered to be a random variable W. Given a range of habitats with the same expected value of W, we find how the variance in W should determine choice of habitat. When there is no predation, then maximizing expected size at a fixed time can result in either choosing the site with the highest variance or the site with the lowest variance, depending on the function that determines growth. Minimizing the expected time to reach a given size always makes it advantageous to choose the site with the lowest variance. When there is size-dependent predation, the results are not quite so clear-cut.

There is scope for further work on the implications of growth for risk-sensitivity. The behaviour of a single growing animal could be analysed as a series of size-dependent decisions. In a model of a parent feeding dependent offspring, energy would be important both in terms of growth and in terms of avoiding starvation.

A different concept of growth is relevant when we model the foraging behaviour of members of a colony of social insects. For an annual colony (e.g. bumble bees) fitness can be evaluated in terms of the number of reproductives produced at the end of the season. Oster and Wilson (1978) argue that there should be risk-sensitive effects in the context of colony growth, but this idea has not been followed up to produce a model that relates risk sensitive foraging decisions to the number of reproductives produced.

We have consistently talked about animals in this paper, but one can apply the theory of risk sensitive behaviour to analyse plant strategies. For further discussion see Caraco and Kelly (1991).

We have tried to emphasize the variety of possible models of risk sensitive foraging. The common feature of the models is an attempt to relate food gain to reproductive success. We have not reviewed models that do not have a direct interpretation in terms of fitness (e.g. models of variance discounting, Real, 1980). Our summary of the predictions of the models (see Tables 1 and 2) makes it clear that there is not a single simple and universal prediction about risk-sensitive foraging. We hope that this summary will encourage people to test a model of risk-sensitive foraging that is appropriate for the animal under study.

We thank John Hutchinson and Marc Mangel for their comments on earlier versions of this paper. A.I.H. was supported by the Natural Environment Research Council.

APPENDIX 1

The Criterion for Long Term Survival. A stationary policy is a rule specifying which option to choose as a function of energy reserves. We adopt the notation that under policy π the animal chooses Option $\pi(x)$ when its reserves are x. For x, y lying between 1 and L we can define $a_{xy}(i)$ to be the probability that an animal with reserves x at time t has reserves y at time $t+1$ if it chooses Option i at time t. A stationary policy π then defines an $L \times L$ matrix A_π whose (x,y)th component is $a_{xy}(\pi(x))$. The matrix A_π determines the time evolution of reserves as follows. Suppose that an animal has reserves x at time 0 with probability α_x. Then the probability it has reserves y at time t is $(\boldsymbol{\alpha}A_\pi^t)_y$, where $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_L)$. Let λ_π be the Perron–Frobenius eigenvalue of A_π (see, for example, Gantmacher, 1959) and let $\boldsymbol{\alpha}_\pi$ be the corresponding left eigenvector, so that $\boldsymbol{\alpha}_\pi A_\pi = \lambda_\pi \boldsymbol{\alpha}_\pi$. We normalize the eigenvector $\boldsymbol{\alpha}_\pi$ so that its components sum to 1. Then for any $\boldsymbol{\alpha}$ we have:

$$\lim_{t\to\infty} \lambda_\pi^{-t}\boldsymbol{\alpha}A_\pi^t = K\boldsymbol{\alpha}_\pi \tag{A1}$$

where K is a constant depending on $\boldsymbol{\alpha}$. Thus for large t the distribution of reserves conditional on the animal being alive is given by $\boldsymbol{\alpha}_\pi$ and the probability that an animal alive at t is still alive at $t+1$ is λ_π. We may interpret $1-\lambda_\pi$ as the equilibrium mortality rate under policy π, and define a policy to be optimal if it maximizes λ_π.

The above criterion characterizes an optimal policy by following reserves forward in time. Another natural optimality criterion is based on working backwards from some final time using dynamic programming. We now describe this second approach and show how the two criteria are related.

The optimality operator A^* is defined as follows. Let f be a non-negative real-valued function on the set of energy levels $\{1, \ldots, L\}$. Then A^*f is defined to be the function on $\{1, \ldots, L\}$ whose value at x is given by:

$$(A^*\mathrm{f})(x) = \underset{i}{\mathrm{maximum}} \sum_{y=1}^{L} a_{xy}(i)\mathrm{f}(y), \tag{A2}$$

where the maximum is taken over all options i. It can be seen that this is the one step dynamic programming equation. Let R be a non-negative valued function on $\{1, \ldots, L\}$, and consider the problem of maximizing the expected value of R at some final time T. Let $n = T - t$ be the number of time units to go until T. Then the maximum expected value of R for an animal which has reserves x at time t is $(A^{*n}R)(x)$. An important special case is when $R(x) = 1$ for all $x = 1, \ldots, L$. In this case $(A^{*n}R)(x)$ is the maximum probability of survival until time T.

Now consider the eigenvalue equation:

$$A^*\phi = \lambda^*\phi \tag{A3}$$

where $\lambda^* \geqslant 0$ and ϕ is a non-negative and non-zero function on $\{1, \ldots, L\}$. It can be shown that this equation always has a solution (Kennedy, 1978). Furthermore, given our assumptions that $p_i, q_i > 0$, λ^* is unique and ϕ is unique up to multiplication by a positive constant (McNamara, 1990). For any terminal reward function R it can be shown that there is a constant C depending on R such that:

$$\operatorname*{limit}_{n \to \infty} \lambda^{*-n} A^{*n} R = C\phi \tag{A4}$$

(McNamara, 1990). Thus the components of ϕ give the asymptotic relative value of the various energy levels since

$$\frac{\phi(y)}{\phi(x)} \simeq \frac{(A^{*n}R)(y)}{(A^{*n}R)(x)} \tag{A5}$$

for large n and for *any R*.

The optimal option to choose when reserves are x with n steps to go maximizes:

$$\sum_{y=1}^{L} a_{xy}(i)\,(A^{*(n-1)}R)(y) \tag{A6}$$

as a function of i. Motivated by this we define a stationary policy π^* whose choice of option $\pi^*(x)$ for reserves x is given by:

$$\sum_{y=1}^{L} a_{xy}(\pi^*(x))\phi(y) = \operatorname*{maximum}_{i} \sum_{y=1}^{L} a_{xy}(i)\phi(y). \tag{A7}$$

It follows from equations (A4) and (A7) that whatever the terminal reward R, $\pi^*(x)$ is asymptotically optimal choice of option for reserves x, and is the optimal choice for all sufficiently large n provided equation (A7) defines π^* uniquely. We may thus think of π^* as defining an optimal stationary policy for long term survival.

To link the backward and forward approaches we note that:

$$\lambda^* = \lambda_{\pi^*} = \operatorname*{maximum}_{\pi} \lambda_\pi, \tag{A8}$$

where the maximum is taken over all stationary policies π (McNamara, 1990). Thus both optimality criteria lead to the same optimal stationary policy π^*.

APPENDIX 2

Long Term Survival when there is a Day/Night Cycle. We introduce an operator A^* which maps functions of state at dusk to function of the state 24 hr before as follows. Let R be a terminal reward function defined on states at dusk on some day. By dynamic programming back through the day one can find the maximum expected reward for all possible states at dawn on that day. The maximum expected reward for an animal with reserves x the preceding dusk can then be found from the overnight energy loss. We define $(A^*R)(x)$ to be this expectation.

Now consider an extended period of many days, such as a winter. Define a terminal reward R_0 on states at dusk on the last day in this period. Set $R_n = A^{*n}R_0$. Then $R_n(x)$ is the maximum expected reward for an animal which has reserves x with n days to go. We can regard the policy adopted while foraging over a daytime period as a single option, and regard options as being chosen for the next day at each dusk. The results applied previously to long-term survival can then be applied to A^* to deduce that:

$$\operatorname*{limit}_{n\to\infty} \lambda^{*-n} R_n = C\phi$$

where C is a constant and

$$A^*\phi = \lambda^*\phi.$$

Without loss of generality we can rescale ϕ so that $\phi(L)=1$. Setting $\tilde{R}_n(x)=R_n(x)/R_n(L)$ we have then:

$$\operatorname*{limit}_{n\to\infty} \tilde{R}_n = \phi.$$

R_n is defined on states at dusk with n days to go and determines the optimal foraging policy during the daytime period preceding this dusk. Replacing R_n with a positive multiple of it will not alter the optimal policy. Since $\tilde{R}_n$ converges to ϕ, the optimal policy over one day converges to a limiting policy $\pi^*(x, t)$ determined by the dusk terminal reward ϕ. $\pi^*(x, t)$ is now a function of both reserves x and time of day t.

LITERATURE

Barnard, C. J., C. A. J. Brown, A. I. Houston and J. M. McNamara. 1985. Risk-sensitive foraging in common shrews: an interruption model and the effects of mean and variance in reward rate. *Behav. Ecol. Sociobiol.* **18**, 139–146.

Caraco, T. 1981. Risk-sensitivity and foraging groups. *Ecology* **62**, 527–531.

Caraco, T. and M. Chasin. 1984. Foraging preferences: response to reward skew. *Anim. Behav.* **32**, 76–85.

Caraco, T. and C. Kelly. 1991. On the adaptive value of physiological integration in clonal plants. *Ecology* **71**, 81–93.

Clark, C. W. 1987. The lazy adaptable lions: a Markovian model of group foraging. *Anim. Behav.* **35**, 361–368.

Clark, C. W. and M. Mangel. 1984. Foraging and flocking strategies: information in an uncertain environment. *Am. Nat.* **123**, 626–641.

Clark, C. W. and M. Mangel. 1986. The evolutionary advantages of group foraging. *Theor. Popul. Biol.* **12**, 45–75.

Charnov, E. L. 1976. Optimal foraging: attack strategy of a mantid. *Am. Nat.* **110**, 141–151.

Ekman, J. and B. Rosander. 1987. Starvation risk and flock size of the social forager: when there is a flocking cost. *Theor. Popul. Biol.* **31**, 167–177.

Ellner, S. and L. R. Real. 1989. Optimal foraging models for stochastic environments: Are we missing the point? *Comments Theor. Biol.* **1**, 129–158.

Gantmacher, F. R. 1959. *Matrix Theory*. New York: Chelsea Publishing Company.

Houston, A. I. and J. M. McNamara. 1982. A sequential approach to risk-taking. *Anim. Behav.* **30**, 1260–1261.

Houston, A. I. and J. M. McNamara. 1985. The choice of two prey types that minimises the probability of starvation. *Behav. Ecol. Sociobiol.* **17**, 135–141.

Houston, A. I. and J. M. McNamara. 1986. Evaluating the selection pressure on foraging decisions. In *Relevance of Models and Theories in Ethology*, R. Campan and R. Zayan (Eds), pp. 61–75. Toulouse, France: Privat.

Houston, A. I. and J. M. McNamara. 1990. The effct of environmental variability on growth. *Oikos* **59**, 15–20.

Kennedy, D. P. 1978. On sets of countable non-negative matrices and Markov decision processes. *Adv. Appl. Prob.* **10**, 63–646.

Mangel, M. and C. W. Clark. 1988. *Dynamic Models in Behavioral Ecology*. Princeton, New Jersey: Princeton University Press.

McNamara, J. M. 1983. Optimal control of the diffusion coefficient of a simple diffusion process. *Math. Oper. Res.* **8**, 373–380.

McNamara, J. M. 1984. Control of a diffusion by switching between two drift-diffusion coefficient pairs. *SIAM J. Cont.* **22**, 87–94.

McNamara, J. M. 1990. The policy which maximizes long-term survival of an animal faced with the risks of starvation and predation. *Adv. Appl. Prob.* **22**, 295–308.

McNamara, J. M. and A. I. Houston. 1982. Short-term behaviour and life-time fitness. In *Functional Ontogeny*, D. J. McFarland (Ed.), pp. 60–87. London: Pitman.

McNamara, J. M. and A. I. Houston. 1986. The common currency for behavioural decisions. *Am. Nat.* **127**, 358–378.

McNamara, J. M. and A. I. Houston. 1987. A general framework for understanding the effects of variability and interruptions on foraging behaviour. *Acta Biotheor.* **36**, 3–22.

McNamara, J. M. and A. I. Houston. 1990a. Starvation and predation in a patchy environment. In *Living in a Patchy Environment*, B. Shorrocks and I. R. Swingland (Eds), pp. 23–43. Oxford: Oxford University Press.

McNamara, J. M. and A. I. Houston. 1990b. The state-dependent ideal free distribution. *Evol. Ecol.* **4**, 298–311.

McNamara, J. M., S. Merad and A. I. Houston. 1991. A model of risk-sensitive foraging for a reproducing animal. *Anim. Behav.* **41**, 787–792.

Merad, S. 1991. Unpublished Ph.D. Thesis. University of Bristol.

Oster, G. F. and E. O. Wilson. 1978. *Caste and Ecology in the Social Insects*. Princeton, New Jersey: Princeton University Press.

Pulliam, H. R. and G. C. Millikan. 1982. Social organization in the non-reproductive season. In *Avian Biology*, Vol. 6, D. S. Farmer, J. R. King and K. C. Parkes (Eds), pp. 169–197. New York: Academic Press.

Real, L. A. 1980. Fitness-uncertainty and the role of diversification in evolution and behaviour. *Am. Nat.* **115**, 623–638.

Real, L. A. and T. Caraco. 1986. Risk and foraging in stochastic environments: theory and evidence. *Ann. Rev. Ecol. Syst.* **17**, 371–390.

Stephens, D. W. 1981. The logic of risk-sensitive foraging preferences. *Anim. Behav.* **29**, 628–629.

Stephens, D. W. and J. R. Krebs. 1986. *Foraging Theory*. Princeton, New Jersey: Princeton University Press.

Sutherland, W. J. and C. W. Anderson. 1987. Six ways in which a foraging predator may encounter options with difference variances. *Biol. J. Linn. Soc.* **30**, 99–114.

Whittle, P. 1983. *Optimization over Time*, Vol. II. Chichester: Wiley.

Zabludoff, S. D., J. Wecker and T. Caraco. 1988. Foraging choice in laboratory rats: constant vs. variable delay. *Behav. Proc.* **16**, 95–110.

Received 25 October 1990
Revised 17 April 1991