Optimality Models in Behavioural Biology

J.M. McNamara[1]
A.I. Houston[2]
E.J. Collins[3]
[1] *School of Mathematics, University of Bristol*
University Walk, Bristol BS8 1TW, UK
Tel: +44 (0)117 928 7986
Fax: +44 (0)117 928 7999
email: John.McNamara@bristol.ac.uk

[2] *School of Biological Sciences, University of Bristol,*
Woodland Road, Bristol BS8 1UG, UK
Tel: +44 (0)117 928 7481
Fax: +44 (0)117 928 7374
email: a.i.houston@bristol.ac.uk

[3] *School of Mathematics, University of Bristol*
University Walk, Bristol BS8 1TW, UK
Tel: +44 (0)117 928 7977
Fax: +44 (0)117 928 7999
email: e.j.collins@bristol.ac.uk

Abstract

The action of natural selection results in organisms that are good at surviving and reproducing. We show how this intuitive idea can be given a formal definition in terms of fitness and reproductive value. An optimal strategy maximizes fitness, and reproductive value provides a common currency for comparing different actions. We provide a broad review of models and methods that have been used in this area, stressing the conceptual issues and exposing the logic of evolutionary explanations.

Key words: Natural selection, fitness, reproductive value, optimization, game theory, evolutionarily stable strategy.

1 Introduction

1.1 Adaptation and optimisation

If you look around you in the natural world you will see animals performing a diverse range of activities. Birds sing, flock, migrate; males fight over access to females (Figure 1); peacocks display; bees and ants bring resources back to their colonies. There are a variety of levels at which explanations of these phenomena can be given (Tinbergen, 1963). For example, suppose that it is observed that a vole tends to avoid foraging away from cover when a hawk is overhead. At one level this occurs because the vole is frightened of the hawk. This is a mechanistic explanation of the vole's behaviour, but begs the question of why the vole has this particular mechanism. An ultimate or functional explanation is that fear of hawks is adaptive; a vole that avoids hawks will have a greater chance of survival and reproduction than a vole that ignores the presence of a hawk. Here we are concerned with such adaptive explanations of behaviour rather than mechanisms such as fear that are used to implement behaviour.

Adaptive explanations view an organism's behaviour as a product of the process of natural selection. They are couched in terms of the behaviour increasing the fitness of the organism, but are really based on the following four assumptions.

1. *Stationary environment.* The ancestors of the current population faced a similar situation to the one analysed.

2. *Inheritance.* The behavioural trait used by an individual to deal with the situation is (in part) specified by its genetics. Thus an individual tends to pass on its trait to offspring through its genes.

3. *Variation.* In the past there was genetic variation in the population so that

population members varied in their behaviour trait.

4. *Differential fitness.* The trait that an individual inherited affected the number of surviving offspring produced by the individual over its lifetime. That is, the trait affected fitness.

Under the above assumptions the action of natural selection tends to lead to the population being dominated by individuals that have high fitness.

Explanations of behaviour in terms of the action of natural selection have been given ever since Darwin. Until quite recently most adaptive explanations were purely verbal. But about 30 years ago biologists started to quantify their arguments and to use mathematics to analyse behaviour. This approach was adopted at around the same time by workers in a number of areas, particularly in foraging theory, evolutionary game theory and life history theory. Whatever the area, the approach entails starting with a specification of those behavioural strategies that can be adopted, together with an appropriate measure of fitness. The dependence of fitness on the behavioural strategy adopted is then quantified. Finally, a suitable mathematical optimisation technique such as dynamic programming is used to find the strategy that maximises fitness. Often the fitness of an organism depends on both its behavioural strategy and the strategies of other population members. In such cases, optimisation is done within a game theoretic context.

Even when assumptions 1-4 hold, natural selection does not always maximise fitness. For example, it is possible to find genetic sytems in which natural selection has the effect of decreasing mean population fitness (Moran, 1964). In response to such examples it has been argued that the long-term evolution of the genetic system needs to be considered. Stable endpoints of evolution will then be optima (Eshel 1996, Hammerstein 1996, Weissing 1996, Eshel et al. 1998). Even if attention is restricted to such endpoints, we do not expect organisms to be exactly optimal. Behavioural strategies are implemented by underlying physiological or psychological mechanisms. These will typically be simple rules of thumb that approximately maximise fitness under natural circumstances, but may produce anomalous behaviour in some circumstances. For example, birds have simple rules for deciding which eggs to incubate and often prefer the largest egg. This preference rule probably has adaptive value in the wild, but under artificial conditions it can result in birds trying to incubate objects which are too large to be eggs. What this means is that if we wish to use optimality models to predict behaviour we must observe behaviour in the environment in which it has evolved. In these circumstances simple optimisation models often make reasonably accurate predictions about behaviour as well as giving insight into the effects of natural selection.

In this paper we focus on adaptations rather than the details of genetics and mechanisms, and review work that uses optimisation to make predictions about behaviour. We use the term ‘behaviour’ not just to refer to actions such as foraging or singing, but also to include decisions about the allocation of resources to growth and other aspects of morphology. The aim of the models we present is not to make detailed quantative predictions about behaviour. Rather, it is to capture the essential underlying biological features, so that analysis of these models can be used to provide general understanding of the behavioural outcomes produced by the process of natural selection.

1.2 Preview of topics

The following illustrates the range of important activities that an organism performs. By introducing these activities we motivate the topics reviewed in this paper. We start by considering the behaviour of members of a bird population over a year. During this period a bird must perform a range of activities. Each activity typically involves one or more trade-offs, so that a change in behaviour that enhances one component of fitness has a deleterious effect on another component of fitness. Activities include:

(1) *Foraging.*

Food supplies the bird with energy that fuels all its activities. If it does not get enough food to balance its energy expenditure then it will die of starvation once its energy reserves are used up. Different activities require different amounts of energy; for example migration and feeding young are particularly expensive. When foraging a bird must decide, for example

(a) whether to spend further time exploiting a given food source or use the time to attempt to find and exploit other food sources.

(b) whether to exploit a reliable but mediocre food source as opposed to a highly variable food source which may provide a lot of food if a bird is lucky or little food if it is unlucky. In other words the bird must decide whether to take a gamble.

(c) whether to exploit a food source with high yield but high predation risk or one of low yield but less risk.

The answers to questions (b) and (c) depend on the value of food to the bird, which in turn depends on the bird's energy reserves and the use to which the food is put.

(2) *Finding a mate.*

It is generally observed that females are more choosy about their mates than males are. As we suggest in Section 4.5, there are good reasons for this difference between the sexes. Females choose mates on the basis of male quality. A female must decide between accepting the current male or continuing to search for a better partner. Choice is more complex if quality cannot be directly observed. There are then two interrelated questions. What indicator of male quality should the female use? How should the resources that a male puts into attracting a female depend on his quality? For example, singing is energetically costly. Should female choice be based on male song, and if so how often and when should a male sing given his foraging ability? A crucial question in this area is whether we expect female preference and male behaviour to co-evolve so that the level of resource allocation by the male (the signal) is an honest indicator of his quality.

(3) *Caring for young.*

When a male and a female have mated and produced young they each face the choice of whether to care for the young or to desert. The decision is based on a trade-off between current and future reproduction. If a parent cares it reduces its chances of further broods that year, but young have greater survival prospects. The best choice for an individual depends on the behaviour of other population members for two reasons. One is that the effect of an individual's care on the survival of the young depends on the partner's behaviour. The other is that the chance of remating to produce a further brood depends on the availability of mates, which depends on the care decisions of all other population members.

The above activities were motivated by bird behaviour. Birds do not continue to grow throughout their life. Many fish and trees do so. These organisms face a trade-off in that energy allocated to reproduction is not available for growth.

Annual routines.

In discussing the above activities we have considered each one in isolation. This may often be a reasonable simplification, but it is important to realise that the activities that are performed over the year are linked. If we look at the whole year, a bird may have to migrate, survive the winter, breed and moult. Performing one activity may preclude another activity at this time. It may also have consequences for the state of the bird in the future. For example, the effort devoted to breeding may influence the state of the immune system and energy reserves at the end of the breeding season. The values of these state variables will affect the bird's ability to survive the winter. Thus the optimal level of reproductive effort depends on winter conditions and the strategy used by the bird during the winter. Conversely, reproductive success in the breeding season depends on the condition at the end of winter. Thus the optimal strategy overwinter depends on behaviour in the breeding season.

From the above examples the following general features of an adequate theoretical framework emerge.

(a) *Decisions should be allowed to depend on state and time.* In the above examples we have already introduced the idea of an organism's state. In the construction of an optimisation model any biologically important aspects of an organism can be included as part of its state. Components of an organism's state vector could be physiological or morphological variables such as energy reserves, body temperature, the state of the immune system or plumage colour. Components could also be external to the organism and include number of young, quality of the mate and size of the territory. The consequences of an action will typically depend on state and hence optimal actions are state dependent. We can also expect optimal actions to depend on time of day or year.

(b) *The behaviour of other population members needs to be taken into account.* In our discussion of mate choice, we raised the question of whether the signal of the male is an honest indicator of quality. To answer this question we cannot consider males in isolation, but need to consider the interaction between male signals and female preference. In this example the fitness of an organism following a given strategy depends on the strategies of other population members. The standard framework for analysing such situations is evolutionary game theory. We discuss this framework in detail in Section 2.4. It is clear that the interactions between animals that we described in the context of caring for young also require a game-theoretic approach.

(c) *Actions cannot be considered in isolation.* Consider the problem of whether an animal should expose itself to a predator in order to obtain an item of food. The best decision depends on how much food the animal will obtain in the future. This depends in part on the animal's future behaviour. Thus it is not possible to specify the best current action until future behaviour has been specified (McNamara and Houston 1986). This dependence makes it natural to consider sequences of decisions and to find optimal strategies rather than single actions. Here by the term 'strategy' we mean a rule $\pi(x, t)$ that specifies the dependence of the action taken on state x and time t. In this context we take the organism's state at time t to be a random variable, $X(t)$, whose time course depends on environmental effects that may be stochastic and on the strategy of the organism. The optimisation problem can then be viewed as choosing a behavioural strategy to optimally control the stochastic process $\{X(t)\}$. The problem can then (often) be solved by standard techniques such as dynamic programming (see Section 2.2 for details).

Throughout our discussion so far we have talked about maximisation of fitness without being specific about the definition of fitness. In some situations the number of offspring produced over an organism's life may be an adequate measure of fitness. This measure ignores the fact that offspring may be in different states so that not all offspring can be considered to be equal. Furthermore, because offspring go on to produce offspring themselves, it may be better to produce offspring earlier rather than later in order to increase the proportion of descendants in the population in the future. One of the central issues that we will address is the general definition of fitness (Sections 5.2, 6.2). Given the definition of fitness we can determine how an animal should behave over its lifetime to maximise fitness. In particular we can determine how growth and reproduction should be scheduled over the lifetime. Life history theory (discussed in Section 5) is concerned with this problem. Central questions addressed by the theory include when to start reproduction, whether to have a single bout of reproduction or to spread reproduction over several bouts and whether growth should stop when reproduction starts.

2 Setting the mathematical framework

2.1 Introduction

In this section we describe some mathematical tools that are needed in analyses of adaptive behaviour. We begin by defining how the "reproductive value" of an individual depends on its state. In the simplest setting reproductive value is defined in terms of future offspring. The general definition is not given until Section 5, but in Section 2.3 we develop the necessary mathematical framework based on long-term growth. Finally, in Section 2.4 we outline the game-theoretical framework that is used to take into account the fact that the fitness of an individual can depend on the strategies adopted by other population members.

2.2 Reproductive value as a common currency

To analyse optimal behaviour requires the advantages and disadvantages of performing an action to be quantified. For example, in modelling whether a vole should come out into the open to obtain a food item, the value of that item and the cost of death to the animal must be quantified. Similarly, in modelling parental effort the advantage to the young of a given level of effort and disadvantages to the parent such as decreased future survival must be specified. For these and other situations we need a common currency to compare the effects of possible actions (McNamara and Houston 1986, Mangel and Clark 1986).

The common currency used depends on the measure of fitness that is appropriate to the situation. We discuss fitness measures in the section on life history theory. For the present we deal with scenarios in which the fitness assigned to a behavioural strategy is the expected number of surviving offspring produced over its lifetime by an individual following that strategy. In this case actions are compared via their effect on surviving offspring. To do this let the reproductive value, $V(x,t)$, of an organism in state x at time t be defined by

$$V(x,t) = \mathbb{E}\{\text{number of surviving offspring produced from time } t \text{ onwards } |X(t) = x\}.$$

Here $\mathbb{E}\{\cdot|X(t) = x\}$ denotes the expectation given $X(t) = x$. Consider an organism that is in state x at time t. We regard the action taken at this time as determining what happens between times t and $t+1$, and ask what action maximises the expected number of surviving offspring produced from time t onwards for given reproductive value $V(\cdot, t+1)$ at time $t+1$. Define the organism's immediate reproductive success under action u to be

$$B_{\text{off}}(x, t; u) = \mathbb{E}_u\{\text{offspring produced between } t \text{ and } t{+}1 \text{ that survive until } t{+}1 | X(t) = x\},$$

where $\mathbb{E}_u$ denotes the expectation under action u. Then the expected total future number of surviving offspring under this action is

$$H(x, t; u) = B_{\text{off}}(x, t; u) + \mathbb{E}_u\{V(X(t+1), t+1) | X(t) = x\}. \tag{1}$$

It is implicitly assumed that the organism has reproductive value 0 at time $t+1$ if it is dead at this time. Future offspring are maximised (for given $V(\cdot, t+1)$) by choosing u to maximise $H(x, t; u)$, and an organism that takes this optimal action has reproductive value

$$V(x, t) = \max_{u \in Q} H(x, t; u), \tag{2}$$

where the maximisation is over the set Q of all possible actions.

Equations (1) and (2) are the dynamic programming equations determining $V(\cdot, t)$ in terms of $V(\cdot, t+1)$.

Example: should the vole emerge from cover?

To illustrate the above framework suppose that a vole has the choice between staying under cover (action u_1) or emerging from cover to obtain a food item of energetic content e (action u_2). Let the state of the vole be its level of energy reserves. Suppose the vole has reserves x at time t. Assume that if the vole stays under cover it is still alive at time $t+1$ and its reserves have decreased to $x-1$. If the vole comes out of cover it is killed by a predator during $(t, t+1)$ with probability z. If it manages to survive its reserves at $t+1$ are $x+e-1$. Thus

$$H(x, t; u_1) = V(x-1, t+1)$$

and

$$H(x,t;u_2) = (1-z)V(x+e-1,t+1).$$

It is optimal to emerge from cover if $H(x,t;u_2) > H(x,t;u_1)$; that is if

$$V(x+e-1,t+1) - V(x-1,t+1) > zV(x+e-1,t+1). \tag{3}$$

The left hand side of inequality (3) is the value, in terms of future offspring, of the food item. The right hand side is the probability of death times the loss in future offspring if the vole is killed. Thus offspring act as a common currency and it is optimal to emerge from cover if the benefit in terms of this currency exceeds the cost.

In many biological scenarios $V(\cdot, t+1)$ is a concave increasing function of reserves. Thus the left hand side of inequality (3) decreases with increasing x while the right hand side increases. It follows that there is a critical level of reserves, $x_c(z)$ such that it is optimal to emerge from cover if and only if reserves are below $x_c(z)$. Not surprisingly $x_c(z)$ decreases as z increases.

This example can be extended or generalised in various ways. For example the food supply could be made stochastic or the number of foraging options increased. We might also be concerned with a sequence of actions rather than a single action. To illustrate this, consider the behaviour of an animal over a day. During daylight the animal forages; at dusk it stops and rests during darkness. If the animal is to survive until the following dawn it must ensure that during daylight reserves never fall to zero (when it starves) and it is not killed by a predator. It must also ensure that reserves at dusk are sufficiently high to survive the night. To model behaviour let time 0 be dawn and time T be dusk. Suppose that the animal must choose a foraging option at each of the times $t = 0, 1, \ldots, T-1$. We assume that reproductive value at dusk is given by a terminal reward function R; specifically if the animal is alive with reserves x at this time its reproductive value is $R(x)$. Here we might, as an approximation, take $R(x) = 0$ for $x \leq x^*$ and $R(x) = K$ for $x > x$*, where x^* is the energy used overnight. For this decision problem we take $V(\cdot, T) = R(\cdot)$ and then use equations (1) and (2) to work backwards over the day. This backwards induction gives, for every x and t, both $V(x, t)$ and the action $u = \pi^*(x, t)$ that achieves the maximum in equation (2). The function π^* specifies how the action chosen should depend on energy reserves and time of day and is called the optimal strategy.

For further examples of reproductive value as a common currency see Houston and McNamara (1999).

2.3 Long-term survival problems

In a number of important biological scenarios the relevant optimisation criterion is the maximisation of some suitable measure of long-term growth. Here we present an analysis of the mathematics of one such scenario: maximisation of the probability of surviving a “long” winter. Later we describe how exactly the same mathematics is used to find daily and annual routines, and is used in life history theory to define fitness.

Consider the behaviour of an animal over a winter in which it must try to avoid starving or being killed by a predator while ensuring that its condition at the end of winter enables it to breed in the spring. To model the animal’s decision problem divide the winter into a finite number of epochs $t = 0, 1, \ldots, T$, where times 0 and T are the start and end of winter respectively. The animal’s state at time t is its level of energy reserves. For simplicity reserves are modelled on a discrete grid $S = \{1, 2, \ldots, L\}$. At each of the decision epochs $t = 0, 1, \ldots, T-1$ the animal must choose an action from a finite set Q. If the animal has reserves

$x \in S$ at time t and chooses action $u \in Q$ then it is alive at time $t+1$ with reserves $y \in S$ with probability $a_{xy}(u)$. In this notation there is no state corresponding to the animal being dead. Instead the probability the animal dies of starvation or predation during $(t, t+1)$ is just $1 - \sum_{y \in S} a_{xy}(u)$. If the animal is alive with reserves x at final time T then its reproductive value is $R(x)$. Here R is a non-negative valued function on the state space S.

For this decision problem a strategy is a map $\pi : S \times \{0, 1, \ldots, T-1\} \to Q$ that specifies, for each state x and time t, the action taken $\pi(x, t)$. An optimal strategy maximises the expected reproductive value at time T and can be found by dynamic programming. Before analysing this strategy, however, we first look at the performance of time-stationary strategies in the limit as the length of winter T tends to infinity.

Consider an animal that follows the stationary strategy π where $\pi(x, t) = \pi(x)$ is a function of reserves alone. Then the animal's state changes according to a stationary Markov chain (e.g. Grimmett and Stirzaker 1992) with transition law

$$P(X(t+1) = y | X(t) = x) = a_{xy}(\pi(x)).$$

Let A_π denote the matrix with (x, y) component $a_{xy}(\pi(x))$. Let $n_x(t) = P(X(t) = x)$ and let $\boldsymbol{n}(t) = (n_1(t), n_2(t), \ldots, n_L(t))$. Then $\boldsymbol{n}(0)$ gives the probability distribution of states of the animal at the beginning of winter, and the distribution after t time periods is

$$\boldsymbol{n}(t) = \boldsymbol{n}(0) A_\pi^t.$$

The animal is alive at time t with probability

$$\text{alive}(t) = \sum_{x \in S} n_x(t).$$

Thus, conditional on it being alive at t it has reserves x at this time with probability $\rho_x(t) = n_x(t)/\text{alive}(t)$, and hence $\boldsymbol{\rho}(t) \equiv (\rho_1(t), \rho_2(t), \ldots, \rho_L(t)) = \boldsymbol{n}(t)/\text{alive}(t)$.

The matrix A_π is non-negative. Following McNamara (1990a) we restrict attention to strategies π for which it is also primitive. That is, there exists a k such that all the components of A_π^k are positive. Standard Perron-Frobenius theory (Gantmacher 1959) then gives the following results. Let λ_π be the eigenvalue of A_π of maximum modulus. Then this eigenvalue is real and positive and all other eigenvalues of A_π have strictly smaller modulus than this Perron-Frobenius eigenvalue. Let $\boldsymbol{\omega}$ be the corresponding left eigenvector; i.e. $\boldsymbol{\omega} A_\pi = \lambda_\pi \boldsymbol{\omega}$. Then we can normalise $\boldsymbol{\omega}$ so that all its components are positive and they sum to 1. We also have

$$\lambda_\pi^{-t} \boldsymbol{n}(t) \to C\boldsymbol{\omega} \text{ as } t \to \infty$$

where the constant C depends on $\boldsymbol{n}(0)$. From this result we then have

$$\boldsymbol{\rho}(t) \to \boldsymbol{\omega} \text{ as } t \to \infty \tag{4}$$

and

$$\text{alive}(t+1)/\text{alive}(t) \to \lambda_\pi \text{ as } t \to \infty. \tag{5}$$

Thus conditional on the animal being alive at t it is still alive at $t+1$ with probability λ_π.

We can apply these same formulae to a large cohort of animals each independently following strategy π. $n_x(t)$ is now interpreted as the number of cohort members in state x at time t and alive(t) is the total number of cohort members still alive at this time. Regardless of the distribution of energy reserves amongst group members at the start of winter, the distribution amongst those still alive settles down to a quasi-equilibrium distribution $\boldsymbol{\omega}$ as winter progresses. At this equilibrium, of those alive at a given time a proportion λ_π are alive one time epoch later. Thus $1-\lambda_\pi$ can be regarded as the equilibrium mortality rate under strategy π.

Of course, whether formulae (4) and (5) are useful in practice depends on whether the length of winter is sufficient for convergence. The rate of convergence is determined by the subdominant eigenvalue. In the case of a small bird computations suggest that around two weeks is usually sufficient for effective convergence.

The reproductive value of an animal following strategy π during the winter is affected by both the long-term performance λ_π of the strategy and the short term effect of the animal's current energy reserves. To analyse the combined effect of these factors let $\mathcal{L} = \{f : S \to [0,\infty)\}$ be the set of non-negative real valued functions on the state space S. Let $T_\pi : \mathcal{L} \to \mathcal{L}$ be given by

$$(T_\pi f)(x) = \sum_{y \in S} a_{xy}(\pi(x)) f(y) \qquad x \in S. \tag{6}$$

Thus the action of T_π is equivalent to that of multiplication of the column vector $\boldsymbol{f} = (f(1), f(2), \ldots, f(L))^T$ by the matrix A_π. Let $V_\pi(x,t)$ denote the reproductive value of an animal with reserves x at time t given that the animal follows strategy π for the remainder of the winter. Then $V_\pi(\cdot,T) = R(\cdot)$ and $V_\pi(\cdot, T-k) = T_\pi V_\pi(\cdot, T-(k-1))$ for $k \geq 1$. Thus

$$V_\pi(\cdot, T-k) = T_\pi^k R \qquad k \geq 0. \tag{7}$$

Let $\widehat{V}_\pi \in \mathcal{L}$ be the eigenvector of T_π given by

$$T_\pi \widehat{V}_\pi = \lambda_\pi \widehat{V}_\pi. \tag{8}$$

As normalisation we set $\widehat{V}_\pi(L) = 1$ where L is the maximum level of energy reserves. Then by the Perron-Frobenius theorems

$$V_\pi(x, T-k) \sim \tilde{C} \lambda_\pi^k \widehat{V}_\pi(x) \quad \text{as} \quad k \to \infty \tag{9}$$

and hence

$$\frac{V_\pi(x, T-k)}{V_\pi(L, T-k)} \to \widehat{V}_\pi(x) \quad \text{as} \quad k \to \infty. \tag{10}$$

This last formula shows that, when the time to go is large, the value of having reserves x relative to the value of having the maximum reserves L is independent of time to go and of the terminal reward R.

Expression (9) can be used to directly compare the performance of different time-stationary strategies. Suppose π_1 and π_2 are stationary strategies whose one-step transition matrices have Perron-Frobenius eigenvalues λ_{π_1} and λ_{π_2} respectively. Let $\lambda_{\pi_1} > \lambda_{\pi_2}$. Then $V_{\pi_1}(x, T-k) > V_{\pi_2}(x, T-k)$ for all $x \in S$ for all k sufficiently large. Motivated by this we define a strategy π^+ to be the best stationary strategy if

$$\lambda_{\pi^+} = \max_\pi \lambda_\pi,$$

where the maximum is over all stationary strategies π. Heuristically this is the stationary strategy that does best in the long term.

We now drop the restriction that strategies are time stationary. Let V^* be the reproductive value function of an animal that follows an optimal strategy. Thus

$$V^*(x, t) = \max_\pi E_\pi\{R(X(T))|X(t) = x\}$$

where the maximum is taken over all strategies π. Let $T^* : \mathcal{L} \to \mathcal{L}$ be the dynamic programming operator given by

$$(T^* f)(x) = \max_{u \in Q} \sum_{y \in S} a_{xy}(u) f(y). \tag{11}$$

Then for $k \geq 0$

$$V^*(\cdot, T-k) = T^{*k} R. \tag{12}$$

Unlike T_π, T^* is a non-linear operator. The mathematics of operators of this type have been studied by a number of authors (e.g. Kennedy 1978, Grey 1984, Whittle 1983). Under suitable assumptions T^* has a real positive eigenvalue λ^* with eigenvector $\widehat{V}^* \in \mathcal{L}$. That is

$$T^* \widehat{V}^* = \lambda^* \widehat{V}^* \tag{13}$$

Furthermore

$$V^*(x, T-k) \sim (\lambda^*)^k \widehat{V}^*(x) K \qquad \text{as} \quad k \to \infty, \tag{14}$$

where the constant K depends on R. If we normalise $\widehat{V}^*$ so that $\widehat{V}^*(L) = 1$ then it can be seen that

$$\frac{V^*(x, T-k)}{V^*(L, T-k)} \to \widehat{V}^*(x) \qquad \text{as} \quad k \to \infty. \tag{15}$$

Thus under the optimal strategy, the value of being in state x relative to the value of being in state L tends to a limit as time to go increases, and the limit is independent of R.

Define the stationary strategy π^+ to satisfy $T_{\pi^+} \widehat{V}^* = T^* \widehat{V}^*$. That is, for each x, $\pi^+(x)$ satisfies

$$\sum_{y \in S} a_{xy}(\pi^+(x)) \widehat{V}^*(y) = \max_{u \in Q} \sum_{y \in S} a_{xy}(u) \widehat{V}^*(y). \tag{16}$$

Thus under π^+ the action taken in each state maximises the expected value one time step later where "value" is assigned by the function $\widehat{V}^*$. Then, provided the matrix A_{π^+} is primitive it can be shown that π^+ is the best stationary strategy (McNamara, 1990a). A stronger version of these results will be used in discussing optimal life histories (Section 5.3).

2.4 Evolutionary game theory

In the preview of topics (Section 1.2) a parent had to decide whether or not to desert the brood of young. We noted that the fitness consequences of its decision depended on the behaviour of its mate and the behaviour of all other members of the breeding population. Such 'frequency dependence' is ubiquitous in the natural world. For example the benefits to an individual of growth to a given size may depend on the size of competitors. If a female can control the sex of her offspring the benefits of producing sons will depend on the proportion of sons produced by other females. Here we outline the theory that has been developed to accomodate this frequency dependence.

Consider a large (essentially infinite) population. If almost all population members follow a strategy π then π is called the resident population strategy. Within such a population the fitness of a single individual that follows strategy π' will be denoted by $W(\pi', \pi)$. We refer to W as the payoff function, and refer to the single individual following a different strategy as a mutant; the idea being that the rare strategy π' has arisen through genetic mutation. It is important to allow for different behaviours within the same population. Given a set of strategies and a probability measure on this set we can define a new strategy under which each of the given strategies is followed with the appropriate probability. This new strategy is refered to as a mixed strategy. If members of a resident population are observed to follow different behavioural rules this could be because all are playing the same mixed strategy or because different members deterministically choose to follow different strategies; that is the population is polymorphic. The payoff to a given mutant is, however, the same regardless of how the different behavioural rules arise. Given two strategies π_0 and π_1 and a real number $\lambda \in [0, 1]$, denote the mixed strategy under which π_0 is followed with probability $1 - \lambda$ and π_1 is followed with probability λ by

$$\pi_\lambda = (1 - \lambda)\pi_0 + \lambda\pi_1.$$

In this notation the payoff function W is linear in its first argument. W is only linear in its second argument in certain special cases such as two-player games.

The physical environment and resident population strategy together determine the fitness of members of the population and of all possible mutants that could arise. As natural selection changes the frequency of the various strategies within the population, the resident strategy changes. This changes the fitnesses of the various possible strategies and hence changes the differential selection on them. One way to analyse this process is to regard the frequencies of the various possible strategies as specifying the state of a dynamical system. The dynamics for this system are typically extremely complex. Furthermore, they depend on the details of the underlying genetics which are often unknown in practice. Thus this is not a feasible approach for most scenarios. Instead the standard biological approach is to assume that the population is at a stable endpoint of the dynamics and attempt to give a purely phenotypic classification of this endpoint, so ignoring the messy details of the genetics. Maynard Smith and Price (1973) were the first to formalise this idea with the concept of an evolutionarily stable strategy (ESS) (reviewed in Maynard Smith 1982). The idea was to look for resident strategies which were stable to invasion by small numbers of identical mutants. A basic requirement for a resident population strategy π^* to be stable is that no single mutant in this population has higher fitness than resident population members; i.e.

$$W(\pi^*, \pi^*) \geq W(\pi, \pi^*) \text{ for all } \pi.$$

This is the Nash equilibrium condition of economics. This condition allows for a mutant to have payoff equal to that of a resident. Such a mutant could potentially increase in numbers due to random drift. To ensure that this does not occur the equilibrium condition is strengthened as follows

Definition. A strategy π^* is an ESS if and only if for every $\pi \neq \pi^*$ one of the following two conditions hold.

(i) $W(\pi^*, \pi^*) > W(\pi, \pi^*)$, or

(ii) $W(\pi^*, \pi^*) = W(\pi, \pi^*)$ and there exists $\delta > 0$ such that for all $\eta \in (0, \delta)$ $W(\pi^*, \pi_\eta) > W(\pi, \pi_\eta)$, where π_η is the mixed strategy $(1 - \eta)\pi^* + \eta\pi$.

Thus, for every mutant strategy either (i) when present singly a mutant does worse than resident population members, or (ii) when mutant numbers increase to a small positive proportion η of the population, mutants do worse than the original residents.

This stability criterion is concerned with whether a population already following π^* is invadable. In contrast the criterion of continuous stability is concerned with whether a population that is perturbed away from π^* will evolve back towards π^* (Eshel 1983, see also Taylor 1989, Christiansen 1991). In the case where the set of strategies form an interval on the real line the definition is as follows.

Definition. A strategy π^* is continuously stable if there exists $\delta > 0$ such that $0 < |\pi - \pi^*| < \delta$ implies there exists $\epsilon > 0$ such that $0 < |\pi' - \pi| < \epsilon$ implies

$$W(\pi', \pi) > W(\pi, \pi) \ \text{ if } \ |\pi' - \pi^*| < |\pi - \pi^*|$$

and

$$W(\pi', \pi) < W(\pi, \pi) \ \text{ if } \ |\pi' - \pi^*| > |\pi - \pi^*|.$$

In one dimension a continuously stable equilibrium is an attractor for the evolutionary dynamics under suitably well-behaved genetics, but may not be an attractor for other specifications of the genetics (Eshel et al. 1997, Weissing 1991). In more than one dimension the situation is even more complex (see Matessi and Di Pasquale 1996).

It is easy to find Nash equilibria that are stable under neither, one or both of these criteria. The evolutionary significance of the combination of criteria that hold is discussed by Geritz et al. (1998). In particular they argue that continuous stability without evolutionary stability can lead to evolutionary bifurcations.

3 Foraging

3.1 Introduction

Foraging is a general term that includes where animals search for food and which sorts of food they eat. The idea that foraging behaviour could be predicted on the basis of maximising fitness was put forward in the 1960s. A large number of models have been based on this idea and constitute what is known as optimal foraging theory (see Stephens and Krebs 1986 for a review).

Exactly how an animal should forage depends on what it is going to do with the energy that it obtains. Sometimes an animal may be building up energy to reproduce, and its reproductive success will increase with the energy that it obtains. In this case it might be reasonable to expect the animal to maximise the amount of energy obtained from a period of foraging. In other circumstances, an animal might benefit from having as much time as possible to devote to activities other than foraging. In this case we might expect the animal to minimize the time to obtain a given amount of energy. Maximising the rate of energetic gain usually achieves both of these goals when behaviour is considered over a long time interval so that transitory effects at the beginning of the interval can be ignored. Thus, many models asume that a forager maximises its fitness by maximising its net rate of energetic gain.

However, there are important classes of models where criteria other than simple rate maximisation are appropriate. Examples include models of foraging under predation risk and models with stochasticity in the food supply. The latter may exhibit risk-sensitive behaviour if there are bounds on the energy reserves or if the terminal reward is a non-linear function of the energy reserves.

We start by reviewing work in the area of rate maximisation and then go on to look at risk-sensitive foraging before finally addressing the energy-predation trade-off.

3.2 Rate maximisation: renewal reward cycles

In this section we describe rate maximisation strategies when food is obtained in bouts and the system "renews" after each bout. Such a renewal cycle can be illustrated in the context of an animal exploiting patches that contain food. We can take a cycle to start when the animal leaves a patch to search for a new one. Once a patch has been found, the animal gains energy at a rate that decreases as the food becomes depleted. Eventually the animal leaves the patch and a new cycle starts. Other examples of renewal cycles are given below. Let G denote the net energy gain on a cycle and let T denote the time taken to complete a cycle. Typically G and T are random variables whose distribution depends on the behavioural strategy adopted by the foraging animal. The mean net rate of energetic gain achieved by the animal is

$$\gamma = \frac{\mathbb{E}(G)}{\mathbb{E}(T)}, \tag{17}$$

(e.g. Johns and Miller 1963). Let

$$\gamma^* = \max \gamma \tag{18}$$

where the maximum is over all possible foraging strategies. Then an optimal strategy achieves mean net rate γ^*.

By equations (17) and (18) $E(G) - \gamma^* E(T) \leq 0$ for all strategies, with equality if and only if the strategy is optimal. Thus a strategy is optimal if and only if it maximises $\mathbb{E}(G - \gamma^* T)$. Furthermore, it can be shown that this is equivalent to maximising

$$H = \mathbb{E}(\text{Gain from remainder of cycle } - \gamma^* \times \text{ Time left on cycle}) \tag{19}$$

at every stage of the cycle (McNamara 1982). Thus γ^* acts as a rate of exchange, converting time spent in the current cycle into energy lost in future cycles. We can thus think of γ^* as the opportunity cost per unit time. This brings out the trade-off between time and energy that is at the heart of rate maximization.

We now give a brief account of two "classic" renewal cycle paradigms and a third, based on diving, that exhibits additional complexity (for simplicity we ignore energy expenditure, i.e. we work with gross rate of energy gain in all examples).

(i) *Patch use*

Consider an animal that finds food in discrete patches. If the animal spends a time t foraging in a patch of type i, then its total expected energetic gain is $G_i(t)$, where $G_i' > 0$ and $G_i'' < 0$. The animal can leave a patch at any time (this is the renewal time) and spend a mean time τ travelling to a new patch. When should the animal leave each patch type if the mean rate of energetic gain is to be maximised? The approach based on maximisation of expression (19) shows that the optimal time t_i^* on a patch of type i maximises

$$G_i(t) - \gamma^* t,$$

and hence satisfies

$$G_i'(t_i^*) = \gamma^*.$$

Thus it is optimal to leave a patch when the marginal rate at which energy is gained falls to the overall mean rate for the environment. This result is the marginal value theorem (Charnov 1976).

The marginal value theorem has been applied in contexts other than gaining energy. For example, it can be used to predict how long a male dung fly should spend copulating with a female if he is to maximize the rate at which he fertilises eggs – see Parker (1978) for details. The migratory behaviour of birds provides another example. Many species alternate between bouts of flying and bouts of feeding to replace the energy lost during flight. It has been suggested that such birds choose their migration speed to maximise their mean rate of travel (Alerstam 1991; see also Hedenström and Alerstam 1995). For this scenario we can take a renewal cycle to comprise a flight together with the subsequent refueling phase, and take the reward on a cycle to be the distance covered during the flight phase. The energy expenditure during flight is a function $P(v)$ of the flight speed v. It is easy to see that the optimal flight speed maximises $\frac{\gamma v}{\gamma + P(v)}$, where γ is the rate of refueling. This optimal speed exceeds the speed that maximises the distance flown per unit of energy (the maximum range speed) but is close to this maximum range speed when the time to refuel is long compared to flight times.

(ii) *Prey choice*

Consider an animal that forages in an environment where there are several types of prey. Prey type i yields an amount of energy e_i when consumed and takes a time h_i to handle and consume. Types are encountered as independent Poisson processes with λ_i being the encounter rate with type i items. The animal cannot encounter prey while handling. Which types should be accepted (i.e. eaten when encountered) and which should be rejected? For this scenario we can take the renewal time as the time at which search recommences after inspecting and possibly eating an item. Assume that a prey item of type i has just been encountered. The optimal decision maximises the value of expression (19). We have

$$H_{(accept)} = e_i - \gamma^* h_i \quad \text{and} \quad H_{(reject)} = 0.$$

It follows that a type i item should be accepted if and only if

$$e_i/h_i > \gamma^*.$$

In both of these examples, we have specified the optimal behaviour in terms of γ^*. But γ^* is the rate given that behaviour is optimal. Thus there is a circularity in the specification, but this is not a problem – the condition uniquely determines both the optimal behaviour and γ^*. The circularity does, however, raise the questions of how to compute γ^* and how an animal might achieve this rate. To investigate these questions McNamara (1985) defines $f : (0, \infty) \to [0, \infty)$ as follows. For each $\gamma > 0$ consider the strategy that maximises $E\{G - \gamma T\}$ and let $f(\gamma)$ equal the mean rate $E\{G\}/E\{T\}$ under this strategy. The properties of f are illustrated in Figure 2. The figure also illustrates one method of computing γ^* by constructing a sequence $\gamma_1, \gamma_2, \gamma_3, \ldots$ as follows. First consider any strategy and let γ_1 be the mean rate under this strategy (it is assumed that the strategy is chosen so that $\gamma_1 > 0$). Then set $\gamma_n = f(\gamma_{n-1})$ for $n \geq 2$. As the figure illustrates, and McNamara (1985) proves, $\gamma_n \to \gamma^*$ as $n \to \infty$.

Figure 2 about here

An animal that encounters a new environment composed of patchily distributed food will at first not even know the distribution of patch qualities in this environment. One simple rule that allows it to asymptotically learn γ^* is as follows. The animal starts with some prior estimate $\gamma_0 = \dfrac{G_{(0)}}{T_{(0)}}$ for γ^*. On the n^{th} patch encountered it behaves as if γ_{n-1} were the true value of γ^*. On leaving this patch it updates its estimate of γ^* to

$$\gamma_n = \frac{G_{(0)} + G_{(1)} + \cdots + G_{(n)}}{T_{(0)} + T_{(1)} + \cdots + T_{(n)}},$$

where $G_{(i)}$ is the gain on the ith patch and $T_{(i)}$ is the duration of the ith cycle. McNamara (1985) shows that $\gamma_n \to \gamma^*$ with probability 1. The difficulty in the proof is in showing that the sequence $\{\gamma_n\}$ converges. The key to this lies in considering the sequence of random variables $\{W_n\}_{n=1}^{\infty}$ where $W_n = G_{(n)} - f(\gamma_{n-1})T_{(n)}$. This sequence has the properties that $E(W_n) = 0$ for all n and there exists a constant K such that $E(W_n^2) \leq K$ for all n. Thus by standard martingale results

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} W_k = 0 \qquad \text{a.s.} \tag{20}$$

Having obtained this convergence result the strong law of large numbers is then used to show that $\limsup_{n\to\infty} \gamma_n \leq \gamma^*$ a.s., and hence that $\liminf_{n\to\infty} \gamma_n \geq \gamma^*$ a.s. as well.

(iii) *Diving for food*

In the above examples of patch use and prey choice the opportunity cost of spending an additional unit of time foraging is just γ^*. As we now explain, for diving animals the opportunity cost contains an additional term.

Birds such as penguins, puffins and cormorants and mammals such as otters, seals and whales hunt for their food while submerged in water. The need to return to the surface to breathe places an upper limit on the time that can be spent underwater, and the recovery time on the surface is liable to be an accelerating function of the oxygen debt on surfacing. We focus on the following renewal cycle. The diver starts a dive with no oxygen debt. It dives to a foraging area, spends a time t there and returns to the surface. The total time spent travelling to and from the foraging area is τ. At the surface, the diver spends a time $S(d)$ recovering, where $d(t, \tau)$ is the oxygen debt incurred while underwater. If the diver has been in the foraging area for a time t, the extra time on the surface, if the diver spends a further time unit foraging, is $S(d(t+1, \tau)) - S(d(t, \tau))$. Thus the opportunity cost of searching for an extra unit of time when time t has been spent foraging is

$$\gamma^*(1 + S(d(t+1, \tau)) - S(d(t, \tau)).$$

If S is an accelerating function and d is linear in t, then this cost increases as t increases. Consequently, it may be optimal to return to the surface before the diver runs out of oxygen. For this diving problem the optimal strategy can again be computed by constructing a sequence $\gamma_1, \gamma_2, \ldots$ where $\gamma_n = f(\gamma_{n-1})$. But in order to evaluate $f(\gamma)$, dynamic programming is used to find the strategy that maximises expected total gains minus expected total costs in the cycle, where the cost incurred in staying between times t and $t+1$ is

$$c_\gamma(t) = \gamma(1 + S(d(t+1, T)) - S(d(t, T))).$$

To explore the implications of this modified opportunity cost, we need to specify the foraging process. Two cases have been examined:

(i) The diver can take at most one prey item to the surface (Houston and McNamara 1985a). Here interest focuses on how the diver's prey choice criterion changes with time spent on the bottom, and the maximum time the diver should be prepared to search before surfacing.

(ii) The diver's energetic gain is proportional to the time spent in the foraging area. This case has also been used to investigate how the time at the surface depends on time underwater if behaviour is optimal. Assume that the oxygen debt is given by $d(t, \tau) = m_1 t + m_2 \tau$. Now consider the effect of increasing τ when $m_1 \neq m_2$. If the animal adopted a fixed value of t, then the resulting time at the surface would reflect the time cost associated with increasing τ. In contrast, if the animal's foraging time $t^*(\tau)$ is optimal given τ, then the resulting time at the surface does not reflect the underlying costs; indeed a plot of $S(d(t^*(\tau), \tau))$ against time underwater $t^*(\tau) + \tau$ may suggest that there is no time cost (Houston and Carbone 1992). An example is given in Figure 3. This illustrates an important general biological point: the behaviour of an organism determines certain costs or consequences, but within a population individuals differ in some confounding factor which affects optimal behaviour; thus the correlation between behaviour and its consequences observed in the population does not reflect the true effect of behaviour on consequences. In other words, natural selection hides the underlying functions.

Figure 3 about here

3.3 Risk-sensitive foraging

If animals are sensitive not only to the mean net rate of energetic gain, but also to the variance, then their foraging is said to be risk-sensitive. Risk-sensitive foraging has been reported in a variety of species (see Kacelnik and Bateson 1996 for a review).

The evolutionary approach to risk-sensitive foraging is straightforward. Assume that two options result in the same mean energetic gain but differ in that one is variable whereas the other yields the mean with certainty. Let the random variable X denote the energetic gain from the variable option, with mean $\mathbb{E}(X)$, so the fixed gain from the other option is also $\mathbb{E}(X)$. If reproductive value V is a non-linear function of the energy gained, then Jensen's inequality (e.g. Feller 1971) can be used to predict whether animals should prefer a gamble to a sure thing. When V is convex, the inequality says

$$\mathbb{E}[V(X)] \geq V(\mathbb{E}(X)),$$

so the expected reproductive value is higher if the animal gambles on the variable option. This preference for variability is known as being risk-prone. When V is concave, the inequality becomes $\mathbb{E}[V(X)] \leq V(\mathbb{E}(X))$, so it is better not to gamble, i.e. it is better to be risk-averse.

From a biological perspective the interesting question is what makes V non-linear. Some general reasons for non-linearity emerge from a consideration of the conditions that underlie the justification of rate maximisation. We pointed out above that rate maximisation will not hold if the terminal reward is nonlinear. Houston and McNamara (1999) show how an upper boundary on possible energy reserves and a lower boundary on reserves at which starvation occurs also makes V a non-linear function of reserves.

The first models of risk-sensitive foraging (Stephens 1981, Houston and McNamara 1982) were based on the foraging behaviour of a bird that stops foraging at dusk. Let the random variable $X(t)$ denote the energy reserves of the bird at times $t = 0, 1, \ldots, T$ during the day (where dawn $= 0$ and dusk $= T$) and let $V(x, t)$ denote the reproductive value of a bird with energy reserves x at time t during the day. The models assume that the bird will die during the night if it has less than a critical level of energy reserves x^* at dusk. As in the discussion of reproductive value as a common currency (Section 2.2), a simple terminal reward that captures this idea is to take

$$V(x,T) = R(x) = \begin{cases} K & x \; > \; x^* \\ 0 & x \; \leq \; x^*, \end{cases}$$

so maximising expected reproductive value at dusk is equivalent to maximising $P(X(T) > x^*)$.

Stephens (1981) considers the optimal decision for an animal with reserves x_0 at time 0. At this time, the animal has a single choice between options that differ in terms of the mean and variance of the energy gained per unit time. Once an option is chosen, it has to be used from 0 until final time T i.e. no subsequent changes in behaviour are allowed. This means that the model is static rather than dynamic (or sequential). Using the central limit theorem, Stephens shows that

$$P(X(T) > x^*) = \Phi\left[\frac{x_0 + \gamma_i T - x^*}{\sigma_i \sqrt{T}}\right]$$

where Φ is the normal distribution function and option i has mean gain γ_i and variance σ_i^2. It follows that the optimal decision is to choose the option that maximises $(x_0 + \gamma_i T - x^*)/\sigma_i$. When there are two options with the same mean, γ, then it is optimal to choose the one with the smaller variance if and only if

$$x_0 + \gamma T > x^*.$$

In other words, the animal should be risk-prone if it does not expect to get enough energy during the day to survive the following night and should be risk-averse if it does expect to get enough energy to survive the night.

However, it is not very realistic to constrain the animal to make just one decision. When the animal makes repeated decisions its level of reserves can be modelled as a diffusion process with the decisions controlling the mean and variance of this process (McNamara 1983, 1984). The general case in which the means are unequal is analyzed by McNamara (1984). Option i has mean γ_i and variance σ_i^2, with $\sigma_1 < \sigma_2$. McNamara shows that it is optimal for an animal with reserves x at time left until dusk $s = T - t$ to choose Option 1 if and only if

$$x + \left[\frac{\gamma_1\sigma_2 - \gamma_2\sigma_1}{\sigma_2 - \sigma_1}\right] s > x^*.$$

(This result makes an accurate estimate of the switching line in models in which reserves do not follow a diffusion process but food items are found as Poisson processes (Houston and McNamara 1985b).) When $\gamma_1 = \gamma_2$ this condition means that the less variable option should be chosen if and only if $x + \gamma s > x^*$. This is exactly the same rule as Stephens (1981) obtains for a problem with fixed time till dusk s. There is, however, a fundamental difference between the rules. In Stephens' model, the animal has to persist with the option originally chosen for the remainder of the time until T, whereas in the dynamic model the animal repeatedly chooses between the options on the basis of the current value of its reserves. This difference in behaviour is reflected in a difference in survival probability. For example, if we consider an animal with reserves at dawn of $x_0 = x^* - \gamma T$, then its probability of having reserves greater than x^* at T, and hence surviving the night, is 0.5 under Stephens' formulation. In contrast, survival probability under the optimal dynamic strategy is $\sigma_2/(\sigma_1 + \sigma_2)$ (see McNamara 1983 for details). We have concentrated on survival in environments with a day-night cycle. Different assumptions about ecology result in different predictions about risk-sensitivity (see McNamara and Houston 1992 for a review).

The above predictions of risk-prone behaviour are not supported by the data (Kacelnik and Bateson 1996). A problem that arises in interpreting experiments on risk-sensitive foraging is that the animal may not have the same view of the experiment as the experimenter. What we mean by this is that an experiment typically offers an animal the choice between two options that do not change over time. During the experiment the animal will not be exposed to bad weather and will not be attacked by predators. Optimal behaviour is calculated on the basis of these facts. But the animal cannot be expected to know that the options will not change over time. It may not even be reasonable to assume that the animal knows the distribution of outcomes associated with each option. The animal's behaviour has evolved in an environment that is much richer than the environment of a typical laboratory experiment. It is a challenge to the theory to explain the behaviour in simple laboratory settings as a byproduct of behaviour that is adaptive in this wider setting. This and other modelling issues are the focus of current work (McNamara 1996, Kacelnik and Bateson 1996).

3.4 The energy-predation trade-off

The maximisation of the net rate of energetic gain is not an appropriate currency when a foraging animal may run the risk of being killed by a predator. There is evidence that a higher rate of energetic gain may be associated with a higher risk (see Lima and Dill 1990, Lima 1998 for reviews). This can arise, for example, because a foraging animal may be able to allocate time to foraging or time to looking around for predators (i.e. being vigilant). As the time that it devotes to foraging increases its intake rate increases, but so does its predation rate.

If we wish to model the optimal foraging decision when an animal is faced with an energy-predation trade-off, we have to compare a gain to the animal as a result of obtaining energy with a loss that results from being killed (Section 2.2). Reproductive value provides a common currency for making this comparison. Let V be the animal's reproductive value and assume that by its choice of foraging option the animal can control both its net rate of energetic gain γ and the predation rate M to which it is subjected. Then it can be shown (Houston and McNamara 1989) that the animal should choose the foraging option that maximises

$$\gamma \frac{\partial V}{\partial x} - MV.$$

To understand this result, note that γ is the rate at which energy reserves x increase, and $\partial V/\partial x$ is the rate at which reproductive value increases with reserves, so $\gamma\partial V/\partial x$ is the rate of increase of reproductive value as a consequence of food intake. M is the rate of mortality as a result of predation, and V is the value of the animal's life, so MV is the rate of decrease of reproductive value as a consequence of predation. Thus the whole expression is the net rate at which reproductive value increases as a consequence of the animal's foraging decision, and the optimal decision maximises this net rate.

We can rewrite the optimality criterion as maximisation of $\gamma - \theta V$, where

$$\theta = \frac{1}{V} \frac{\partial V}{\partial x}$$

is the marginal rate of substitution of predation risk for energy. The concept of marginal rate of substitution is a standard part of economic theory and has been used by Caraco (1979) and Brown (1988) to analyse the energy-predation trade-off.

To apply this analysis we need to know θ. We can find θ from V and get V from an appropriate model of the future. Some examples are given in Figure 4.

Figure 4 about here

The above analysis deals with a single decision for given future expectations. We now describe a dynamic optimisation problem in which a bird makes a sequence of decisions, each of which involves an energy-predation trade-off.

Fat levels of small birds in winter

Winter can be a difficult time for small birds at high latitudes. The days are short and the nights are long and cold. Many species of birds avoid these conditions by migrating to somewhere warmer. The species that don't migrate adopt various strategies (both behavioural and physiological) that help them to survive. In this section we concentrate on the adaptive regulation of fat reserves. The basic idea is that a bird's level of fat can be understood in terms of the trade-off between starvation and predation (cf. the vole example of Section 2.2). A bird's fat reserves help it to survive a period when it cannot forage. The night is one such period, but it is also likely that foraging will be interrupted by bad weather. The higher the level of fat, the better is the bird's chance of avoiding starvation during an interruption. On this view, the bird should carry the highest possible level of fat. Of course a predation risk incurred in obtaining food may mean that it is not optimal to put on fat reserves in the first place. However, McNamara (1990a) shows that if risks do not depend on current reserves, then in the long term these initial costs do not matter (see also Houston et al. 1997a). Specifically McNamara (1990a) shows that under the strategy that maximises long-term survival, the equilibrium distribution of reserves conditional on being alive (Section 2.3) assigns positive probability to every level of reserves. Since birds are observed to regulate their fat levels below the maximum possible levels this suggests that predation costs increase with fat levels. Costs associated with high levels of fat are reviewed by Witter and Cuthill (1993). McNamara and Houston (1990) identified two costs that result in an increase in predation. Both are based on the fact that an increase in fat (all else being equal) will increase the bird's mass. An increase in mass is likely to increase the rate of energy expenditure during flight. This will mean that the bird has to spend more time foraging to maintain its reserves at a high as opposed to a low level. A longer time spent foraging means a longer time exposed to predators. The second cost is a direct increase in predation risk as a result of increased mass. This effect is expected because an increase in mass will reduce the bird's flight performance and hence its ability to escape from a predator (Hedenström 1992).

Given that fat has costs as well as benefits, how should a bird regulate its level of fat so as to minimise total mortality (starvation plus predation)? The first models to explore this question considered a single level of fat (Lima 1986, McNamara and Houston 1990). In contrast, Houston and McNamara (1993) found the optimal behaviour and the resultant level of fat as a function of time during the day. The mathematics describing these optimal daily foraging routines is a direct extension of that relevant to the long-term survival problem of Section 2.3. Unlike that problem the scenario here envisages a day-night cycle, so that the actions available (and possibly their consequences) depend on the time within a 24 hour cycle. However, we can embed the current problem into the previous time-homogeneous framework by defining the decision epochs $t = 0, 1, 2, \ldots$ of Section 2.3 to be at dusk on successive days. The state- and time-dependent strategy used over the 24 hour period between times t and $t{+}1$ is then regarded as a single action taken at time t, and the action that maximizes expected reproductive value in one day's time is found by dynamic programming over the 24 hour period. The best stationary strategy π^+ of Section 2.3 now specifies the optimal choice of action as a function of reserves and time of day. Having found this strategy Houston and McNamara (1993) calculated the distribution of energy reserves expected in a large group of birds that all follow this strategy. Figure 5a illustrates the optimal strategy and resultant levels of reserves for a specific case. The proportion of birds foraging when the optimal strategy is followed is shown in Figure 5b.

Figure 5 about here

Levels of starvation and predation

As well as making predictions about behaviour, models of optimal fat reserves also specify the levels of mortality as a result of starvation and predation if the optimal strategy is followed. An interesting general feature of the models is that the level of starvation tends to be much lower than the level of predation. Natural selection leads to the marginal levels being equal rather than the levels themselves. As a consequence the absolute level of a factor does not indicate its importance. This has been explored by Abrams (1993) who shows that the magnitude of the effect that a factor has on fitness does not predict how much of an influence the factor has on optimal behaviour. Related to this point, it is easy to construct models that exhibit "paradoxical" effects. For example, McNamara (1990b) illustrates that an increase in the food supply can increase levels of starvation. In this example animals divide their time between two food sources. Source A is poor but safe from predators. Source B is good but risky. A small improvement in Source A leads to animals spending a greater proportion of their time on this source. The predation rate thus decreases. But since Source A is still not as good as Source B, mean intake rate decreases and the starvation rate increases. McNamara (1990b) also shows that an increase in the predation risk can decrease the number of individuals that die of predation. See McNamara and Houston (1994) for a review.

4 Reproduction

Reproduction is central to an evolutionary account of behaviour: no matter how well it forages, an animal that fails to reproduce leaves no descendants. From our human perspective, we automatically think of reproduction as involving sex, but there are other forms of reproduction, and why sexual reproduction occurs is a deep and fundamental question (see Maynard Smith 1978, Michod and Levin 1988 for reviews). Sexual reproduction involves the fusion of gametes. In many species there are two types of gametes: large ones that contain nutrients (eggs) and small mobile ones (sperm). This condition is known as anisogamy and the circumstances in which it is stable have been considered (e.g. Parker et al. 1972, Matsuda and Abrams 1999). In species with separate sexes, females are defined as the sex that produces eggs and males are defined as the sex that produces sperm. There may be competition between sperm from different males to fertilise a female. Models of sperm competition are reviewed by Parker (1998). A fundamental feature of sexual reproduction in diploid species is that every individual has two parents, one male and one female. With characteristic insight Fisher (1930) realised that this simple fact has profound implications for our understanding of the allocation of resources to male and female reproduction. The total reproductive success of males has to equal the total reproductive success of females. This means that if one sex is rare, then it pays to invest in it, and so there is a frequency-dependence that can maintain a particular allocation of resources to the production of males and females. Although not phrased in terms of ESS's, this is the first example of a model using the principle of evolutionary stability. See Charnov (1982), Frank (1990) for reviews.

In this section we are primarily concerned with two aspects of reproduction: obtaining a mate and investing in offspring. These two aspects are closely related. A cost of caring for young is a loss of future matings. Future matings depend on the number of males and females seeking mates and the rules used by population members in deciding whether to mate with a particular individual. But numbers of males and females searching for mates depends on patterns of care, and the choice rules that should be used depend on the value of mating with a particular individual, which in turn depends on the individual's care behaviour. It is clear from this that a complete analysis would be very complicated. What is usually done to simplify the problem is to consider various components in relative isolation. We follow this approach by first concentrating on decisions about care and then decisions about the choice of mate.

4.1 Limiting rates of reproduction

Maynard Smith (1977; model 3) and Grafen and Sibly (1978) independently presented models in which the dependence of the time to find a mate on population decisions about care is explicit. The models are based on a species in which breeding is continuous. There is no mortality, so the only cost of care is time wasted (cf. prey choice, Section 3.2), and fitness can be measured by the long-term average rate at which offspring are produced. The essence of the model is that the potential rate of reproduction of one sex may be higher than that of the other sex. As a result, this sex has to "wait" for members of the other sex to become available. To see this principle in action, we make the following definitions. Let R_m be the reproductive rate achieved by males and R_f be the reproductive rate achieved by females, and assume that in the breeding population there are S males for every female. Then the Fisher condition of equal reproduction by males and females means that

$$R_m S = R_f. \tag{21}$$

Let τ_m be the time required by a male to recover from one breeding attempt before it searches for a mate again. Let τ_f be the corresponding recovery time for a female. These recovery times are taken as a constraint. The times that are allowed to evolve are the times spent caring by males, t_m, and the time spent caring by females, t_f. The number of offspring that survive from a breeding attempt V depends on t_m and t_f. Let the time spent searching for a mate be w_m for males and w_f for females. Then from these definitions

$$R_m = \frac{V}{\tau_m + t_m + w_m}$$

and

$$R_f = \frac{V}{\tau_f + t_f + w_f}.$$

We now seek ESS values of t_m, t_f, w_f and w_m, subject to the constraint imposed by equation (21) together with the assumption that only one sex is waiting so that minimum $\{w_m, w_f\} = 0$.

The analysis of Grafen and Sibly (1978) brings out the importance of the sex ratio. In particular, they showed that if two parents are much better than one in terms of caring for the young, then desertion will only occur at extreme sex ratios. They also show that if the sex ratio is one, then desertion first occurs when the value of both parents caring is twice that of the parent that deserts.

Yamamura and Tsuji (1993) simplified the problem by assuming that the time that the parents devote to their offspring can only take two values, 0 (i.e. no care is given) and T (corresponding to a fixed duration of care). This simplification enabled them to give a complete characterisation of the stable patterns of care as a function of the model's parameters. An interesting finding was that for some parameter values, the stable outcome involves some individuals of at least one sex caring, while other individuals of the same sex do not care. It is a fundamental result of Selten (1980) that in a two-player game with asymmetric rôles (for example, male and female) an ESS cannot involve a random choice of action. As Webb et al. (1999) point out, a mixed strategy ESS is possible here because the game between the parents is not really a two-player game but implicitly involves the whole population. Specifically, the payoffs in the game between the parents depend on the time it would take a deserting parent to find a new mate. But the time to find a mate depends on the availability of unmated individuals, which in turn depends on the outcome of the desertion games between all parents in the population. Once this feedback is taken into account mixed strategy solutions are possible.

The above models are based on rate maximisation over an infinite time horizon. McNamara et al. (2000) consider a variant of the model in which there is a breeding season of finite length. It is then possible to obtain oscillations in the pattern of care that propagate back from the end of the breeding season (Figure 6). One of the reasons for these oscillations is the difference in mating opportunities for the two sexes, and the reason for the difference in mating opportunities is the pattern of care. Another reason for oscillations is changing future expectations but expectations change because of the oscillations themselves, resulting overall in a complex pattern of behavior and rewards.

Figure 6 about here

In addition to models that investigate care and desertion, the concept of rates of reproduction forms the basis of attempts to understand competition for mates and sex rôles (e.g. Clutton-Brock and Parker 1992). The basic idea is that the sex with the lower potential rate of reproduction limits the rate of reproduction of members of the opposite sex who then compete for access to the sex with the limiting rate. In many species females care; they have a lower potential rate of reproduction as a result, and competition among males for access to females is usually observed. In some species (e.g. phalaropes, Figure 7), males have the lower rate of reproduction and females compete for access to males; a condition known as reversed sex rôles (Clutton-Brock and Vincent 1991).

4.2 Parental effort games

An alternative approach to parental care is to focus on a single reproductive bout and embed the game between parents in a life-history model. We start by outlining a simple, but widely quoted, standard model based on Houston and Davies (1985). Under the model, each parent can choose the level of effort that it devotes to the current offspring. Let the effort of the male be x and the effort of the female be y. The the total reproductive success of the male is assumed to be given by

$$B(x+y)+f_1(x),$$

where B is the success of the current brood (an increasing function of total effort) and $f_1(x)$ is the male's future reproductive success (a decreasing function of his current effort). Similarly the total reproductive success of the female is

$$B(x+y)+f_2(y).$$

The evolutionary stable efforts x^* and y^* are best responses to each other and hence emerge as joint solutions to the equations

$$B'(x^*+y^*)+f_1'(x^*)=0 \quad \text{and} \quad B'(x^*+y^*)+f_2'(y^*)=0.$$

The model above assumes that future reproductive success depends only on the level of effort of the parents and ignores any variation in the quality in either population. As a result, the equilibrium strategy for each parent is to adopt a fixed level of effort, independently of the behaviour of its partner. It is clear, however, that in practice parents respond to each other's behaviour. Responding alters the analysis since a change in the effort of one parent alters the effort of the other. Specifically if the effort of the female is a function $r_f(x)$ of the male effort x, then the payoff to the male if he adopts fixed effort x is

$$B(x + r_f(x)) + f_1(x).$$

McNamara et al. (1999) point out that responding is likely to be widespread, particularly if individuals differ in quality and these differences cannot be immediately recognised by the partner. To model this they assume that the future reproductive success of an individual depends on both its effort and its quality. Parents respond iteratively to each other until their efforts settle down to negotiated values. McNamara et al. seek a pair of negotiation rules r_m^* and r_f^* for males and females respectively, such that each is the best rule given the rule employed by the other sex. For a particular cost structure they are able to find the evolutionarily stable negotiation rule. Following this rule a male of quality q_m should respond to female effort y by adopting effort

$$r_m^*(y, q_m) = x^* + \rho + \mu q - \lambda(y - y^*),$$

wheren ρ, μ and λ are constants. Females have an analogous rule. The parameter λ specifies how responsive one parent should be to the effort of its partner. As McNamara et al. show, it is not evolutionarily stable to respond to one's partner by adopting the effort that is the best given that of the partner, since this can be exploited by a lazy partner. McNamara et al. also show that the efforts reached by a pair as a result of using these negotiation rules are not best responses to one another.

4.3 The cost of obtaining a mate

Finding a mate may be costly both in terms of energy and an associated risk of mortality (Magnhagen 1991). For example, male birds sing to attract a mate. Song costs energy and may be incompatible with foraging. McNamara et al. (1987) investigate daily routines of singing by allowing a male to choose between singing and foraging throughout the day. During the night, the bird rests. The resulting optimal routines are driven by the possibility of starvation, even though starvation is very unlikely under the optimal strategy. Because the food supply is unpredictable birds build up reserves well before dusk, typically achieving their target dusk levels before dusk. They then sing until dusk. Because overnight temperature and hence energy used overnight are unpredictable, birds cater for a worse-case scenario, typically having excess reserves at dawn and singing until the excess is used up. Thus realistic routines with a dusk and dawn chorus can be obtained without the need to invoke any variation in the food supply or the probability of attracting a female as a function of time of day.

In general, males have to compete with each other for access to females. Realistic models will typically involve repeated decisions, and hence will be modelled as a dynamic game. Sometimes such games can be solved analytically (e.g. Iwasa and Odendaal 1984), but usually this is not possible. For example, Lucas and Howard (1995) and Lucas et al. (1996) model the calling decisions of male frogs over the season. Alonzo and Warner (2000) analyse the competitive interactions between territorial and non-territorial male fish over the breeding season. These are complex game-theoretical models involving state- and time-dependent behaviour that must be solved numerically. Even a numerical solution may be difficult because the best response is typically a discontinuous function of the resident population strategy. One way in which this difficulty can be overcome is the realistic assumption that animals make errors in their decisions, where the probability of an error decreases as the cost of the error increases (McNamara et al. 1997).

4.4 Why should individuals be selective in their choice of mate?

Being selective about a mate only makes sense if potential mates differ in the benefit that results if they are chosen. For convenience we consider the possible benefits in the context of females choosing between males. (The reasons why females are likely to be choosy are discussed in the next section.) Broadly speaking, the benefits that a female obtains from mate choice are either direct or indirect (e.g. Kirkpatrick and Ryan 1991).

(i) *Direct benefits.* These are benefits that have a direct effect on a female's survival or fecundity. They include the food that a male brings to a female or that the female obtains from a male's territory. They also include any care of the offspring that the male might provide.

(ii) *Indirect effects.* Indirect effects are based on the male's genetic contribution to the female. Two sorts of genetic effect have been considered. Fisher (1930) argued that given an initial female preference for a male trait, the preference and the trait could both become exaggerated because of a positive feedback process in which females that choose males with a large value of the trait have an advantage because their sons have the exaggerated trait and hence are preferred. This is known as Fisher's runaway process. The alternative is that females get genes from the male which result in better offspring. When this is so, females may be unable to determine male quality directly, but have to rely on some signal from the male. Interest in this idea was stimulated by Zahavi's suggestion that a reliable signal of quality must carry a cost i.e. be a handicap (Zahavi 1975), so his claim is known as the handicap principle. Zahavi did not present a formal model and the principle was not initially accepted. The consensus now is that genetic models can show the effects that Zahavi claimed: to quote the title of Pomiankowski (1987) "The handicap principle does work - sometimes". Pomiankowski concluded that although these handicaps in isolation cannot result in an initial increase in female preference for the handicap when it is rare, they can substantially increase the probability that the handicap and the female preference become fixed in the population (see Andersson 1994 for a review).

Discussions of the handicap principle and the idea that a male's trait should depend on his condition led to an important general idea: honest signalling. Grafen (1990a,b) has done much to bring this idea to prominence. Grafen (1990a) considers a model in which males vary in quality q, independently of genotype. A male strategy is an advertising function A, so a male of quality q advertises himself to females at level $a = A(q)$. A female strategy is a function $D(a, t)$ specifying the probability that a female accepts as a mate a male with level of advertisement a met at time t in the season. A female's fecundity is an increasing function of her mate's quality and an eventually decreasing function of the time in the season at which she mates. While higher quality increases survival and more advertising decreases survival, an important feature of Grafen's model is the assumption that the survival disadvantage of increasing advertising is greater for lower quality males, i.e.

$$\frac{\alpha(q_1, a_2)}{\alpha(q_1, a_1)} < \frac{\alpha(q_2, a_2)}{\alpha(q_2, a_1)} \qquad \text{for } a_1 < a_2, \ \ q_1 < q_2,$$

where $\alpha(q, a)$ is the probability that a male survives to the breeding season as a function of q and a. Grafen concluded that there were essentially only two possible ‘smooth’ equilibria – a degenerate equilibrium in which $A(q) = 0$, and a unique non-degenerate equilibrium in which $A(q)$ was increasing in q (i.e. better quality males have higher levels of advertisement) and females prefer males with high levels of advertisement to males with low levels. The males develop a costly signal because of the choice behaviour of the females, and the female’s choice behaviour gives them males of higher quality that they would otherwise obtain. Although female choice is costly, the benefit obtained makes it worthwhile.

Honest signalling is now a major area of interest. Conditions for signals to be honest are discussed by Grafen (1990b), Getty (1998a,b), Számadó (1999) and Eshel et al. (2000).

Because indirect effects are based on genes, it is natural to investigate them with genetic models. This area has been the principal battleground in disputes over the strengths and weaknesses of genotypic as opposed to phenotypic models. The evolutionary process can be regarded as a dynamic system whose change over time depends on the details of the underlying genetics. The debate is over whether the stable endpoints of this dynamic system can be characterised in simple phenotype terms, independently the genetic details. This debate is ongoing. (See, for example, Kirkpatrick 1985, Grafen 1990a pp 481-2, Grafen 1990b p. 541, Kirkpatrick 1992, Andersson 1994 p. 33, Iwasa and Pomiankowski 1999 p. 106. For a general review see Gomulkiewicz 1998).

So far we have considered the benefits that a female may obtain from mate choice. It is possible that the female gains no benefit, but that the signals used by males have evolved to exploit the sensory system of females (e.g. Ryan 1998 – but see Bradbury and Vehrencamp 2000).

4.5 Optimal mate choice

Our discussion of the benefits of choice has led us to models in which female choice can evolve. We now look in more detail at non-genetic models of optimal mate choice.

In the most basic class of models the female must choose a single mate from an infinite sequence of males whose qualities are the observed values of a sequence of independent identically distributed random variables. On inspecting a male she must decide whether to accept that male or reject him and move on to the next. There is no recall of previously rejected males. To illustrate a specific model in this class, assume that between each inspection the female is killed with probability $1-\theta$, and that if she mates with a male her reproductive value equals his quality. For this model let W be the reproductive value of the female on rejecting a male, assuming that the female behaves optimally. Then the female accepts the first male whose quality exceeds W. Thus W is the unique solution of the equation

$$W = \theta E\{\max(X, W)\},$$

where the random variable X has the male quality distribution. (See Real 1990 for a review of this problem.)

The problem becomes more interesting both biologically and mathematically if interactions between the animals are present. This means that a game-theoretic approach is required. One sort of interaction, analysed in Collins and McNamara (1993), arises if more than one female is choosing, and chosen mates are no longer available to other choosers. Thus the quality of mates available to a female declines over time due to the actions of other females. (This would arise if animals form a pair to raise young, which illustrates the interactions between patterns of care and patterns of mating.) Let β denote the overall ratio of males to females and let $L(x,t)$ denote the rate at which a searching female finds males of quality greater than x at time t (so $L(\cdot,0)$ models the initial quality distribution of available males). Collins and McNamara show that, for an arbitrary smooth male availability function $L(\cdot,\cdot)$, an optimal female choice rule is a control-limit rule of the form 'accept a male of quality x found at time t if and only if $x \geq f(t)$', where f is a non-negative, non-increasing function satisfying

$$f'(t) = -\int_{f(t)}^{\infty} L(x,t)dx \qquad (t \geq 0).$$

Moreover, if all females use a common non-negative, non-increasing control-limit choice rule, then the actual male availability rule resulting from their actions can be characterised by a set of linked differential equations. Combining these results, they show that there is a unique evolutionarily stable female choice rule (i.e. one that is an optimal choice rule for precisely the the pattern of male availability created when all females use this rule). The control-limit function c^* for this rule is a solution to the equation

$$c'(t) = \beta[c(0) - h(c(t))],$$

where $h(x) = x + \beta^{-1}\int_x^{\infty} L(v,0)dv$, with $c^*(0)$ determined by the values of β and $L(x,0)$, $x \geq 0$. From this they are able to derive a number of qualitative and quantitative properties of the expected quality of the mate chosen by a given female.

Another sort of interaction arises if both sexes make choices. It is usually assumed that mating only occurs if each member of the pair finds the other acceptable. If there is a time cost associated with searching, then the best acceptance rule for one sex depends on the acceptance rule adopted by the other sex, so even if mates are not removed from the population, game theory must be used to predict behaviour. McNamara and Collins (1990) establish that the general form of the ESS in this case involves assortative mating by class. Specifically male quality can be partitioned into discrete classes $M_1, M_2, \ldots$, and female quality can be partitioned into discrete classes $F_1, F_2, \ldots$. At evolutionary stability males in class M_i mate only with females of class F_i. One example is given in Figure 8.

Figure 8 about here

Variants of this model have been developed in economics, with the interpretation of a game between a population of employers and a population of candidates seeking jobs (e.g. Burdett and Coles 1999, Mailath et al. 2000, see also Burdett and Coles 1998). In a biological context Johnstone et al. (1996) analyze a model in which, on mating, individuals spend a fixed time caring before searching for the next mate. Each population member's mate choice strategy maximises its long-term rate of reproductive success (cf. the models of Section 4.2). Given that male and female care times are different, which sex is predicted to be the more choosy? If one sex cares and the other does not then the caring sex gets less from mate choice because all that the mate provides is good genes. On the other hand since members of the opposite sex are more common amongst single individuals, the caring sex has lower mate choice costs. The numerical results obtained by Johnstone et al. suggest that the latter effect is more important and the sex that cares more should be more choosy. When both sexes care for the young it can be optimal for each sex to be selective if quality of care is sufficiently variable. This provides an explanation for species in which biparental care is associated with bright plumage in both sexes. We note that in this model time devoted to the young is fixed. It would be interesting to look at the joint evolution of mate choice and care times.

4.6 Models involving decisions about mating and care

We conclude this section with a discussion of models that consider both mating behaviour and parental care. Houston et al. (1997b) assumed that a female can control the matings that she has with two males. The matings that the males obtain determine their paternity (i.e. the probability that they are the father of the female's offspring). Given the paternity of each male, the evolutionarily stable level of effort for each male and the female can be found. Once these efforts are known, it is possible to determine the mating behaviour of the female that maximises the total parental effort. Houston et al. showed that whether a female should mate with both males or with just one male depends on the function f that specifies how future reproductive success depends on effort. In particular, if f'' is strictly increasing, the female does best by mating with just one of the males, whereas if f'' is strictly decreasing, then the female does better if she shares matings equally between the two males.

In contrast to a situation in which a female might obtain help in raising her young from two males, Kokko (1999) investigates the case in which only one male (the 'social partner') will care. The female can cuckold this male by mating with other males. This gives the female a benefit (perhaps an indirect genetic one). A male may suspect that his female has mated with another male. Although a male may mistakenly believe that he has been cuckolded when he has not, the male's perception is accurate enough for the probability of cuckoldry to be greater when a male suspects that it has occurred. As a result, males that suspect cuckoldry put less effort into caring for the young. If a male cares for the young, he reduces his opportunities for extra-pair matings. The magnitude of the cost of care thus depends on the survival of young from these matings, which depends on the care that they receive, which in turn depends on the behaviour (i.e. cuckoldry by females, suspicion by males) of other paired animals.

As with the Houston and Davies model in Section 4.2, given that the social partner contributes care x, the reward to the female for contributing care y is

$$B(x+y)-c(y)$$

and the optimal level of care y^* depends only on the given x. Assume first that the proportion of cuckolding females is fixed at z_C and assume the current (population) levels of care for suspicious and non-suspicious males are x_S and $x_{\bar{S}}$ respectively. Then the expected reward to an individual suspicious male who provides care x is

$$W_S(x;x_S,x_{\bar{S}})=H_S(x,y^*(x))-c(x)+G_S(x;x_S,x_{\bar{S}}),$$

reflecting the reproductive success of the brood, the immediate cost of care and the reward from extra-pair matings respectively. There is a similar expression for $W_{\bar{S}}(x;x_S,x_{\bar{S}})$ for non-suspicious males. For this fixed z_C, the male equilibrium care efforts x_S^* and $x_{\bar{S}}^*$ then satisfy

$$W_S(x_S^*;x_S^*,x_{\bar{S}}^*)\geq W_S(x;x_S^*,x_{\bar{S}}^*) \qquad \text{for all } x$$

$$W_{\bar{S}}(x_{\bar{S}}^*;x_S^*,x_{\bar{S}}^*)\geq W_{\bar{S}}(x;x_S^*,x_{\bar{S}}^*) \qquad \text{for all } x$$

and can be found numerically by iterating the best response map. The overall expected reward $W_C(z_C)$ to a cuckolding female is then a weighted combination of the two terms $B(x_S^*+y^*(x_S^*))-c(y^*(x_S^*))$ and $B(x_{\bar{S}}^*+y^*(x_{\bar{S}}^*))-c(y^*(x_{\bar{S}}^*))$, where the weights reflect the probability that her social partner becomes suspicious given that she is cuckolding him. There is a similar expression for $W_{\bar{C}}(z_C)$ for non-cuckolding females.

Finally, let z_C (and hence x_S and $x_{\bar{S}}$) vary. Then z_C will increase whenever $W_C(z_C) > W_{\bar{C}}(z_C)$ and a necessary condition for an overall equilibrium cuckolding and care effort behaviour at z_C^* is then $W_C(z_C^*) > W_{\bar{C}}(z_C^*)$.

Kokko found that in the non-degenerate case there were two evolutionarily stable outcomes in this system: either animals form pairs in which there is a limited amount of extra-pair mating and both parents feed the young, or all females mate with more than one male and males do not provide care. In addition to the light that it sheds on the possibility of a stable level of extra-pair mating, Kokko's analysis is also of interest in showing that only some possible patterns of mating and care may be stable. This opens up the possibility of characterising possible social systems.

5 Life history theory

5.1 Introduction

Organisms are not infinite in their capacity to produce young. There are trade-offs so that current reproductive success can only be increased at the cost of a decrease in future success. A trade-off can be environmental; for example increased time spent feeding young may increase predation risk and hence reduce the probability of being able to breed again. As discussed by McNamara and Houston (1996), many trade-offs are physiological and are mediated through the action of current behaviour on future state. For example, a tree can increase current reproductive success by increasing the allocation of resources to seed production, but this reduces growth rate and hence future reproductive success since this success increases with size.

Life history theory (see Lessells 1991, Roff 1992, Stearns 1992 for reviews) is concerned with reproductive behaviour over the whole of an organism's lifetime. The aim of the theory is to identify strategies that maximise some suitable measure of total lifetime reproductive success. The measure of lifetime reproductive success depends on the definition of fitness that is appropriate. The expected lifetime number of surviving offspring may not be an adequate fitness measure for a number of reasons including the following.

(i) How quickly the young are produced may be important. If the expected number of offspring produced by a genotype exceeds 1 then the quicker these are produced the greater the growth rate in genotype numbers.

(ii) Surviving offspring may be in different states. This is particularly important when an individual can increase the quality of offspring produced by reducing their number. It is also important when the state of an offspring is correlated with maternal state (maternal effect).

(iii) An individual can increase the rate of increase of its genotype or specific genes by helping related individuals (kin selection).

(iv) When there are large scale fluctuations of the environment that affects all population members it is not appropriate to measure reproductive success by simply taking an average over fluctuations.

In this section we concentrate on the complications (i) and (ii), with a brief description of the effects of kin selection. Environmental fluctuations are dealt with in Section 6.

5.2 Fitness

Here we define fitness when there may be seasonal effects but the environment as a whole is not subject to fluctuations from year to year. As will be seen, the mathematics required to define fitness and to characterise optimal strategies in this case is directly analogous to that of Section 2.3. In Section 2.3 individuals are maximising the probability that they survive a long winter; here individuals are maximising the number of descendants left far into the future. The two cases differ in that the former is concerned with both time-dependent and stationary strategies, whereas here we are concerned solely with stationary strategies.

Consider a large population of organisms of a given species. For simplicity assume that the species is asexual. A census is made of population members at times $t = 0, 1, 2, \ldots$ that are one year apart. At a census time the state of each population member lies in the finite state space S. Consider a given genotype within this population, and focus on a member of this genotype that is alive at time t. Define a genotype member present at time $t + 1$ to be a one-year descendant of the focal individual if either (i) it is the focal individual (having survived for the one year period), (ii) it is an offspring of the focal individual that was born between times t and $t+1$, or (iii) it is a direct descendant of an offspring of the focal individual, where the offspring was born between times t and $t + 1$. Let a_{xy} be the expected number of one-year descendants left in state y at time $t + 1$ given that the focal individual is in state x at time t. Note that, since we are assuming that there are no environmental fluctuations so that all years are the same, a_{xy} does not depend on calendar time t. The matrix $A = (a_{xy})$ is called the projection matrix for the genotype. We assume the matrix is primitive as defined in Section 2.3. Let $N(t)$ be the number of genotype members present at time t. The analogue of $N(t)$ in Section 2.3 is alive(t). Assume that numbers remain large so that we can average over the variation in numbers of descendants left by different individuals of the genotype. Then by analogy with equation (5)

$$\frac{N(t+1)}{N(t)} \to \lambda \quad \text{as} \quad t \to \infty,$$

where λ is the Perron-Frobenius eigenvalue of A. Thus λ is the asymptotic annual growth rate in genotype numbers. We define the fitness of the genotype to be this growth rate.

The above analysis assumes that the projection matrix A does not change over time. However, individuals of one genotype compete for resources with other members of that genotype and with members of other genotypes. Thus the projection matrix of a genotype depends on the numbers of all genotype members present in the population. If numbers change or the composition of the population changes the projection matrix will change and $N(t+1)/N(t)$ need not tend to a limit. Nevertheless, if we analyse the evolutionary stability of a population this is not a problem. Consider a population whose members are all of the same genotype. Suppose that the population has reached a stable size with numbers in each state stable. Then $N(t+1)/N(t) = 1$ and this resident genotype has fitness $\lambda^* = 1$. Now consider a mutation that gives rise to a new genotype. While mutant numbers are small compared with the resident population size the projection matrix for this mutant genotype can be taken to be constant. Let λ be the Perron-Frobenius eigenvalue of this matrix. If $\lambda < 1$ then mutant numbers will not grow. If this is true for all possible mutants then the resident genotype is evolutionarily stable. Conversely suppose that $\lambda > 1$ for some mutant. If there is just a single mutant individual, then the descendants of this individual may die out due to stochastic effects. If the mutation is common, however, then eventually the line of descent from one mutant individual will not die out, numbers will increase, and the per year growth rate in mutant numbers will be greater than 1 provided mutant numbers are small compared with the size of the whole population. Thus the resident genotype is not evolutionarily stable.

The above arguments show that the fitness of a genotype should be defined in the context of a resident population and gives the rate of invasion of the genotype into the population (Metz et al. 1992, see also Mylius and Diekmann 1995). At evolutionary stability the resident genotype has maximal fitness. In this sense natural selection leads to maximisation of fitness. The reasoning can be extended to various other cases. If a population is not subject to density-dependent effects fitness is again maximised at evolutionary stability but the maximum value λ^* need not be one. For sexually reproducing populations, fitness is defined on strategies rather than genotypes, and numbers of one-year descendants of a focal individual are discounted to take into account the relatedness of descendants to the focal individual. In this case, provided certain assumptions about the genetics are made, at evolutionary stability the resident strategy maximises fitness.

5.3 The optimisation problem

We now assume that there is some background resident population in which the numbers of population members in each state are constant over time. Given the 'environment' created by the physical environment and the resident population we consider how to characterise and find the strategy that maximises fitness.

In order to put the problem of fitness maximisation into a decision-theoretic framework suppose that an individual that is in state x at time t must choose an action from a set $Q(x)$ of available actions. If action u is chosen then the individual leaves $a_{xy}(u)$ expected one-year descendants in state y at time $t + 1$. We suppose that the genotype of an individual codes for the rule determining the individual's choice of actions. In other words the genotype specifies the strategy, where by a strategy π we mean a function that, for each x in S, assigns an action $\pi(x)$ in $Q(x)$. Note that strategies are defined to be time-stationary; it would not be realistic to allow the action of an individual to depend explicitly on calendar year. It is assumed that the underlying genetics are such that every possible such function can arise in the population.

The projection matrix under strategy π is $A_\pi = (a_{xy}(\pi(x)))$. Let λ_π be the Perron-Frobenius eigenvalue of the matrix; i.e. the fitness of strategy π. Set

$$\lambda^* = \max_\pi \lambda_\pi.$$

Then a strategy π^* is optimal if $\lambda_{\pi^*} = \lambda^*$. In what follows we assume that an optimal strategy exists. This will certainly hold when the set of possible actions is finite, and will hold under suitable regularity conditions when action sets are compact. In what follows we will also assume that the projection matrix A_{π^*} for some optimal strategy π^* is primitive. If this is does not hold in any particular biological scenario it can usually be made to hold by omitting redundant states from the state space S.

Given a strategy π we can define the relative reproductive value $\widehat{V}_\pi$ under π by equation (8). Then by equation (10), $\widehat{V}_\pi(x)$ measures the expected number of descendants left far into the future by an individual in state x relative to an individual in some reference state L, given both individuals follow strategy π.

Let $\widehat{V}^*$ be given by equation (13) and let strategy π^* be given by

$$T_{\pi^*}\widehat{V}^* = T^*\widehat{V}^*. \tag{22}$$

Thus under π^* the action chosen in every state maximises the expected value of one-year descendants (cf. equation (16)). Here the "value" of a descendant in state y is taken to be $\widehat{V}^*(y)$. Then as in Section 2.3 we can deduce that π^* is an optimal strategy. Furthermore, it can be shown that strategy π is optimal if and only if $T_\pi \widehat{V}_\pi = T^* \widehat{V}_\pi$ where $\widehat{V}_\pi$ is the reproductive value under π (McNamara 1991, 1993a). This condition is especially useful because it can be used to easily determine whether a given strategy is optimal. The condition says that π is optimal if and only if all actions taken under π maximise the expected value of one-year descendants, where "value" is assigned by $\widehat{V}_\pi$.

Finally, the argument leading up to equation (15) gives a robust method of computing optimal strategies. Dynamic programming is used to determine how maximum expected numbers of descendants left far into the future depends on current state. Equations (15) and (22) then determine $\widehat{V}^*$ and π^* respectively.

5.4 Lifetime reproductive success

The analysis of optimal life histories given above is based on projection matrices. The elements of a projection matrix are defined in terms of one-year descendants. But in counting descendants of an individual, no distinction is made between the focal individual itself and its offspring. Thus the notation loses track of what happens to an individual over its lifetime. Here we reformulate the above optimality criteria in terms of the maximisation of a suitable measure of total lifetime reproductive success.

Let the annual census times be as before. Since we wish to follow an individual over its lifetime we now include an organism's age as a component of its state at a census time. For simplicity we assume that offspring born between times t and $t+1$ do not reproduce before time $t+1$. We adopt the convention that offspring that survive until time $t+1$ are classified as age 1 at this time. The age of an individual at a census time is then a positive integer.

Consider the offspring that are born to a parent between times t and $t+1$.

If an offspring fails to survive until time $t+1$ it is given value 0. If an offspring survives until $t+1$ and is in state x at this time it is given value $\widehat{V}^*(x)$, where $\widehat{V}^*$ is the relative reproductive value under the optimal strategy. We define the reproductive success at time t of the parent as the total expected value of all such offspring. Let l_k be the probability that an individual survives from age 1 to age k and let b_k denote the reproductive success of an individual given that it survives to this age. The total (discounted) lifetime reproductive success of an individual is then defined to be

$$\sum_{k=1}^{\infty} (\lambda^*)^{-k} l_k b_k \tag{23}$$

where the discount factor λ^* is the growth rate under the optimal strategy. A strategy is optimal if and only if it maximises this quantity (Taylor et al. 1974, McNamara 1993b). A proof of this result in a special case is given below. The discount factor λ^* quantifies the value of producing offspring early in the lifetime when expected lifetime number of offspring exceeds 1.

Formula (23) hides much complexity. The probability of survival until age k depends on the state at age 1 and the strategy followed up until age k. The state of an individual at age k also depends on these two factors. Thus reproductive success b_k at age k depends not only on the behaviour at this age but state at age 1 and behaviour up to age k. The formula is nevertheless conceptually useful.

5.5 Age-dependent life history theory

Historically, life history theory emphasised age as an organism's state variable. Early models focussed on the scheduling of reproduction over the lifetime of an organism. Specifically, let the action u represent reproductive effort, and suppose that age is the only state variable. Suppose that an organism of age k expends reproductive effort u. Let $b_k(u)$ be the resultant expected number of offspring that survive until next year (when they are defined to be age 1). An increase in u will increase $b_k(u)$ but will typically decrease the probability $s_k(u)$ that the focal organism survives until next year (when it will be age $k+1$). Suppose that there is a maximum possible age K in this species. Then a life-history strategy can be taken to be a vector $\pi = (u_1, \ldots, u_K)$ where u_k is the reproductive effort at age k. The projection matrix under this strategy is

$$A(\pi) = \begin{bmatrix} b_1(u_1) & s_1(u_1) & 0 & 0 & \cdots & 0 & 0 \\ b_2(u_2) & 0 & s_2(u_2) & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & & \vdots & \vdots \\ b_{K-1}(u_{K-1}) & 0 & 0 & 0 & & 0 & s_{K-1}(u_{K-1}) \\ b_K(u_K) & 0 & 0 & 0 & & 0 & 0 \end{bmatrix}$$

To derive formula (23) in this case let π^* be an optimal strategy. Let λ^* be the fitness and $\widehat{V}^*(k)$ the relative reproductive value at age k under this strategy. Then $T_{\pi^*}\widehat{V}^* = \lambda^*\widehat{V}^*$. Let $\pi = (u_1, u_2, \ldots, u_K)$ be any strategy, then equation (22) implies that

$$(T_\pi \widehat{V}^*)(k) \leq (T_{\pi^*}\widehat{V}^*)(k) = \lambda^*\widehat{V}^*(k)$$

for all k, with equality at all k if and only if π is optimal. Thus

$$\widehat{V}^*(k) \geq (\lambda^*)^{-1}[b_k(u_k)\widehat{V}^*(1) + s_k(u_k)\widehat{V}^*(k+1)] \qquad 1 \leq k \leq K-1$$

with

$$\widehat{V}^*(K) \geq (\lambda^*)^{-1} b_K(u_K)\widehat{V}^*(1).$$

By repeated substitution and cancellation of $\widehat{V}^*(1)$ this yields

$$\sum_{k=1}^{K} (\lambda^*)^{-k} \ell_k(u_1, \ldots, u_{k-1}) b_k(u_k) \leq 1$$

where

$$\ell_k(u_1, \ldots, u_{k-1}) = \prod_{i=1}^{k-1} s_i(u_i)$$

is the probability of survival from age 1 to age k. Since there is equality here if and only if π is optimal, we see that π is optimal if and only if expression (23) is maximised.

One fundamental question addressed by this simplified theory is whether an organism should put all its effort into one reproductive bout that kills it (semelparity) or should attempt to breed several times (iteroparity). A general conclusion is that high background mortality selects for semelparity. For further discussion of this topic see Schaffer (1974), Bulmer (1985) and Orzack and Tuljapurkar (1989).

Given an organism that is iteroparous, we can ask about how reproductive effort should be scheduled over the organism's life. When reproductive success $b_k(u)$ does not depend on age k, we might expect a decrease in annual survival $s_k(u) = \ell_{k+1}/\ell_k$ to select for increasing reproductive effort with age. When reproductive success decreases, results are less clear.

5.6 Growth

The only state variable other than age that has frequently been used in life history theory is size. In models based on size, the energy that an organism obtains typically depends on its size, and the organism decides whether to allocate this energy to growth (thus increasing size) or to immediate reproduction. The central problem in this context is whether the organism should initially allocate all energy to growth until a critical size is reached and then allocate all energy to reproduction. Mammals and birds adopt this strategy. The alternative is that, for at least part of its life the organism allocates energy to both growth and reproduction. Fish and trees adopt this strategy, which is known as indeterminate growth. Most of the models are deterministic and Pontryagin's Maximum Principle is often used to find the optimal strategy (see Perrin and Sibly 1993 for a review). Given the usual assumptions about the trade-off between growth and reproduction, indeterminate growth can only be obtained in a constant environment when the control vector follows a singular arc (Perrin et al. 1993). Models of intermediate growth are reviewed by Heino and Kaitala (1999). Optimal growth in a seasonal environment is analyzed by Kozłowski and Teriokhin (1999).

5.7 Maternal effects

Within a hyena population females are grouped together into clans. The females in a clan have a dominance hierarchy. High ranking females produce more offspring than low ranking females. Female offspring remain in the same clan as their mother and because of the social structure, tend to have a similar dominance rank to her. Such phenotypic inheritance of maternal state is referred to as a maternal effect. Other examples include the inheritance of social status in primates and the inheritance of territories in red squirrels.

If reproductive success depends on state, and state tends to be passed on to descendants, as in hyenas, then it is clear that fitness cannot be measured in terms of numbers of offspring; even numbers of grandchildren may be a poor measure. In taking fitness to be the eigenvalue of the projection matrix, asymptotic effects far into the future are automatically included.

To illustrate the long-term consequences of maternal effects consider whether a female of high quality or high status should produce sons or daughters. The theory of parental investment presented by Trivers and Willard (1973) is based on maternal effects: females pass on their quality to both sons and daughters but males do not pass on their quality. It is argued that high quality females should produce sons since empirical evidence indicates that a high quality son can produce more offspring than a high quality daughter. This line of reasoning is based on counting numbers of grandchildren and ignores the fact that these grandchildren may vary in quality. As Leimar (1996) shows, using the correct approach based on the eigenvalue as fitness measure can mean that high quality females should produce daughters.

5.8 Kin selection

Natural selection acts on gene frequencies. In an asexual population, it tends to produce genotypes that maximize their number of descendants left far into the future. Similarly, for sexual populations natural selection tends to produce genes that maximize their number of copies left far into the future. By its effect on an organism's behaviour, a gene can influence its own spread in two different ways. One is the direct effect on the reproductive success of the organism that carries it. The other is the effect of this organism's behaviour on the reproductive success of other organisms that carry the same gene. The topic of kin selection takes account of both these effects.

The theory of kin selection was developed by Hamilton (1964a,b). Hamilton characterised the extent to which an organism should sacrifice its own reproduction in order to enhance the reproductive success of relatives.

We make no attempt to summarize the vast literature on kin selection (see Grafen 1984, Frank 1998 for an introduction). Our aim here is to illustrate how kin selection can be understood in terms of life-history theory. As before, consider an asexual organism which has a choice of actions in each of its possible states. Suppose that the organism is in state x. In the absence of kin selection, the organism's choice of action affects the projection matrix only through its effect on a_{xy} for all y. When the organism is interacting with another member of its genotype in state x', the focal organism's choice of action influences not only a_{xy} but also $a_{x'y}$. We still expect selection to maximize λ, but because of the non-localised effect of an action, we can no longer characterise the optimal strategy by the criteria given in Section 5.3. To model the non-linear effect of an action we can suppose that the projection matrix elements are functions $a_{xy}(\alpha)$ of a parameter (or vector of parameters) α. Let $A(\alpha) = (a_{xy}(\alpha))$ have Perron-Frobenius eigenvalue $\lambda(\alpha)$. Then a necessary condition for optimality is that $\lambda'(\alpha) = 0$. Let the vectors $\boldsymbol{\omega}(\alpha)$ and $\widehat{\boldsymbol{V}}(\alpha)$ satisfy

$$\boldsymbol{\omega}(\alpha)A(\alpha) = \lambda(\alpha)\boldsymbol{\omega}(\alpha) \tag{24}$$

and

$$A(\alpha)\widehat{\boldsymbol{V}}^T(\alpha) = \lambda(\alpha)\widehat{\boldsymbol{V}}^T(\alpha) \tag{25}$$

respectively. Then differentiating equation (24) with respect to α gives

$$\boldsymbol{\omega}'(\alpha)A(\alpha) + \boldsymbol{\omega}(\alpha)A'(\alpha) = \lambda'(\alpha)\boldsymbol{\omega}(\alpha) + \lambda\boldsymbol{\omega}'(\alpha).$$

Multiplying terms in this equation on the right by $\widehat{\boldsymbol{V}}^T(\alpha)$ then gives

$$\boldsymbol{\omega}'(\alpha)[A(\alpha)\widehat{\boldsymbol{V}}^T(\alpha)] + \boldsymbol{\omega}(\alpha)A'(\alpha)\widehat{\boldsymbol{V}}^T(\alpha) = \lambda'(\alpha)\boldsymbol{\omega}(\alpha)\widehat{\boldsymbol{V}}^T(\alpha) + \lambda(\alpha)\boldsymbol{\omega}'(\alpha)\widehat{\boldsymbol{V}}^T(\alpha).$$

Setting $\lambda'(\alpha) = 0$ and using equation (25) gives $\boldsymbol{\omega}(\alpha)A'(\alpha)\boldsymbol{V}^T(\alpha) = 0$. Thus, in terms of the original life history problem, we see that a necessary condition for π to be optimal is that

$$\boldsymbol{\omega}_\pi A_{\pi'}\widehat{\boldsymbol{V}}_\pi$$

has a local maximum at $\pi' = \pi$ (Taylor 1990, see also Taylor and Frank 1996) where $\boldsymbol{\omega}_\pi$ is the stable distribution of states under π and $\widehat{\boldsymbol{V}}_\pi$ is the vector of reproductive values under π. In other words, when actions are not local in their effect, optimal actions maximise average reproductive value next year.

5.9 Annual routines

Life-history models have often been used to predict a single annual decision such as the level of effort that is allocated to reproduction. There are, however, many circumstances in which the interesting biological question is how the animal should organise a series of activities over the year. The problem with predicting an annual routine is that the performance of one activity may preclude another at the same time and, by affecting future state, affect the ability to perform other activities in future. It is straightforward, however, to extend the iterative technique used to find daily routines (Section 3.4) to this case. The only minor complication is that the behaviour of young that are produced must also be explicitly taken into account (Houston and McNamara 1999).

Bird migration exhibits a particularly rich spectrum of annual routines. Some species migrate while others do not. Within a species some migrate while others (e.g. juveniles) do not, and there may be differences between the sexes in their wintering grounds and in the timing of migration. There is enormous variation across species in the timing of breeding, moult and migration. A complete moult immediately after breeding is the norm in non-migratory species, and in some migratory species. Some long-distance migrants moult on the wintering ground after the autumn migration, others start moult before the autumn migration then interupt it and resume just before the spring migration, while a few species moult twice a year. Finally, there are systematic changes in the morphology and physiology of many birds over the annual cycle.

It is a challenge to see if the above diverse patterns can be understood in adaptive terms. As a prelimiary exploration of this field McNamara et al. (1998) present a simple model of a bird that can move between two habitats, an equatorial site at which food availability is constant throughout the year and a temperate site at which food varies with season, and is worse than the equatorial site in winter and better than the equatorial site in summer. The optimal strategy typically involves breeding in the temperate site in summer and spending the winter at the equator. Figure 9 illustrates the effect of seasonality on breeding.

Figure 9 about here

Even this simple model can predict the qualitative features of known phenomena, such as a difference in the timing of adult and juvenile migration and the fact that the spring migration is less spread out than the autumn migration.

6 Fluctuating environments

6.1 Introduction

One common feature of all the models considered so far is that they only involve *demographic stochasticity*, i.e. stochasticity due to random events that affect each individual in a population largely independently of the other population members. For example, the value of food items found by a forager during a given foraging bout may be appropriately modelled as independent of the value of food items found by other foragers. In contrast, this section will focus on models where some random events may simultaneously affect *all* the individuals in a population. Typical examples are large-scale fluctuations in the overall environmental conditions (the time of the onset of winter, the general availability of food in a given season etc.) where all population members are subject to identical environmental conditions at a given time. We refer to this form of stochasticity as *environmental stochasticity*. Models where the dominant form of stochasticity is environmental (or both environmental and demographic) rather than just demographic are distinguished by referring to the former as being set in a *fluctuating environment* and to the latter as taking place in a *constant environment.* Referring back to the assumptions in Section 1.1, in a fluctuating environment it is the probability distribution over environmental conditions which is stationary over time.

The following simple example illustrates the comparative effects of demographic and environmental stochasticity on the rate of spread of a gene and the need for an appropriate definition of fitness.

Example Consider an asexual species. Assume generations are non-overlapping so individuals born in year t reach maturity and reproduce one year later and then die. Focus on one particular genotype and let the number of individuals of this genotype be $N(t), t = 0, 1, 2, \ldots$.

First suppose there is demographic stochasticity, but no environmental stochasticity. Assume each individual of the genotype produces 0 offspring with probability 0.5 and 4 offspring with probability 0.5, independently of the other individuals of the genotype. Thus the mean number of offspring per individual is 2. If $N(0)$ is large the independence of different lines means that

$$N(t) \simeq \mathbb{E}[N(t)] = 2^t N(0)$$

so $N(t)$ roughly doubles at each generation and the genotype certainly has high fitness.

Now suppose there is environmental stochasticity but not demographic stochasticity. Each year is 'good' with probability 0.5 and 'bad' with probability 0.5, independently of previous years. Each individual produces exactly 4 offspring in a good year and 0 offspring in a bad year. Now numbers quadruple in size each year until the first bad year, when all members of the genotype die:

$$P(N(t) = 4^t N(0)) = 2^{-t} \qquad P(N(t) = 0) = 1 - 2^{-t}.$$

Averaging over environmental stochasticity we again have $\mathbb{E}[N(t)] = 2^t N(0)$ but now $N(t) \to 0$ with probability 1 as $t \to \infty$. Even though mean numbers are increasing rapidly, the fitness of this genotype is, by any reasonable measure, very low. Thus, in a fluctuating environment, the mean number of offspring is not a good indicator of genotype numbers.

As with life history theory, much of the analysis in fluctuating environments is motivated by models in which there is a trade-off, either for the individual or for the genotype, between current and future reproductive success, but now future conditions are uncertain. A simple illustration of the way uncertainty affects this trade-off for the genotype is provided by Cohen's (1966) model of seed germination (with later variants considered by Bulmer (1984) and Ellner (1985a,b)). In the basic model a seed can either germinate immediately (i.e. in the current season) or remain dormant until next year. If all seeds each year follow the strategy of germinating immediately, then the entire genotype is susceptible to the potentially devastating effect of a bad season; however, if they all follow the strategy of remaining dormant, then growth of genotype numbers is delayed indefinitely. Thus an optimal strategy involves some proportion, p^*, of seeds germinating immediately and the remainder staying dormant. Again, a precise definition and characterisation of fitness in a fluctuating environment is required before we can calculate the appropriate value of p^*.

The general framework of assumptions for models in this section is as follows:

1. We have a large asexual population.

2. At each relevant time point individuals in the focal genotype can be classified

as falling into one of L states ($1 \leq L < \infty$). As before, state can be age, size, territory quality etc., or, more generally, some combination of any or all of these. For each x, we let ρ_x denote the proportion of members of the genotype in state x and we call $\boldsymbol{\rho}$ (i.e. the vector with components ρ_x) the *structure vector* for the genotype.

3. The growth of the population is described by its size at census times $t = 0, 1, 2, \ldots$ (say one year apart).

4. Between annual census times t and $t + 1$ each individual in the population makes a sequence of one or more possibly state- and time-dependent choices.

5. The state of the environment between annual census times t and $t + 1$ can be represented by a (possibly vector-valued) random variable S_t (for simplicity we assume $S_0, S_1, S_2, \ldots$ are independent random variables each with the same distribution as a random variable S, though this assumption can be considerably relaxed (Haccou and Iwasa, 1995)).

This framework is applicable to many of the biological scenarios considered in previous sections, except that the definition of fitness and hence the optimisation problems are now different. In this section we will focus on the implication of this on the mathematical techniques required.

6.2 Single-decision model with no after-effects

We start by considering a simple basic model in which the simplifying assumptions are (i) each individual in the population is required to make only one relevant decision each year and (ii) all members of the population are in the same fixed state at each census time (e.g. all eggs).

Definition of fitness

Consider members of a particular genotype in a population and let $N(t)$ denote the numbers of this genotype at census time t. We assume throughout that $N(t)$ is large. Let $r(s)$ denote the expected number of descendants left by an individual member of this genotype which are still alive at $t+1$, when the environmental state following time t is $S_t = s$. We refer to $r(s)$ as the *profile* of this genotype. Because the population is assumed to be large, we ignore demographic stochasticity and take $N(t+1) = r(S_t)N(t)$.

The mean number of offspring per individual (averaged over S) is called the *arithmetic mean fitness* and denoted by $\mathbb{E}[r(S)]$. We have seen above that this is not the appropriate measure of fitness in a fluctuating environment. Instead, following Lewontin and Cohen (1969), we note that

$$N(t) = r(S_{t-1})r(S_{t-2})\dots r(S_0)N(0)$$

so by the law of large numbers

$$\frac{1}{t}\log[N(t)] = \frac{1}{t}\log[N(0)] + \frac{1}{t}\sum_{n=0}^{t-1}\log[r(S_n)] \to g \quad \text{as} \quad n \to \infty$$

where

$$g = \mathbb{E}[\log(r(S)].$$

In Section 5.2 we saw that the value of the Perron-Frobenius eigenvalue λ formed a criterion for invasibility, that fitness could be defined in terms of invasibility, and that natural selection could be taken as leading to maximisation of fitness. Similar comments apply here, except that now the value of g forms the appropriate criterion for invasibility (Metz *et al*, 1992) in that, at stability, g is maximised. Thus, in a fluctuating environment, the standard measure of fitness of a genotype is taken to be the corresponding value of g, or equivalently the value of the *geometric mean fitness* G, where $G = e^g = \exp\{\mathbb{E}[\log(r(S))]\}$.

Characterisation of an optimal strategy

Now assume that at each census time point t each individual in the population must choose an action from a given set of possible alternatives (e.g. clutch size) without knowing the future environmental states. Let $d(u, s)$ denote the expected number of descendants left at the census time one year hence by an individual, when that individual chooses action u and the subsequent environmental state is s. We will see later that it is necessary to allow for randomised strategies. Thus a strategy π is taken to specify a probability distribution $P_\pi(u)$ over actions u. Each (distinct) genotype specifies a (distinct) strategy, which is followed by all members of that genotype. We equate the genotype with the corresponding strategy and henceforth refer to them interchangeably.

The profile r_π of a strategy π specifies the expected number of descendants left under π as a function of the environmental state, so for a genotype using π

$$r_\pi(s) = \int d(u, s) dP_\pi(u)$$

and the fitness of the genotype is

$$g(\pi) = \mathbb{E}[\log(r_\pi(S))].$$

We call π^* an optimal strategy if it satisfies

$$g(\pi^*) = g^* = \sup_\pi g(\pi).$$

The nonlinear nature of the objective function means that one cannot in general find an explicit expression for the optimal strategy and its corresponding profile and fitness. The following equivalent implicit characterisations were developed by Haccou and Iwasa (1995) and Sasaki and Ellner (1995) and, in a more general context, by McNamara (1995).

Theorem 6.1 *Let π^* be a strategy with profile r_{π^*}. Then*

(i) π^* *is optimal if and only if* $\mathbb{E}[r_\pi(S)/r_{\pi^*}(S)] \leq 1$ *for all strategies* π*,*

(ii) π^* *is optimal if and only if* $\mathbb{E}[d(u,S)/r_{\pi^*}(S)] \leq 1$ *for all pure actions* u*.*

Note that there must be equality in (ii) above for almost all actions u chosen under π^*. Furthermore although an optimal strategy need not always exist, the profile of an optimal strategy (when it exists) is unique in the sense that all optimal strategies must have the same profile (McNamara 1995).

In some cases it may be easy to find the best strategy in the class of pure strategies, say u^*. The above theorem then provides a means of checking if this pure strategy is optimal in the class of all strategies, since u^* will be optimal if and only if

$$\mathbb{E}[d(u,S)/d(u^*,S)] \leq 1 \text{ for all pure actions } u$$

(see Haccou and Iwasa 1995).

Explicit optimal strategies have been found by Haccou and Iwasa (1995) and Sasaki and Ellner (1995) for particular cases, where the distribution of S and the form of $d(u,s)$ are essentially conjugate. A more general iterative method of finding solutions was developed in McNamara (1998). Models with partial information were studied by Cohen (1967) and Haccou and Iwasa (1995).

Game-Theoretic Interpretation

Theorem **??** has the following game-theoretic interpretation. Given any strategy π, and an environment S which is a continuous random variable with density f, we can define a *modified density* f_π for S by setting $f_\pi(s) = f(s)/(K_\pi r_\pi(s))$, where K_π is chosen so that $\int f_\pi(s)ds = 1$. Similar definitions, with appropriate modifications, hold in the discrete case.

Now let π_1 and π_2 be strategies with profiles r_1 and r_2 respectively. Set

$$W(\pi_1, \pi_2) = \mathbb{E}[r_1(S)/r_2(S)]. \tag{26}$$

Following McNamara (1995), $W(\pi_1, \pi_2)$ can be interpreted as the payoff to a single individual playing strategy π_1 in a particular constant-environment game in which the resident population plays strategy π_2 and in which S is treated as a demographic variable with the modified distribution f_{π_2}. Our previous criterion for a strategy π^* to be optimal can now be interpreted as saying that π^* is optimal if and only if it is a Nash equilibrium strategy for this constant-environment game, i.e.

$$W(\pi^*, \pi^*) \geq W(\pi, \pi^*) \quad \text{for all strategies} \quad \pi.$$

Moreover, if there is a unique optimal strategy then π^* is the optimal strategy if and only if it is an ESS for the corresponding constant-environment game.

Thus, even for fluctuating environments, we can recover the concept of individual maximisation in the following sense: given the optimal profile r_{π^*} we can define a 'modified' distribution f_{π^*} for the environmental state giving rise to a modified random variable S^* such that a strategy maximising geometric mean fitness maximises arithmetic mean fitness with respect to this modified distribution.

In general, the optimal strategy π^* may well be a mixed strategy. One interpretation of this is that the fitness of the genotype is determined by the, perhaps complementary, actions of all members of the genotype present in the population and that it may be optimal for different individuals of the same genotype to take different actions, thereby using diversity to maintain a balance between expected growth and risk in much the same way that optimal portfolio selection does in economics (for a general review of the mean variance approach to balancing growth and risk see Steinbach 2001). For example, we saw in the seed germination example that the best action for one member of the genotype depends on the actions of the others. The above game-theoretic interpretation of the strategy maximising geometric mean fitness precisely quantifies this dependence on kin. Note also that such risk spreading means some individuals of the genotype may lower their own chances of producing descendants for the good of the genotype as a whole (Cooper and Kaplan 1982, Ellner 1986).

State-dependent structured populations

The results for the previous model have a straightforward extension to the case where, at each census time point t, the members of the population are allowed to differ in terms of a state variable x, but where the distribution over these states is described by a probability density or mass function which remains fixed from census point to census point irrespective of the action of individuals or the environmental conditions in the previous years. See McNamara (1998) and Haccou and McNamara (1998) for details and applications to optimal state-dependent clutch size decisions.

Variance reduction

It can be shown (Lewontin and Cohen 1969) that, when two strategies have the same arithmetic mean fitness $\mathbb{E}[r(S)]$ (as in the example above), the one with the smaller variance $\mathrm{Var}[r(S)]$ generally has greater geometric mean fitness $g = \mathbb{E}[\log(r(S)]$. A number of authors have used this as an alternative approach to comparing strategies in a fluctuating environment (see Seger and Brockmann 1987 for a review).

6.3 Dynamic optimisation within each year

Consider now the case where each individual in the population makes a sequence of $T (> 1)$ state- and time-dependent choices between each annual census point t and $t + 1$, say at decision-times $n = 0, 1, 2, \dots, T - 1$. We assume the strategy followed affects the state of each individual during the year, and that this, together with the ensuing environmental state, affects the number of members of the genotype present at the next census point. However, we assume the strategy followed within each year and sequence of the environmental states does not affect the *distribution* of genotype members into the various possible states at each annual census point. Thus the strategy followed may affect annual and long term growth in numbers, and may also affect how the structure vector $\boldsymbol{\rho}$ evolves within each year, but it will not have any long-term consequence on the value of the structure vector since its value is assumed to be reset to some fixed value at each census point.

McNamara (2000) discusses the three generic cases possible within this framework, and the various solution techniques required: (A) demographic stochasticity alone, (B) environmental stochasticity but no demographic stochasticity, and (C) both environmental and demographic stochasticity. In case A, the optimal strategy for the genotype is one under which each individual maximises the expected number of descendants left next year, and the strategy that achieves this maximum can be found simply by working backward using dynamic programming. In case B, a simple solution is possible if we further assume: (i) all members of a genotype are in the same state at decision-time 0 in each season, (ii) genotype members follow a deterministic strategy. Under these assumptions all individuals will be in the same state as one another at each subsequent time in the season. Thus, what happens to one individual in a given year represents what happens to all genotype members within that year. Provided these assumptions hold at least for the optimal genotype, each member of this genotype behaves to maximise the expected logarithm of the number of descendant that it leaves next year and this behavioural strategy can again be found by dynamic programming. Finally, in case C, we may suppose that one of the two assumptions above fails. Thus there will be a time n in the season when members of the genotype will be in a range of states. In this case, the best action of an individual at this time implicitly depends on this distribution of states of relatives. But this distribution depends on the behavioural strategy followed by genotype members before time n. Thus it is not possible to specify optimal behaviour at time n until optimal behaviour has been specified at all previous times as well as future times, since the states of relatives can only be reconstructed from knowledge of their previous behaviour. The optimal strategy cannot now be found by simple dynamic programming (Collins and McNamara 1998).

The above discussion has several immediate implications.

1. The genotypic fitness under the best population-based strategy will generally be greater than under the best individual-based strategy — where we call a strategy π an *individual-based* strategy if the probability distribution of possible actions at a given time employed by an individual in a given state x depends only on x, and call π a *population-based* strategy if this distribution may depend on both x and the current value of $\boldsymbol{\rho}$.

2. In case C above, the computation of the best population-based strategy, even when straightforward in principle, will always be computationally complex, since the space Ω of possible $\boldsymbol{\rho}$ values is a $L-1$ dimensional simplex. Moreover, even the formal characterisation of the best individual-based strategy can be problematic

since it cannot be characterised explicitly in the standard way as the optimal solution to a dynamic programming problem.

3. Issues of biological interpretation arise if the best population-based strategy differs from the best individual-based strategy, since it will generally be unrealistic to assume each individual member of a genotype will have available to them enough information about all the fellow members of their genotype to compute $\boldsymbol{\rho}$ and follow a population based strategy.

6.3.1 Diapause: a population-based approach

The direct calculation of an optimal population-based strategy is exemplified in McNamara (1994), where a complete analytic solution is obtained to the problem of finding the optimal proportion of a population entering diapause (i.e. a dormant state which allows for overwinter survival) at any given time in an unpredictable environment. The original model was first introduced by Cohen (1970). Others have since considered this question (see Hanski 1988 for a review) and the related question of optimal allocation to growth and reproduction in annual plants (Cohen 1971,1976, King and Roughgarden 1982, and Amir and Cohen 1990).

In the model of McNamara (1994), the variable length of the time till the weather turns bad plays the rôle of the environmental state variable S and the start of each season marks successive census points. Within each season, time is scaled so that the time points $n = 0, 1, 2, \ldots$ (up to a maximum of T) mark the start of successive non overlapping generations of an organism. At each time $(n-1)$, the nth generation is composed entirely of eggs. Some of these enter diapause; the remainder die if the season ends before time n, otherwise they grow to maturity at time n, lay on average $\alpha(> 1)$ new eggs and die. All eggs entering diapause survive till the start of the next season with a probability that is independent of the time of entering diapause and the season length. The trade-off is between a strategy of growth, which risks the genotype being wiped out if the season ends before there are any eggs in diapause, and the conservative strategy of all eggs entering diapause, under which genotype numbers will dwindle over successive seasons.

Let $r(S)$ denote the number of eggs in diapause when the season ends and let $p(n)$ denote the proportion of new eggs at time n which enter diapause. The problem then is to choose a strategy $\pi = (p(0), \ldots, p(T-1))$ to maximise $\mathbb{E}[\log(r(S))]$. What makes the model particularly tractable is that the structure vector $\boldsymbol{\rho}(n)$ is essentially a scalar; its two components are $\rho_D(n)$, corresponding to the proportion of the genotype currently in diapause, and $\rho_N(n) = 1-\rho_D(n)$, corresponding to the proportion in the form of new eggs.

The measure of fitness can be reformulated so that maximising $\mathbb{E}[\log(r(S))]$ is shown to be equivalent to maximising an expected total reward over time of the form

$$\mathbb{E}\left[\log \tilde{r}(1, \rho_D(1)) + \sum_{n=1}^{S-1} \tilde{r}(\rho_D(n), \rho_D(n+1))\right]$$

where

$$\tilde{r}(x, y) = \log\{\alpha y/[x(1 + (\alpha - 1)y)]\}.$$

Let $H(n) = P(S = n \mid S > n - 1)$ denote the probability that bad weather first occurs during the growth of generation n, given that it has not previously occurred. Assume $H(1) \leq H(2) \leq \ldots \leq H(T) = 1$ and set

$$c^*(n) = \min\left\{1, \left(\frac{1}{\alpha - 1}\right)\left[\frac{H(n)}{1 - H(n)}\right]\right\}.$$

Then the optimal strategy is to control the process in such a way that $\rho_D(n)$ is as close to $c^*(n)$ as possible.

An explicit solution for the optimal strategy π^* in terms of $H(\cdot)$ is found using dynamic programming arguments. Set $\tilde{n}$ as the first time at which $c^*(n) = 1$. The optimal strategy maximising $\mathbb{E}[\log(r(S))]$ is

$$p^*(n) = \frac{\alpha P(S = n+1) - P(S = n)}{\alpha P(S \geq n+1) - P(S \geq n)}$$

for $n \leq \tilde{n}$ and $p^*(\tilde{n}) = 1$.

Note that the optimal population-based strategy specifies a graded optimal response, with various non-zero proportions entering diapause up to a given time $\tilde{n}$, when all the surviving eggs enter diapause. In contrast, the 'bang-bang' strategy optimal for the corresponding constant-environment formulation specifies that no eggs enter diapause before $\tilde{n}$ and all eggs enter diapause at $\tilde{n}$. Note also that the optimal population-based strategy has a particularly simple form: each egg in generation n should independently enter diapause with the same probability $p^*(n-1)$. This optimal strategy can be applied by each individual ignoring the actions of related individuals. Nevertheless, in a large population of individuals all carrying a control gene coding for this individual-based rule, the proportion of individuals in diapause at time n will attain the optimal value $c^*(n)$, and any one individual will be behaving as if it were responding optimally to this value of the population structure vector $\boldsymbol{\rho}$. Thus the optimal population-based strategy π^* can be achieved by an individual-based rule (see above).

6.3.2 Dynamic optimisation: an individual-based approach

McNamara et al. (1995) and Collins and McNamara (1998) use an individual-based strategy formulation to analyse a general class of problems for which the annual reward to the genotype depends only on its state (i.e. the structure vector) at the end of each season. The approach is illustrated by an application to optimal parasitoid foraging and subsequent egg laying over a season.

In their model each season is composed of decision points $n = 0, 1, 2, \ldots, T-1$ at which each individual makes general state- and time-dependent transitions according to an individual-based strategy. Let $d(x, s)$ be the expected number of descendants left by an individual whose final state at time T is x when the state of the environment is $S = s$. The distribution of the state of these descendants is taken to be independent of both x and s. Let ρ_x^π denote the probability that a randomly chosen individual will be in state x at time T when all individuals in the population use the same individually-based strategy π, and let $\boldsymbol{\rho}^\pi$ be the corresponding end of season structure vector with components ρ_x^π. The population size is assumed to be sufficiently large that the profile of a given strategy π can be taken to be $r_\pi(s) = \Sigma_x \rho_x^\pi d(x, s)$ and its fitness is then $g^\pi = \mathbb{E}[\log r_\pi(S)]$.

Using a Markov decision process formulation, Collins and McNamara (1998) show that an optimal individual-based strategy will, in general, be in the form of a mixed strategy, i.e. one in which each individual choses a pure deterministic strategy from a given set $\pi_1, \pi_2, \pi_3, \ldots$ with respective probabilities $\theta_1, \theta_2, \theta_3, \ldots$ and then uses that strategy throughout the season, so that $\rho_x^\pi = \Sigma_i \theta_i \rho_x^{\pi_i}$. They demonstrate that a direct dynammic programming approach cannot be used to find an optimal individual-based strategy. Instead, analysis is based on extending the characterisation above of π^* as a Nash equilibrium strategy in a constant-environment game.

For any two strategies π_1 and π_2 again set

$$W(\pi_1, \pi_2) = \mathbb{E}[r_{\pi_1}(S)/r_{\pi_2}(S)] = \Sigma_x \ \rho_x^{\pi_1} \mathbb{E}[d(x, S)/\Sigma_y \{d(y, S)\rho_x^{\pi_2}\}]$$

and call π_1 a *best response* to π_2 if

$$W(\pi_1, \pi_2) \geq W(\pi, \pi_2) \quad \text{for all mixed strategies } \ \pi.$$

Then a strategy π^* is an optimal strategy if and only if it is a best response to itself. Moreover, the form of $W(\pi_1, \pi_2)$ shows that best response to a given strategy π_1 *can* now be found by direct dynamic programming. This motivates an iterative procedure in which each iteration corresponds to (i) computing the end of season structure vector $\boldsymbol{\rho}^{\pi_k}$ corresponding to a given strategy π_k (by projecting forward using the known initial structure vector $\boldsymbol{\rho}_0$ and the known individual transition probabilities under π_k), (ii) using standard dynamic programming methods to find the best response $\hat{\pi}_k$, and (iii) the construction of an updated strategy π_{k+1} based on $\hat{\pi}_k$ and previous strategies. In practice a simple updating in which $\pi_{k+1} = \lambda\hat{\pi}_k + (1-\lambda)\pi_k$ for given fixed λ is often sufficient to ensure convergence (McNamara et al. 1995), but convergence can always be guaranteed by use of an appropriately rigorous algorithm (Collins and McNamara 1998).

Finally Collins and McNamara (1998) show that the optimal individual-based strategy is also an optimal strategy within the wider class of population-based strategies.

6.4 Optimal life histories for structured populations

The case of structured populations with effects persisting into the future is treated in McNamara (1997), illustrated by applications to optimal state-dependent offspring number decisions. Again the population is censused at time $t = 0, 1, 2, \ldots$. At a census time each individual can be classified by state and must choose a single (possibly randomised) action from a set of actions, with the choice being made before the environmental state S_t between t and $t+1$ is known. The individual's state x, the action chosen u and the environmental condition s determine the expected number of descendants $a_{xy}(u, s)$ left by this individual in state y at time $t+1$. Thus in this formulation the structure vector $\boldsymbol{\rho}(t+1)$ depends explicitly on $\boldsymbol{\rho}(t)$ and S_t as well as the strategy employed between times t and $t+1$, whereas in the previous formulations discussed $\boldsymbol{\rho}(t)$ was assumed to revert to some fixed constant value at the start of each year.

Let $r(\boldsymbol{\rho}, S)$ be the expected number of descendants left by a randomly selected genotype member when the structure vector is $\boldsymbol{\rho}$ and the ensuing environmental state is s, and let

$$m(\boldsymbol{\rho}) = \mathbb{E}[\log(r(\boldsymbol{\rho}, S))].$$

Then the fitness of a given population-based strategy is taken to be

$$\log \lambda = \int_{\Omega} m(\boldsymbol{\rho}) h(\boldsymbol{\rho}) d\boldsymbol{\rho} = \lim_{T\to\infty} \frac{1}{T} \mathbb{E}\left\{ \sum_{t=0}^{T} m(\boldsymbol{\rho}(t)) \right\}$$

where $h(\cdot)$ is the long-term stationary distribution of $\boldsymbol{\rho}$ under the given strategy and Ω is the standard $L-1$ dimensional simplex. This extends results of Tuljapurkar (1989) for individual-based strategies.

This leads to an interpretation of $\log \lambda$ as the 'average reward' each year and motivates the use of standard Markov decision process theory to provide the following characterisation of the fitness of a given strategy π.

Theorem 6.2 *There is a real valued function v^{π} such that*

$$v^{\pi}(\boldsymbol{\rho}) + \log \lambda^{\pi} = m^{\pi}(\boldsymbol{\rho}) + \mathbb{E}^{\pi}\{v^{\pi}(\boldsymbol{\rho}(1)) \mid \boldsymbol{\rho}(0) = \boldsymbol{\rho}\} \tag{27}$$

holds for all structure vectors $\boldsymbol{\rho}$.

Note that equation (27) gives the basis for a method of iteratively calculating the dominant Lyapunov exponent of a random matrix without recourse to calculating the stationary measure.

Equation (**??**) determines $\log \lambda^\pi$ uniquely and determines v^π up to an additive constant. Furthermore, $v^\pi(\boldsymbol{\rho})$ can be interpreted as the *reproductive value* of a state $\boldsymbol{\rho}$, in that the difference $v^\pi(\boldsymbol{\rho}_1) - v^\pi(\boldsymbol{\rho}_2)$ gives the relative advantage for a population starting with structure vector $\boldsymbol{\rho}_1$ rather than $\boldsymbol{\rho}_2$, when both populations use the same strategy π.

Based on Theorem **??**, policy iteration and value iteration methods can be used to characterise and (at least formally) compute an optimal population-based strategy π^*.

Although optimal individual-based strategies are much harder to characterise, McNamara (1997) obtains the following:

Theorem 6.3 *Let π^* be an optimal individual-based strategy. Let v^{π^*} be the reproductive value function under π^* and let h^{π^*} be the equilibrium density function for the process $\{\boldsymbol{\rho}(t) : t = 0, 1, 2, \ldots\}$ under π^*. Then the expression*

$$\int_\Omega h^{\pi^*}(\boldsymbol{\rho})[m^\pi(\boldsymbol{\rho}) + \mathbb{E}^\pi\{v^{\pi^*}(\boldsymbol{\rho}(1)) \mid \boldsymbol{\rho}(0) = \boldsymbol{\rho}\}]d\boldsymbol{\rho}$$

is maximised when $\pi = \pi^*$, i.e. π^* maximises in the mean the quantity $m^\pi(\rho) + \mathbb{E}^\pi\{v^{\pi^*}(\boldsymbol{\rho}(1)) \mid \boldsymbol{\rho}(0) = \boldsymbol{\rho}\}$.

7 Further topics

The topic of optimality models in behavioural biology is vast. It has been impossible for us to cover all of it. Important topics omitted include the following.

(i) The conflict of interest between offspring and their parents and between siblings (Mock and Parker, 1997).

(ii) The distribution of animals across habitats. The central concept here is that of an ideal feee distribution – a distribution such that no animal can do better by moving to another location (see Milinski and Parker 1991, Tregenza, 1995 for reviews).

(iii) Building up from optimal individual behaviour to population level processes (Abrams 1992, 1997, Matsuda and Abrams 1994, Sutherland 1996, Abrams and Matsuda 1997, Houston and McNamara 1997, Brown 1998).

(iv) Social behaviour, including co-operation (Axelrod and Hamilton 1981; see Dugatkin 1997 for a review) and social foraging (Giraldeau and Caraco 2000).

Acknowledgements

We thank Philip Drazin, Jon Keating and Guy Nason for comments on a previous version of the manuscript. Our work on care and desertion was supported by NERC (grant GR3/10957).

REFERENCES

P.A. ABRAMS (1992), *Adaptive foraging by predators as a cause of predator-prey cycles*, Evol. Ecol., 6, pp. 56-72.

P.A. ABRAMS (1993), *Optimal traits when there are several costs: the interaction of mortality and energy costs in determining foraging behavior*, Behav. Ecol., 4, pp. 246-253.

P.A. ABRAMS (1997), *Evolutionary responses of foraging-related traits in unstable predator-prey system*, Evol. Ecol., 11, pp. 673-686.

P.A. ABRAMS, AND H. MATSUDA (1997), *Fitness minimization and dynamic instability as a consequence of predator-prey coevolution*, Evol. Ecol., 11, pp. 1-20.

T. ALERSTAM (1991), *Bird flight and optimal migration*, Trends Ecol. Evol., 6, pp. 210-215.

S.H. ALONZO AND R.R. WARNER (2000), *Dynamic games and field experiments examining intra- and intersexual conflict: explaining counterintuitive mating behavior in a Mediterranean wrasse, Symphodus ocellatus*, Behav. Ecol., 11, pp. 56-70.

S. AMIR AND D. COHEN (1990), *Optimal reproductive efforts and the timing of reproduction of annual plants in randomly varying environments*, J. Theor. Biol., 147, pp. 17-42.

M. ANDERSSON (1994), *Sexual Selection*, Princeton University Press, Princeton.

R. AXELROD AND W.D. HAMILTON (1981), *The evolution of cooperation*, Science, 211, pp. 1390-1396.

J.W. BRADBURY AND S.L. VEHRENCAMP (2000), *Economic models of animal communication*, Anim. Behav., 59, pp. 259-268.

J.S. BROWN (1988), *Patch use as an indicator of habitat preference, predation risk, and competition*, Behav. Ecol. Sociobiol., 22, pp. 37-47.

J.S. BROWN (1998), *Game theory and habitat selection*, in Game Theory and Animal Behavior, L.A. Dugatkin and H.K. Reeve, eds, Oxford University Press, New York.

M.G. BULMER (1984), *Delayed germination of seeds: Cohen's model revisited*, Theor. Pop. Biol., 26, pp. 367-377.

M.G. BULMER (1985), *Selection for iteroparity in a variable environment*, Am. Nat., 126, pp. 63-71.

K. BURDETT AND M.G. COLES (1998), *Separation cycles*, J. Econ. Dyn. Cont., 22, pp. 1069-1090.

K. BURDETT AND M.G. COLES (1999), *Long-term partnership formation: Marriage and employment*, Econ. Jour., 109, pp. F307-F334.

T. CARACO (1979), *Time budgeting and group size: a theory*, Ecology, 60, pp. 611-617.

E.L. CHARNOV (1976), *Optimal foraging: the marginal value theorem*, Theor. Pop. Biol., 9, pp. 129-136.

E.L. CHARNOV (1982), *The Theory of Sex Allocation*, Princeton University Press, Princeton.

F.B. CHRISTIANSEN (1991), *On conditions for evolutionary stability for a continuously varying character*, Am. Nat., 138, pp. 37-50.

T.H. CLUTTON-BROCK AND G.A. PARKER (1992), *Potential reproductive rates and the operation of sexual selection*, Quart. Rev. Biol., 67, pp. 437-456.

T.H. CLUTTON-BROCK AND A.C.J. VINCENT (1991), *Sexual selection and the potential reproductive rates of males and females*, Nature, 351, pp. 58-60.

D. COHEN (1966), *Optimizing reproduction in a randomly varying environment*, J. Theor. Biol., 12, pp. 119-129.

D. COHEN (1967), *Optimizing reproduction in a randomly varying environment when a correlation may exist between the conditions at the time a choice has to be made and the subsequent outcome*, J. Theor. Biol., 16, pp. 1-14.

D. COHEN (1970), *A theoretical model for the optimal timing of diapause*, Am. Nat., 104, pp. 389-400.

D. COHEN (1971), *Maximizing final yield when growth is limited by time or by limiting resources*, J. Theor. Biol., 33, pp. 299-307.

D. COHEN (1976), *The optimal timing of reproduction*, Am. Nat., 110, pp. 801-807.

E.J. COLLINS AND J.M. McNAMARA (1993), *The job search problem with competition: an evolutionarily stable dynamic strategy*, Adv. in Appl. Probab., 25, pp. 314-333.

E.J. COLLINS AND J.M. McNAMARA (1998) *Finite-horizon dynamic optimisation when the terminal reward is a concave functional of the distribution of the final state*, Adv. in Appl. Probab., 30, pp. 122-136.

W.S. COOPER AND R.H. KAPLAN (1994), *Adaptive "coin-flipping": a decision-theoretic examination of natural selection for random individual variation*, J. Theor. Biol., 94, pp. 135-151.

L.A. DUGATKIN (1997), *Cooperation Among Animals: An Evolutionary Perspective*, Oxford University Press, Oxford.

S. ELLNER (1985a), *ESS germination strategies in randomly varying environments 1. Logistic-type models*, Theor. Pop. Biol., 28, pp. 50-79.

S. ELLNER (1985b), *ESS germination strategies in randomly varying environments 2. Reciprocal yield-law models*, Theor. Pop. Biol., 28, pp. 80-116.

S. ELLNER (1986), *Germination dimorphisms and parent-offspring conflict in seed-germination*, J. Theor. Biol., 123, pp. 173-185.

I. ESHEL (1983), *Evolutionary and continuous stability*, J. Theor. Biol., 103, pp. 99-111.

I. ESHEL (1996), *On the changing concept of evolutionary population stability as a reflection of a changing point of view in the quantitative theory of evolution*, J. Math. Biol., 34, pp. 485-510.

I. ESHEL, M.W. FELDMAN AND A. BERGMAN (1998), *Long-term evolution, short-term evolution and population genetic theory*, J. Theor. Biol., 191, pp. 391-396.

I. ESHEL, U. MOTRO AND E. SANSONE (1997), *Continuous stability and evolutionary convergence*, J. Theor. Biol., 185, pp. 333-343.

I. ESHEL, I. VOLOVIK AND E. SANSONE (2000), *On Fisher-Zahavi's handicapped sexy son*, Evol. Ecol. Res., 2, pp. 509-523.

W. FELLER (1971), *An Introduction to Probability Theory and its Applications, Vol. 2*, John Wiley, New York.

R.A. FISHER (1930), *The Genetical Theory of Natural Selection*, Clarendon Press, Oxford.

S.A. FRANK (1990), *Sex allocation theory for birds and mammals*, Ann. Rev. Ecol. Syst., 21, pp. 13-56.

S.A. FRANK (1998), *Foundations of Social Evolution*, Princeton University Press, Princeton.

F.R. GANTMACHER (1959), *Applications of the Theory of Matrices*, Interscience, New York.

S.A.H. GERITZ, E. KISDI, G. MESZENA AND J.A.J. METZ (1998), *Evolutionarily singular strategies and the adaptive growth and branching of the evoutionary tree*, Evol. Ecol., 12, pp. 35-57.

T. GETTY (1998a), *Handicap signalling: when fecundity and viability do not add up*, Anim. Behav., 56, pp. 127-130.

T. GETTY (1998b), *Reliable signalling need not be a handicap*, Anim. Behav., 56, pp. 253-255.

L.-A. GIRALDEAU AND T. CARACO (2000), *Social Foraging Theory*, Princeton University Press, Princeton.

R. GOMULKIEWICZ (1998), *Game theory, optimization and quantitative genetics*, in Game Theory and Animal Behavior, L.A. Dugatkin and H.K. Reeve, eds, Oxford University Press, New York, pp. 283-303.

A. GRAFEN (1984), *Natural selection, kin selection and group selection*, in Behavioural Ecology: an Evolutionary Approach 2nd ed, J.R. Krebs and N.B. Davies, eds, Blackwell Scientific Publications, Oxford, pp. 62-84.

A. GRAFEN (1990a), *Sexual selection unhandicapped by the Fisher process*, J. Theor. Biol., 144, pp. 473-516.

A. GRAFEN (1990b), *Biological signals as handicaps*, J. Theor. Biol., 144, pp. 517-546.

A. GRAFEN AND R.M. SIBLY (1978), *A model of mate desertion*, Anim. Behav., 26, pp. 642-652.

D.R. GREY (1984), *Non-negative matrices, dynamic programming and a harvesting problem*, J. Appl. Probab., 21, pp. 685-694.

G.R. GRIMMETT AND D.R. STIRZAKER (1992), *Probability and Random Processes*, Clarendon Press, Oxford.

P. HACCOU AND Y. IWASA (1995), *Optimal mixed strategies in stochastic environments*, Theor. Pop. Biol., 47, pp. 212-243.

P. HACCOU AND J.M. McNAMARA (1998), *Effects of parental survival on clutch size decisions in fluctuating environments*, Evol. Ecol., 12, pp. 459-475.

W.D. HAMILTON (1964a), *The genetical evolution of social behaviour I*, J. Theor. Biol., 7, pp. 1-16.

W.D. HAMILTON (1964b), *The genetical evolution of social behaviour II*, J. Theor. Biol., 7, pp. 17-52.

P. HAMMERSTEIN (1996), *Darwinian adaptation, population genetics and the streetcar theory of evolution*, J. Math. Biol., 34, pp. 511-532.

I. HANSKI (1988), *Four kinds of extra long diapause in insects: a review of theory and observations*, Annal. Zool. Fenni., 25, pp. 37-53.

A. HEDENSTRÖM (1992), *Flight performance in relation to fuel load in birds*, J. Theor. Biol., 158, pp. 535-537.

A. HEDENSTRÖM AND T. ALERSTAM (1995), *Optimal flight speed of birds*, Phil. Trans. Roy. Soc., B348, pp. 471-487.

M. HEINO AND V. KAITALA (1999), *Evolution of resource allocation between growth and reproduction in animals with indeterminate growth*, J. Evol. Biol., 12, pp. 423-429.

A.I. HOUSTON AND C. CARBONE (1992), *The optimal allocation of time during the diving cycle*, Behav. Ecol., 3, pp. 255-265.

A.I. HOUSTON, and N.B. DAVIES (1985), *The evolution of cooperation and life history in the dunnock, (Prunella modularis)*, in Behavioural Ecology, R.M. Sibly and R.H. Smith, eds., Blackwell Scientific Publications, Oxford, pp. 471-487.

A.I. HOUSTON, C.E. GASSON AND J.M. McNAMARA (1997b), *Female choice of matings to maximize parental care*, Proc. Roy. Soc. B, 264, pp. 173-179.

A.I. HOUSTON AND J.M. McNAMARA (1982), *A sequential approach to risk-taking*, Anim. Behav., 30, pp. 1260-1261.

A.I. HOUSTON AND J.M. McNAMARA (1985a), *A general theory of central-place foraging for single-prey loaders*, Theor. Pop. Biol, 28, pp. 233-262.

A.I. HOUSTON AND J.M. McNAMARA (1985b), *The choice of two prey types that minimises the probability of starvation*, Behav. Ecol. Sociobiol, 17, pp. 135-141.

A.I. HOUSTON AND J.M. McNAMARA (1989), *The value of food: effects of open and closed economies*, Anim. Behav., 37, pp. 546-562.

A.I. HOUSTON AND J.M. McNAMARA (1993), *A theoretical investigation of the fat reserves and mortality levels of small birds in winter*, Ornis Scand., 24, pp. 205-219.

A.I. HOUSTON AND J.M. McNAMARA (1997), *Patch choice and population size*, Evol. Ecol., 11, pp. 703-722.

A.I. HOUSTON, and J.M. McNAMARA (1999), *Models of Adaptive Behaviour*, Cambridge University Press, Cambridge.

A.I. HOUSTON, N.J. WELTON AND J.M. McNAMARA (1997a), *Acquisition and maintenance costs in the long-term regulation of avian fat reserves*, Oikos, 78, pp. 331-340.

Y. IWASA AND F.J. ODENDAAL (1984), *A theory on the temporal pattern of operational sex ratio: the active-inactive model*, Ecology, 65, pp. 886-893.

Y. IWASA AND A. POMIANKOWSKI (1999), *Good parent and good genes models of handicap evolution*, J. Theor. Biol., 200, pp. 97-109.

M. JOHNS AND R.G. MILLER (1963), *Average renewal loss rate*, Annal. Math. Stat., 34, pp. 396-401.

R.A. JOHNSTONE, J.D. REYNOLDS AND J.C. DEUTSCH (1996), *Mutual mate choice and sex differences in choosiness*, Evolution, 50, pp. 1382-1391.

A. KACELNIK AND M. BATESON (1996), *Risky theories - the effects of variance on foraging decisions*, Am. Zool., 36, pp. 402-434.

D.P. KENNEDY (1978), *On sets of countable non-negative matrices and Markov decision processes*, Adv. in App. Probab., 10, pp. 633-646.

D. KING AND J. ROUGHGARDEN (1982), *Graded allocation between vegetative and reproductive growth for annual plants in growing seasons of random length*, Theor. Pop. Biol., 22, pp. 1-16.

M. KIRKPATRICK (1985), *Evolution of female choice and male parental investment in polygynous species - the demise of the sexy son*, Am. Nat., 125, pp. 788-810.

M. KIRKPATRICK (1992), *Direct selection of female mating preferences: comments on Grafen's models*, J. Theor. Biol., 154, pp. 127-129.

M. KIRKPATRICK AND M.J. RYAN (1991), *The evolution of mating preferences and the paradox of the lek*, Nature, 350, pp. 33-38.

H. KOKKO (1999), *Cuckoldry and the stability of biparental care*, Ecol. Lett., 2, pp. 247-255.

J. KOZŁOWSKI AND A.T. TERIOKHIN (1999), *Allocation of energy between growth and reproduction: the Pontryagin Maximum Principle solution for the case of age- and season-dependent mortality*, Evol. Ecol. Res., 1, pp. 423-441.

O. LEIMAR (1996), *Life-history analysis of the Trivers and Willard sex-ratio problem*, Behav. Ecol., 7, pp. 316-325.

C.M. LESSELS (1991), *The evolution of life histories*, in Behavioural Ecology, J.R. Krebs and N.B. Davies, eds., Blackwell Scientific Publications, Oxford, pp. 32-68.

R.C. LEWONTIN AND D. COHEN (1969), *On population growth in randomly varying environments*, Proc. Nat. Acad. Sc., 62, pp. 1056-1060.

S.L. LIMA (1986), *Predation risk and unpredictable feeding conditions: determinants of body mass in birds*, Ecology, 67, pp. 377-385.

S.L. LIMA (1998), *Stress and decision making under the risk of predation: recent development from behavioral, reproductive, and ecological perspectives*, Advan. Study Behav., 27, pp. 215-290.

S.L. LIMA AND L.M. DILL (1990), *Behavioral decisions made under the risk of predation: a review and prospectus*, Can. J. Zool., 68, pp. 619-640.

J.R. LUCAS AND R.D. HOWARD (1995), *On alternative reproductive tactics in anurans: dynamic games with density and frequency dependence*, Am. Nat., 146, pp. 365-397.

J.R. LUCAS, R.D. HOWARD AND J.G. PALMER (1996), *Callers and satellites: chorus behaviour in anurans as a stochastic dynamic game*, Anim. Behav., 51, pp. 501-518.

C. MAGNHAGEN (1991), *Predation risk as a cost of reproduction*, Trends Ecol. and Evol., 6, pp. 183-186.

G.J. MAILATH, L. SAMUELSON AND A. SHAKED (2000), *Endogenous inequality in integrated labor markets with two-sided search*, Amer. Econ. Rev., 90, pp. 46-72.

M. MANGEL AND C.W. CLARK (1986), *Towards a unified foraging theory*, Ecology, 67, pp. 1127-1138.

C. MATESSI AND C. DI PASQUALE (1996), *Long-term evolution of multilocus traits*, J. Math. Biol., 34, pp. 613-653.

H. MATSUDA AND P.A. ABRAMS (1994), *Timid consumers: self-extinction due to adaptive change in foraging and anti–predator effort*, Theor. Pop. Biol., 45, pp. 76-91.

H. MATSUDA AND P.A. ABRAMS (1999), *Why are equally sized gametes so rare? The instability of isogamy and the cost of anisogamy*, Evol. Ecol. Res., 1, pp. 769-784.

J. MAYNARD SMITH (1977), *Parental investment: a prospective analysis*, Anim. Behav., 25, pp. 1-9.

J. MAYNARD SMITH (1978), *The Evolution of Sex*, Cambridge University Press, Cambridge.

J. MAYNARD SMITH (1982), *Evolution and the Theory of Games*, Cambridge University Press, Cambridge.

J. MAYNARD SMITH AND G.R. PRICE (1973), *The logic of animal conflict*, Nature, 246, pp. 15-18.

J.M. McNAMARA (1982), *Optimal patch use in a stochastic environment*, Theor. Pop. Biol., 21, pp. 269-288.

J.M. McNAMARA (1983) *Optimal control of the diffusion coefficient of a simple diffusion process*, Math. Oper. Res., 8, pp. 373-380.

J.M. McNAMARA (1984), *Control of a diffusion by switching between two drift coefficient pairs*, SIAM J. Cont., 22, pp. 87-94.

J.M. McNAMARA (1985), *An optimal sequential policy for controlling a Markov renewal process*, J. Appl. Probab., 22, pp. 324-335.

J.M. McNAMARA (1990a), *The policy which maximises long–term survival of an animal faced with the risks of starvation and predation*, Adv. in Appl. Probab., 22, pp. 295-308.

J.M. McNAMARA (1990b), *The starvation-predation trade-off and some behavioural and ecological consequences*, in Mechanisms of food selection, R.N. Hughes, ed., Springer-Verlag, Berlin, pp. 39-58.

J.M. McNAMARA (1991), *Optimal life histories: a generalisation of the Perron-Frobenius theorem*, Theor. Pop. Biol., 40, pp. 230-245.

J.M. McNAMARA (1993a), *State-dependent life-history equations*, Acta Biotheor., 41, pp. 165-174.

J.M. McNAMARA (1993b), *Evolutionary paths in strategy space: an improvement algorithm for life-history strategies*, J. Theor. Biol., 161, pp. 23-37.

J.M. McNAMARA (1994), *Timing of entry into diapause: optimal allocation to "growth" and "reproduction" in a stochastic environment*, J. Theor. Biol., 168, pp. 201-209.

J.M. McNAMARA (1995), *Implicit frequency dependence and kin selection in fluctuating environments*, Evol. Ecol., 9, pp. 185-203.

J.M. McNAMARA (1996), *Risk-prone behaviour under rules which have evolved in a changing environment*, Amer. Zool., 36, pp. 484-495.

J.M. McNAMARA (1997), *Optimal life histories for structured populations in fluctuating environments*, Theor. Pop. Biol., 51, pp. 94-108.

J.M. McNAMARA (1998), *Phenotypic plasticity in flutuating environments: consequences of the lack of individual optimisation*, Behav. Ecol., 9, pp. 642-648

J.M. McNAMARA (2000), *A classification of dynamic optimization problems in fluctuating environments*, Evol. Ecol. Res., 2, pp. 457-471.

J.M. McNAMARA AND E.J. COLLINS (1990), *The job search problem as an employer-candidate game*, J. Appl. Probab., 28, pp. 815-827.

J.M. McNAMARA, C. GASSON AND A.I. HOUSTON (1999), *Incorporating rules for responding into evolutionary games*, Nature, 401, pp. 368-371.

J.M. McNAMARA AND A.I. HOUSTON (1986), *The common currency for behavioral decisions*, Am. Nat., 127, pp. 358-378.

J.M. McNAMARA AND A.I. HOUSTON (1990), *The value of fat reserves and the trade-off between starvation and predation*, Acta Biotheor., 38, pp. 37-61.

J.M. McNAMARA AND A.I. HOUSTON (1992), *Risk-sensitive foraging: a review of the theory*, Bull. Math. Biol., 54, pp. 355-378.

J.M. McNAMARA AND A.I. HOUSTON (1994), *The effect of a change in foraging options on intake rate and predation rate*, Am. Nat., 144, pp. 978-1000.

J.M. McNAMARA AND A.I. HOUSTON (1996), *State-dependent life histories*, Nature, 380, pp. 215-220.

J.M. McNAMARA, R.H. MACE AND A.I. HOUSTON (1987), *Optimal daily routines of singing and foraging in a bird singing to attract a mate*, Behav. Ecol. Sociobiol., 20, pp. 399-405.

J.M. McNAMARA, T. SZÉKELY, J.N. WEBB AND A.I. HOUSTON (2000), *A dynamic game theoretic model of parental care*, J. Theor. Biol., 205, pp. 605-623.

J.M. McNAMARA, J.N. WEBB AND E.J. COLLINS (1995), *Dynamic optimisation in fluctuating environments*, Proc. Roy. Soc. B, 261, pp. 279-284.

J.M. McNAMARA, J.N. WEBB, E.J. COLLINS, T. SZÉKELY AND A.I. HOUSTON (1997), *A general technique for computing evolutionarily stable strategies based on errors in decision-making*, J. Theor. Biol., 189, pp. 211-225.

J.M. McNAMARA, R.K. WELHAM AND A.I. HOUSTON (1998), *The timing of migration within the context of an annual routine*, J. Avian Biol., 29, pp. 416-423.

J.A.J. METZ, R.M. NISBET, AND S.A.H. GERITZ (1992), *How should we define "fitness" for general ecological scenarios*? Trends Ecol. and Evol., 7, pp. 198-202.

R.E. MICHOD AND B.R. LEVINS (eds), *The Evolution of Sex, an Examination of Current Ideas*, Sinauer, Sunderland MA.

M. MILINSKI AND G.A. PARKER, *Competition for resources*, in Behavioural Ecology, (3rd ed), J.R. Krebs and N.B. Davies eds, Blackwell Scientific Publications, Oxford, pp. 137-168.

D.W. MOCK AND G.A. PARKER (1997), *The Evolution of Sibling Rivalry*, Oxford University Press, Oxford.

P.A.P. MORAN (1964), *On the nonexistence of adaptive topographies*, Ann. Hum. Genet., 27, pp. 383-393.

S.D. MYLIUS AND O. DIEKMANN (1995), *On evolutionarily stable life histories, optimization and the need to be specific about density dependence*, Oikos, 74, pp. 218-224.

S.H. ORZACK AND S. TULJAPURKAR (1989), *Population dynamics in variable environments VII. The demography and evolution of iteroparity*, Am. Nat., 133, pp. 901-923.

G.A. PARKER (1978), *Searching for mates*, in Behavioural Ecology, J.R. Krebs and N.B. Davies, eds., Blackwell Scientific Publications, Oxford, pp. 214-244.

G.A. PARKER (1998), *Sperm competition and the evolution of ejaculates: towards a theory base*, in Sperm Competition and Sexual Selection, T.R. Birkhead and A.P. Møller, eds, Academic Press, San Diego, pp. 3-54.

G.A. PARKER, R.R. BAKER AND V.G.F. SMITH (1972), *The origin and evolution of gamete dimorphism and the male-female phenomenon*, J. Theor. Biol., 36, pp. 529-533.

N. PERRIN AND R.M. SIBLY (1993), *Dynamic models of energy allocation and investment*, Ann. Rev. of Ecol. and Syst., 24, pp. 379-410.

N. PERRIN, R.M. SIBLY AND N.K. NICHOLS (1993), *Optimal growth strategies when mortality and production rates are size-dependent*, Evol. Ecol., 7, pp. 576-592.

A. POMIANKOWSKI (1987), *Sexual selection: the handicap principle does work – sometimes*, Proc. Roy. Soc. B, 231, pp. 123-145.

L. REAL (1990), *Search theory and mate choice. I. Models of single-sex discrimination*, Am. Nat., 136, pp. 376-404.

D.A. ROFF (1992), *The Evolution of Life Histories*, Chapman and Hall, New York.

M.J. RYAN (1998), *Sexual selection, receiver biases, and the evolution of sex differences*, Science, 281, pp. 1999-2003.

A. SASAKI AND S. ELLNER (1995), *The evolutionary stable phenotype distribution in a random environment*, Evolution, 49, pp. 337-350.

W.M. SCHAFFER (1974), *Selection for optimal life histories: the effects of age structure*, Ecology, 55, pp. 291-303.

J. SEGER AND H.J. BROCKMANN (1987), *What is bet-hedging?* in Oxford Surveys in Evolutionary Biology, Vo. 4, P.H. Harvey and L. Partridge, eds, Oxford University Press, Oxford, pp. 182-211.

R. SELTEN (1980), *A note on evolutionarily stable strategies in asymmetric animal conflicts*, J. Theor. Biol., 84, pp. 93-101.

S.C. STEARNS (1992), *The Evolution of Life Histories*, Oxford University Press, Oxford.

M.C. STEINBACK (2001), *Markowitz revisited: Mean-variance models in financial portfolio analysis*, SIAM Review 43, pp ??-??

D.W STEPHENS (1981), *The logic of risk-sensitive foraging preferences*, Anim. Behav., 29, pp. 628-629.

D.W. STEPHENS AND J.R. KREBS (1986), *Foraging Theory*, Princeton University Press, Princeton.

W.J. SUTHERLAND (1996), *From Individual Behaviour to Population Ecology*, Oxford University Press, Oxford.

S. SZÁMADÓ (1999), *The validity of the handicap principle in discrete action-response games*, J. Theor. Biol., 198, pp. 593-602.

H.M. TAYLOR, R.S. GOURLEY, C.E. LAWRENCE AND R.S. KAPLAN (1974), *Natural selection of life history attributes: an analytical approach*, Theor. Pop. Biol., 5, pp. 104-122.

P.D. TAYLOR (1989), *Evolutionary stability in one-parameter models under weak selection*, Theor. Pop. Biol., 36, pp. 125-143.

P.D. TAYLOR (1990), *Allele frequency change in a class-structured population*, Am. Nat., 135, pp. 95-106.

P.D. TAYLOR AND S.A. FRANK (1996), *How to make a kin selection model*, J. Theor. Biol., 180, pp. 27-37.

N. TINBERGEN (1963), *On aims and methods of ethology*, Z. Tierpsychol., 20, pp. 410-433.

T. TREGENZA (1995), *Building on the ideal free distribution*, Adv. Ecol. Res., 26, pp. 253-307.

R.L. TRIVERS AND D.E. WILLARD (1973), *Natural selection of parental ability to vary the sex ratio of offspring*, Science, 179, pp. 90-92.

S.D. TULJAPURKAR (1989), *An uncertain life: Demography in random environments*, Theor. Pop. Biol., 35, pp. 227-294.

J.N. WEBB, A.I. HOUSTON, J.M. McNAMARA AND T. SZÉKELY (1999), *Multiple patterns of parental care*, Anim. Behav., 58, pp. 983-993.

F. WEISSING (1991), *Evolutionary stability and dynamic stability in a class of evolutionary normal form games*, in Game Equilibrium Models. I. Evolution and Game Dynamics, R. Selten, ed., Springer-Verlag, Berlin, pp. 29-97.

F.J. WEISSING (1996), *Genetic versus phenotypic models of selection: can genetics be neglected in a long-term perspective*? J. Math. Biol., 34, pp. 533-555.

P. WHITTLE (1983), *Optimization over Time*, *Volume 2*, Wiley, Chichester.

M.S. WITTER AMD I.C. CUTHILL (1993), *The ecological costs of avian fat storage*, Phil. Trans. Roy. Soc., B, 340, pp. 73-92.

N. YAMAMURA AND N. TSUJI (1993), *Parental care as a game*, J. Evol. Biol., 6, pp. 103-127.

A. ZAHAVI (1975), *Mate selection – a selection for a handicap*, J. Theor. Biol., 53, pp. 205-214.

Figure Captions

Figure 1. Male southern elephant seals (*Mirounga leonina*) fighting for access to females during the breeding season. (Photograph by courtesy of Ian Boyd.)

Figure 2. An illustration of the constructive computation of the optimal net rate of energetic gain γ^*.

Figure 3. An illustration of the difference between the time at the surface $S(d(t^*(\tau),\tau))$ as a function of time underwater $t^*(\tau)+\tau$ when the diver chooses the optimal time foraging t^* and the relationship when $S(d(t,\tau))$ is plotted against $t+\tau$ for a fixed value of τ. The figure shows $\tau = 2$ and $\tau = 6$; $d(t,\tau) = t + 4\tau$, $S(d(t,\tau)) = 2\ln[K/(K - d(t,\tau))]$, $K = 20$.

Figure 4. The marginal rate of substitution of predation risk for energy, θ, as a function of energy reserves in four cases. (i) Deterministic foraging with reproduction at a critical level L. (ii) Stochastic foraging with reproduction at the critical level L. (iii) Stochastic foraging with no reproduction; food that would take reserves above L is lost. (iv) Stochastic foraging with no reproduction; food that would take reserves above L is consumed in a safe refuge. $L = 15$, $\gamma = 0.5$, rate of mortality $= 0.5$. See Houston and McNamara (1999) for details.

Figure 5. (a) The optimal foraging strategy together with the distribution of reserves. It is optimal to forage if reserves are below the bold line and to rest if reserves are above it. The broken line gives the mean level of reserves. (b) The resulting proportion of birds that are foraging. See Houston and McNamara (1999) for details.

Figure 6. Pattern of care over the breeding season together with various aspects of the rewards in the model of McNamara et al. (2000). (a) The pattern of care together with the remating probabilities for each sex; (C, C) = care by both the male and female, (D, C) = desertion by the male and care by the female. (b) The advantage to the female of deserting as opposed to caring (future reproductive success if the female deserts minus future reproductive success if the female cares) for each of the possible male behaviors. (c) The payoff to the male at the ESS, minus the payoff if both parents desert. See McNamara et al. (2000) or Houston and McNamara (1999) for details.

Figure 7. A pair of red-necked phalaropes (*Phalaropus lobatus*). In the phalarope, the males care for the young and have a lower potential rate of reproduction than females. Sex rôles are reversed: the females are brightly coloured and compete for access to males. (Photograph by courtesy of John Reynolds.)

Figure 8. An illustration of the partition of male and female qualities into discrete classes $M_1, M_2, \ldots$ and $F_1, F_2, \ldots$ respectively. Females in each class accept males above the corresponding female acceptance threshold. Similarly, males in each class accept females above the corresponding male acceptance threshold.

Figure 9. Effect on breeding of the amplitude ϵ of the seasonal variation in the food supply at the temperate site. When $\epsilon = 0.4$, there is one breeding attempt per year, whereas when $\epsilon = 0.8$ there are two attempts. See McNamara et al. (1998) for details.