



## REVIEWS

## Ways to test stochastic dynamic programming models empirically

JOHN M. C. HUTCHINSON &amp; JOHN M. McNAMARA

School of Mathematics, University of Bristol

*(Received 15 April 1999; initial acceptance 26 May 1999;  
final acceptance 7 December 1999; MS. number: RV-17)*

Stochastic dynamic programming (SDP) models are widely used to predict optimal behavioural and life history strategies. We discuss a diversity of ways to test SDP models empirically, taking as our main illustration a model of the daily singing routine of birds. One approach to verification is to quantify model parameters, but most SDP models are schematic. Because predictions are therefore qualitative, testing several predictions is desirable. How state determines behaviour (the policy) is a central prediction that should be examined directly if both state and behaviour are measurable. Complementary predictions concern how behaviour and state change through time, but information is discarded by considering behaviour rather than state, by looking only at average state rather than its distribution, and by not following individuals. We identify the various circumstances in which an individual's state/behaviour at one time is correlated with its state/behaviour at a later time. When there are several state variables the relationships between them may be informative. Often model parameters represent environmental conditions that can also be viewed as state variables. Experimental manipulation of the environment has several advantages as a test, but a problem is uncertainty over how much the organism's policy will adjust. As an example we allow birds to use different assumptions about how well past weather predicts future weather. We advocate mirroring planned empirical investigations on the computer to investigate which manipulations and predictions will best test a model.

© 2000 The Association for the Study of Animal Behaviour


Stochastic dynamic programming (SDP) is a computational technique that finds optimum sequences of actions. Since its promotion in behavioural ecology by McNamara & Houston (1986) and Mangel & Clark (1986), SDP modelling has been applied to a wide variety of phenomena, for instance: foraging (McNamara & Houston 1992); food caching (Lucas & Walter 1991); antipredation behaviour (Houston et al. 1993); intraspecific fighting and other games (Enquist & Leimar 1987; Houston & McNamara 1988); movement and migration (Weber et al. 1998; Hutchinson 1999); mating tactics (Lucas & Howard 1995; Galvani & Johnstone 1998); and life history decisions about growth, reproduction and sex change (Iwasa 1991; McNamara & Houston 1996; Hutchinson et al. 1997). Mangel & Clark (1988) and Houston & McNamara (1999) provide further examples. One strength of SDP models is that they can be quite complex, yet computationally tractable.

Correspondence: J. M. C. Hutchinson, School of Biological Sciences, University of Bristol, Woodland Road, Bristol BS8 1UG, U.K. (email: [john.hutchinson@bristol.ac.uk](mailto:john.hutchinson@bristol.ac.uk)).

The philosophy and aims of optimality modelling in behavioural ecology have been discussed elsewhere (e.g. Stearns & Schmid-Hempel 1987; Parker & Maynard Smith 1990). The purpose of some SDP models has been only to prove the logical coherence of verbal arguments. But often we want to test whether models fit observed biological phenomena. In this paper we aim to show that in this last respect SDP has not been used to its full potential. SDP can provide a variety of predictions that, if tested in combination, would yield more compelling evidence of whether a model is realistic.

SDP models have stimulated empirical experiments designed to test their predictions (e.g. Ekman & Hake 1990; Lucas & Walter 1991; Witter et al. 1995; Reinhardt & Healey 1999; Thomas 1999a, b). However, the number of parameters in some more realistic models can make them daunting to apply. Fortunately it is often unnecessary to know every parameter value accurately. Complex SDP models can produce simple predictions that are robust to a wide range of parameter values.

We start this paper with a nontechnical explanation of SDP models, including the definition of some standard



**Figure 1.** A daily routine of avian song generated by an SDP optimality model. Each day is divided into 150 intervals (dawn at 0, dusk at 75). —: Singing routine as a percentage of birds singing; ---: switchpoint levels of reserves above which it is optimal to sing; ···: mean value of reserves for those birds that are alive and unpaired at the time. The units of reserves are arbitrary. The parameter values, using the symbols defined in [Hutchinson et al. \(1993\)](#), are:  $e=28$ ,  $p_{\text{food}}=0.8$ ,  $S_0=10$ ,  $S_1=1 \times 10^{-3}$ ,  $S_2=2 \times 10^{-6}$ ,  $\eta_s=4.2$ ,  $q_s=0.5$ ,  $F_0=10$ ,  $F_1=2 \times 10^{-3}$ ,  $F_2=4 \times 10^{-6}$ ,  $\eta_r=6.4$ ,  $q_r=0.5$ ,  $R_{\text{cold}}=4.1$ ,  $R_{\text{normal}}=3.6$ ,  $R_{\text{hot}}=3.3$ ,  $q_{\text{temp}}=0.8$ ,  $\eta_r=3.2$ ,  $q_r=0.5$ ,  $p_{\text{pair}}=4 \times 10^{-3}$ , mean duration of pairing window=10 days.

terms. We then describe a model of daily singing routines in birds, which we use to illustrate points of general applicability throughout the rest of this paper. We also detail what sorts of empirical data can be used to test SDP models.

### AN EXPLANATION OF SDP MODELS

SDP models characterize individuals by their 'state'. Examples of state variables include an individual's energy reserves, its size, the number of its dependent offspring, current environmental temperature, and its estimate of the current food availability. Several state variables may be required to describe state adequately for the purpose of the model, but in many SDP models a single state variable has sufficed.

SDP models divide time into discrete intervals. In each the organism chooses an 'action'. The actions might be behavioural (e.g. to forage or to sing), or physiological or morphological (e.g. to moult). The action chosen, together with the individual's state, affects the probability that it survives the time interval, its state in the next interval, and the number of surviving offspring produced during the interval. Some state variables are specified to change stochastically.

Constructing an SDP model thus requires specifying the state variables, the available actions, and the short-term consequences of taking each action in each state. The role of dynamic programming is to find the optimum behaviour given an additional specification of what should be maximized in the long term (the 'terminal reward', e.g. long-term survival, number of offspring, or energy reserves). Dynamic programming works backwards in time to calculate the optimal 'policy' ([Mangel & Clark 1988](#); [Houston & McNamara 1999](#)). This policy is the rule specifying the optimal action when in each state at each time. Using this policy we can follow an individual forward in time; at each time interval its current state

determines what action it takes, and thus the probability of being in each state in the next time interval. The sequence of actions constitutes the behavioural 'routine'. Because state changes stochastically, routines vary even though all individuals use the same policy.

The routine also depends on the state in which an individual started. But when animals have settled into a stable or cyclic environment, stochasticity will have caused the distribution of states to converge on a distribution determined by the policy and the environment, and independent of initial state. (In the case of a cyclic environment, such as a day–night cycle, this distribution will change within a cycle, but recur at any particular stage of the cycle.) Modellers often report only an average routine after this convergence, giving the proportion of the surviving population taking each action at each time.

### SONG EXAMPLE

We use as an example a simple SDP model that explains why birds sing most at dawn and dusk ([Houston & McNamara 1987](#); [McNamara et al. 1987](#); [Hutchinson et al. 1993](#)). The model assumes that during the daytime a male bird chooses between two actions, singing and foraging. Singing gives a chance of attracting a mate, which brings a higher terminal reward than mere survival. At night birds can only rest, and lose energy keeping warm. The transition between night and day is sudden, and within each the environment remains constant.

The single state variable is energy reserves. Reserves accumulate through foraging, and fall while birds rest or sing. These changes include a stochastic component. Thus, even if a bird forages it may fail to find food, and its reserves fall. The unpredictability of overnight energy requirements is accentuated by temperature variation between nights. If reserves fall to zero the bird has starved to death. High reserves have a cost, making the bird

more prone to predation, or making activity (particularly flight) energetically more expensive.

Figure 1 shows the optimal policy. It specifies singing when reserves are above a 'switchpoint' and foraging when below. This switchpoint changes with time of day. In the morning the switchpoint is low, keeping reserves just sufficient almost to eliminate the possibility of starvation should a run of feeding attempts prove unsuccessful. Later in the day the switchpoint rises steadily to ensure that the bird has sufficient reserves at dusk not to starve overnight. Dusk reserves are enough to survive an almost worst-case night, so most nights leave reserves at dawn well above the switchpoint, causing birds to sing a dawn chorus.

In this paper we use two embellishments to this basic model to illustrate particular points. The first supposes that females prefer continuous bouts of song over the same amount of song delivered intermittently. This would enable females to select better mates (Hutchinson et al. 1993). An extra state variable represents a female's memory of how much a male has sung in the recent past. The variable increases when the male sings and decreases when he forages. The optimal policy now specifies how behaviour depends on both energy reserves and the memory variable.

The other embellishment concerns overnight temperature. The basic model supposed that three different overnight temperatures were possible and that the temperature of the following night was unpredictable. The optimal policy did not depend on overnight temperature, because during the day the temperature of the preceding night predicted nothing about the succeeding conditions, and at night the only action available was to rest. In our embellished model temperatures of successive nights are correlated. The optimal policy during the day then depends on the previous night's temperature because this is a predictor of the next night's temperature. Overnight temperature must now be modelled as a state variable.

## TESTING MODEL ASSUMPTIONS

Modellers typically have to make informed guesses about the biology (often the first suggestion that particular facets of the biology are worth measuring comes from the model). Subsequently showing that these assumptions are inappropriate for particular organisms is as logical a reason to reject the model's relevance as its predictions not fitting. Sometimes a key assumption is completely wrong, but lack of realism is often a matter of degree, so checking assumptions means measuring parameter values.

Models should always undergo a sensitivity analysis in which the effects of altering assumptions and parameter values are systematically investigated (Gladstein et al. 1991; Houston et al. 1992). This might concentrate on the least well established parameter values (e.g. Clark & Ydenberg 1990). The sensitivity analysis may reveal that output is very insensitive to variation in some parameters, so measuring their values would be pointless. Conversely we might learn that accurate predictions are impossible without gaining more information about

other parameters, which should be the stimulus to measure them.

Often the sensitivity analysis shows that realistic predictions occur only with particular parameter values. For instance in the singing routine model, with some parameter values for the food supply birds sang little at dawn and song steadily increased through the day (Figs 4 and 5 in Hutchinson et al. 1993), which is not typical for real birds. For pairing and mortality rates to be realistic a restricted range of values of other parameters was required. But choosing parameters a posteriori to achieve realism allows only a demonstration that the model is a possible explanation of an observed phenomenon; evidence for the model is weak unless other predictions of the model are confirmed. One sort of test is to measure parameter values hitherto selected on this basis of producing realistic output. This is a form of inverse optimality argument.

We may sometimes reject the parameter set without measuring any parameters individually. For instance, we could directly measure metabolic rates, the calorific value of different food items and the probability of finding each. Alternatively, as an easier but less complete check, we could measure how reserves change when foraging, and on this basis reject the parameter set as a whole.

Some parameters can be very difficult to measure. An important example is parameters describing how state and actions affect predation rate. Many SDP models of bird behaviour depend on the assumption that the predation risk associated with a behaviour increases with energy reserves. Unfortunately models show that the form of the relationship can be crucial (e.g. whether state and vigilance interact additively or multiplicatively to affect predation rate: Houston et al. 1993). One difficulty in quantifying predation is simply its rarity. Yet, despite their rarity, unusually bad events (such as predation, an unusually cold night, or a run of bad luck when foraging) often have a dominant effect on the policy. Experimental manipulations that increase predation may be considered unethical, and direct measurements of predation rate confound predation risk with strategic alterations of behaviour by the bird (McNamara & Houston 1987; McNamara 1990). An alternative approach has been to estimate from aerodynamic principles, or from aviary trials, how reserves affect flight performance. But flight performance is an indirect and incomplete estimate of predation risk, so the relationships of interest are only crudely inferred (Witter et al. 1994; Bednekoff 1996).

How predation risk relates to reserves also depends on the numbers of different types of predator in the environment. This is itself difficult for us to estimate but to predict behaviour successfully may require knowing the bird's own estimate, which may still be adapted to predation risk in the past (e.g. Newman et al. 1995 discussed this with respect to domestic sheep, *Ovis aries*). Generally parameter values in a laboratory study may be easier for us to measure; but the organism's policy may still be partially adapted to the different parameter values in the wild (Houston & McNamara 1989).

Another difficulty with measuring parameter values is that it is usually not possible to make the organism use a

different policy. For instance, one cannot check how foraging would change reserves when reserves exceed the level at which birds choose to forage. Thus we can check only a subset of the state dynamics assumed. However, sometimes behaviour can be manipulated; for instance, we can delay an animal mating by preventing access to a mate. Even if behaviour cannot itself be manipulated, manipulating state away from its usual values does help to provide a more extensive check of state dynamics. Avian clutch size manipulations are a classic example. However, changing the state (number of eggs) should not be equated with a change in the behaviour that determines the state (egg laying), since egg laying affects other state variables, such as energy reserves (Monaghan et al. 1998).

It has rarely been possible for all parameter values in an SDP model to be based on accurate estimates (Mace 1993 and Heimpel et al. 1998 are possible exceptions). Furthermore, even a model using only accurate values would still disregard other aspects of the biology: models are simplifications. So we probably should not expect a quantitative fit of the predictions, and thus accurate measurements of parameter values may be inappropriate. For these reasons, and also so as not to restrict their generality, most SDP models have not even set out to be quantitatively accurate. Thus our singing routine model was not based on any particular species; we tried merely to ensure that the parameters' relative values were reasonable for a small bird. Overnight temperature was considered to take only three possible values, which is obviously unrealistic, but it captured the essential idea that overnight energy consumption could occasionally and unpredictably be higher than average. The model's purpose was to demonstrate a principle, not produce quantitative predictions, but nevertheless our sensitivity analyses showed that some qualitative predictions were robust over a broad range of parameter values. Throughout the rest of this paper it is the testing of such robust qualitative predictions that we envisage. Sometimes the predictions may concern merely the direction of a relationship, but in other cases more specific predictions can be made about the relationship's shape, as with the relationship of mass with time shown in Figure 1. Even so, it is harder to falsify qualitative predictions than quantitative ones, which is a reason to test several different predictions.

### THE RANGE OF TESTABLE PREDICTIONS

Checking parameter values can only be a partial test of a model: even if all parameter values are correct, the organism may not use the predicted optimal policy. This might be because of the time lag in the evolutionary process, because the actions have extra unrecognized consequences, or because of unrecognized constraints in what actions are possible or to what state variables the organism can respond. So one must also check the model's predictions. If these are confirmed, this also provides evidence that the parameter values were realistic. On the other hand, without checking parameter values the model could be giving the right results for the wrong

reasons, especially if parameter values have been rigged a posteriori to engineer a realistic output.

We have explained above how SDP models can predict the average behavioural routine. This has been the prediction most often tested, probably because this is the sort of empirical observation that originally inspires a model's construction. We now explain the value of testing other predictions.

### Testing Policy

The policy specifies what action to take when in each particular state at each particular time. All SDP models make the fundamental assumptions that individuals consistently follow such a rule and that different individuals use the same rule. This consistency is not true of the behavioural routine, which varies between individuals and from day to day in the same individual. Consequently, predicted routines are given as averages, which means that such predictions can only be probabilistic, and therefore may require more observations to disprove.

SDP models normally predict deterministic policies. (Probabilistic policies may be optimal in game-theoretical situations, but seem likely to be realized via deterministic strategies contingent on other cues.) When actions are all or nothing (such as whether to reproduce in a season, or whether to lay one egg or two), a deterministic policy typically entails one or more switchpoints: one action is optimal if state lies above this value, another action if below. With other behaviours, such as vigilance, the optimal policy can be a continuous adjustment of behaviour in response to state (Houston et al. 1993). One might envisage that song rate in birds can be continuously adjusted as a function of reserves. However, here we allow only two options, to sing or to forage, and this approximation makes little difference to the routines.

Whether the options available are discrete or continuous, the form of the policy is often monotonic (e.g. feed if reserves are below a switchpoint, otherwise sing). So the simplest sort of qualitative, testable prediction about the policy is that at a given time of day birds are more likely to sing if they have higher reserves.

A more critical evaluation would examine whether the observed policy is consistent. This is easiest when the predicted policy has a sharp switchpoint. If the switchpoint appears to vary between individuals, or when the same individual is tested repeatedly, the most likely conclusion is that the model ignores a state variable that influences the organism's decision. (For instance, a bird's decision whether to sing might be influenced by food availability in its territory, its neighbour's song output, or the temperature expected the next night.) This conclusion is strongly suggested if individuals in a particular state that choose action A consistently turn out to have a lower fitness than individuals in the same state that choose action B. For example, Clutton-Brock et al. (1996) found that Soay sheep of the same age and weight class were more likely to survive some winters if they reproduced. It seems implausible that reproduction itself enhances survival, so probably the sheep differed in some unconsidered aspect of state that affected their decisions.



Even if we do measure all relevant state variables, difficulties in measuring state and behaviour completely accurately and synchronously will mean that empirical data will rarely suggest a perfectly sharp switchpoint. Also the organism itself may be mismeasuring state variables. For instance, models often predict a dependence of behaviour on time since an event, time of day, or time of year, but we know empirically that animals do not measure time with perfect accuracy (Gibbon & Allan 1984; Brunner et al. 1992).

If we feel satisfied that a model has captured the essence of the relationship between behaviour and one state variable at a particular time, we should go on to examine how this relationship changes depending on time of day or on the value of other state variables. For instance our singing routine model predicts that the switchpoint is low in the morning and rises through the afternoon (Fig. 1), so we would predict that at some reserve levels at which the bird sings in the morning it would forage in the afternoon. If the overnight temperature shows a correlation between nights, we predict that the switchpoint at dusk should be higher after cold nights (Table 1). Similarly Bull et al. (1996) experimentally confirmed a prediction of an SDP model that the readiness of salmon, *Salmo salar*, to feed when lipid levels were low was higher in early winter than late. Reinhardt & Healey (1999) also used a salmon species, *Onchorhynchus kisutch*, to test predictions that size and nutritional state should both affect willingness to run a predation risk.

To observe the policy directly both state and behaviour must be measured, but the advantages of doing this have not always been appreciated. Several tests of models of winter fattening in birds (e.g. Ekman & Hake 1990; Witter et al. 1995) have been laboratory based and measured the state variable (weight), but ignored the behaviour (foraging). In contrast, tests of the similar models of singing routines (Thomas 1999a, b, and references therein) have been field based and measured one easy-to-monitor behaviour (song), but ignored the other behaviour (foraging) and the state (weight). In both situations it would be worth trying to measure both state and behaviour, so as to gauge the policy. Laboratory studies of hoarding in birds (Lucas & Walter 1991; Hurly 1992) have measured both behaviour and state. For instance, Lucas & Walter related the percentage of seeds cached not merely to time of day, but also to dawn mass, the food access regime, season and body size.

The values of certain parameters may have little effect on the optimal policy, so even if the policy observed fits the predictions, the model need not necessarily fit in other respects. Such parameter values must be checked by direct measurement, or by examining other predictions such as the routine.

### Distribution of State as a Function of Time

The average behavioural routine describes how the percentage of the population carrying out each action changes over time. Calculating this routine entails first calculating the distribution of states in the population at each time. But this distribution of states can be a

more revealing prediction than the behavioural routine when the same action is appropriate for a range of states (e.g. if the choice of actions is not continuous). Then, even though the distribution of states might not agree with the model, the proportion of animals performing each behaviour at a particular time might not disagree significantly.

The kinds of qualitative predictions that we might test are whether the distribution of states at a particular time is skewed or bimodal, and how the average state, or the variance in state, changes from one time to another. In the singing routine example, it is too trite a prediction that reserves are higher at dusk than at dawn and fall further during the dawn chorus (reserves inevitably fall while the animal is not foraging), but predictions about the variance and skew of reserves at different times of day are less obvious.

Figure 2a shows the distributions of reserves associated with the singing routine in Figure 1. Variation at dawn is large and symmetrical because it is mostly generated by variation in overnight energy requirements, which the model specifies to follow a roughly normal distribution. By mid-morning the variance in reserves has decreased because overweight birds have lost their excess by singing, and underweight birds have been foraging. The distribution is now less symmetrical, because singing always lowers the reserves of overweight birds, whereas underweight birds are sometimes unsuccessful in increasing their reserves by foraging. In the afternoon the asymmetry and variance increase, because birds are more frequently choosing to forage, the action with the more variable outcome. Birds achieve a much narrower range of reserve levels with which to start the night.

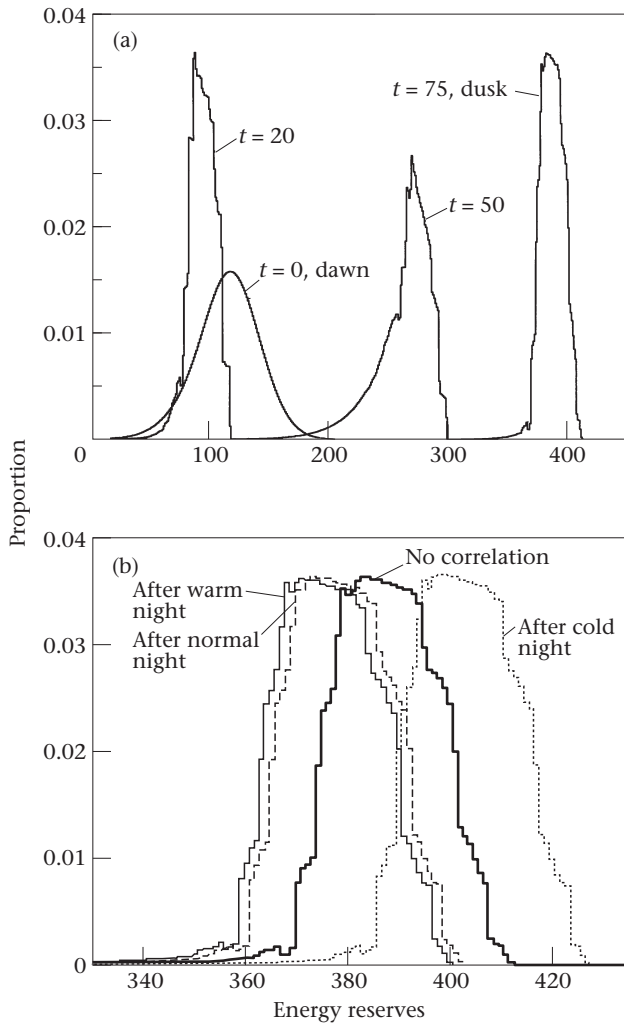
These distributions are averages over nights of all temperatures. Many birds measured after the same night should show a much lower variance in their reserves at dawn. In contrast, the distributions of reserves later in the day are unaffected by the temperature of the preceding night.

The latter is not true if temperatures of successive nights correlate, allowing birds to improve their estimate of the forthcoming overnight energy loss. In that case the model predicts that on most days their afternoon and evening reserves are lower than in Figure 2a, except after a cold night when their reserves are higher because they predict another cold night (Fig. 2b). Indeed some species of bird do change their evening body mass in response to recent temperatures (Wansink & Tinbergen 1994; Rogers 1995).

### Correlations between Different State Variables

In the last example overnight temperature acts as a state variable. Averaging over all values of this state variable gives valid predictions for data collected over many days; but more critical tests are possible if we measure temperatures of particular nights in order to make predictions about the distribution of reserves specific to the following day (Fig. 2b).

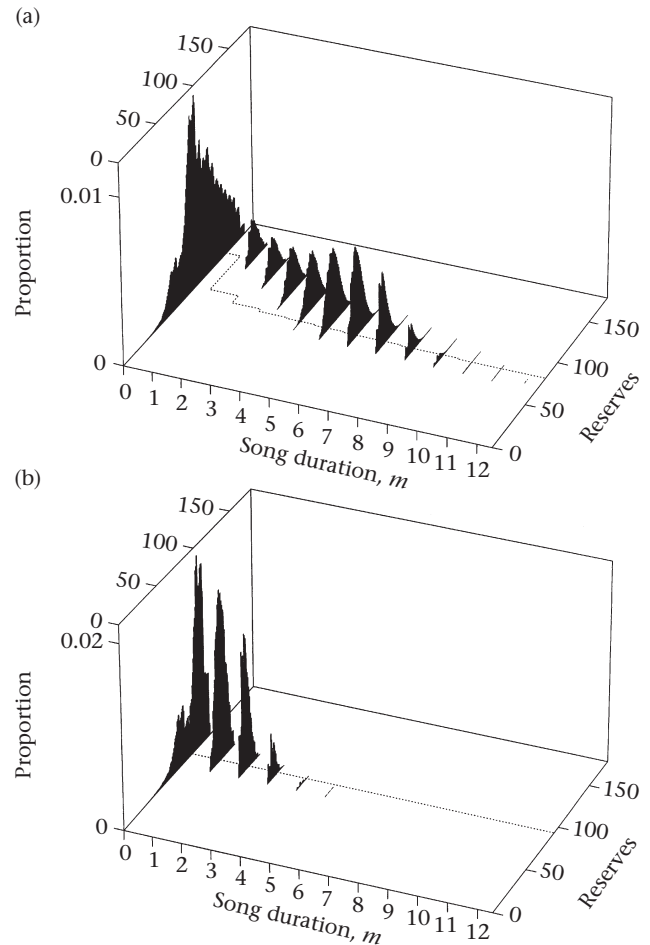
In general when there are two state variables  $x$  and  $y$ , SDP models can predict not only the distribution of  $x$



**Figure 2.** The distribution of energy reserves (shown as the proportion of the surviving unpaired population in each state) for birds following the optimal routine shown in Figure 1. (a) Distributions at different times of day ( $t$ ). Each curve is a weighted average of the three distributions after nights of different temperatures. (b) The distribution of reserves at dusk under various different assumptions. The thick solid line is for the same model as (a), i.e. no correlation between temperatures of successive nights; at dusk this distribution is the same whatever the past night-time temperature. Other lines are for a model in which successive night-time temperatures are highly correlated. (The rule is effectively this. First toss a biased coin. If it is heads ( $P=0.25$ ), the probability of each temperature follows the same overall probabilities as in (a); if it is tails ( $P=0.75$ ), the temperature is the same as last night.)

given a specific value of  $y$ , or averaged over all values of  $y$ , but also the bivariate distribution of  $x$  and  $y$ . In the above example one might display the data collected after nights of different temperatures as a scattergram of reserves plotted against temperature. As with the univariate distributions, the bivariate distribution will depend on which time is considered.

As another example of two state variables we consider energy reserves and an index that remembers song output over the recent past. Here we take this memory variable ( $m$ ) to be simply the number of time intervals of



**Figure 3.** Bivariate plots of energy reserves against the duration of unbroken song ( $m$ ), measured at the start of the 20th time interval after dawn. The vertical axis shows the proportion of the surviving unpaired population in each state. In (a) the probability of pairing  $= 0.004 m^{0.01}$ , whereas in (b) it always  $= 0.004$  if  $m > 0$ . Other parameters and model details are as in Figure 1. The dotted lines show the switchpoint levels of reserves above which birds should sing.

continuous song leading up to the present interval (one period of foraging resets  $m$  to 0). So we could directly measure  $m$ . We suppose that the probability of pairing is slightly higher if the male has been singing for a long time continually (probability  $\propto m^{0.004}$ ). Figure 3a is a plot of reserves against  $m$  at one time mid-morning. Unsurprisingly, the longer a bird has been singing the lower its reserves. When  $m=0$  reserves can take a wide range of values as birds feed in order to be able to sing for a long time subsequently. When  $m=0$  they forage at levels of reserves at which they continue to sing when  $m$  is larger.

Compare Figure 3b, which is generated by parameters identical to Figure 3a except that the probability of pairing is now independent of  $m$ . Reserves still tend to be lower in birds that have been singing for longer, but this is much less marked because the switchpoint level of reserves at which the bird sings instead of forages is now independent of  $m$ . Because Figure 3a and b differ, plotting

empirical data in this way could provide evidence for whether birds do attempt to sing in continuous bouts.

Note that the shape of the relationship between two state variables will depend on the source of variation. In Figure 3 the main source of variation is the stochasticity in foraging success and metabolic costs. But one could also imagine that some individuals hold territories enabling a consistently better foraging success. A bivariate plot of reserves against  $m$  will be transposed downwards for higher-quality territories. Amalgamating the plots from a variety of territory qualities will thus at least smudge the distributions in Figure 3, and could change the main axis of variation.

Another common sort of plot is between the timing of some event and the state in which it occurs. For instance with life histories one can use dynamic programming models to produce norms of reaction relating the age at maturation against the size at maturation (e.g. Rowe & Ludwig 1991; Hutchinson et al. 1997). Typically the shape of that relationship will depend on the source of variation. In one model of arthropod growth, variation in foraging success caused maturation at an old age to be associated with large adult size, whereas variation in egg size caused maturation at an old age to be associated with small adult size (Fig. 5c in Hutchinson et al. 1997). Empiricists measuring age and size at maturation might not know the main source of variation in their system. One way round this difficulty is an experimental manipulation.

### Comparing the Same Individual at Different Times

To test an average behavioural routine qualitatively we would test whether an action occurs more often at one time than another; similarly, we have discussed testing whether a state variable is higher at one time than another. In both cases we envisaged comparing population averages. An alternative approach is to measure state or behaviour twice in each individual and then to plot the later measurement against the earlier, with each point representing one individual. This reveals information that is not apparent from average behavioural routines, or from changes in the average state.

For instance, when birds have to sing a very long time continuously for pairing to be possible, they should sing for long bouts on alternate days and on the intervening days rebuild their reserves (Fig. 4). The average singing routine cannot reveal this alternation: the days on which birds sing are not synchronized, so every day roughly half the population sings and half does not. Nor could one guess from looking casually at the policy, which is also the same every day.

An obvious way to make predictions about individual routines is by simulation. The state of an individual is followed forward, at each time interval using the optimal policy to determine the actions taken and random numbers to decide the effect of stochastic factors. Many such independent simulations produce a data set that we can analyse statistically to establish how behaviour or state at one time depends on earlier behaviour or state. Houston & McNamara (1999) explained an alternative approach, using dummy state variables instead of simulations, that

calculates the same sort of predictions ( $m$  acts as such a dummy state variable in Fig. 3b).

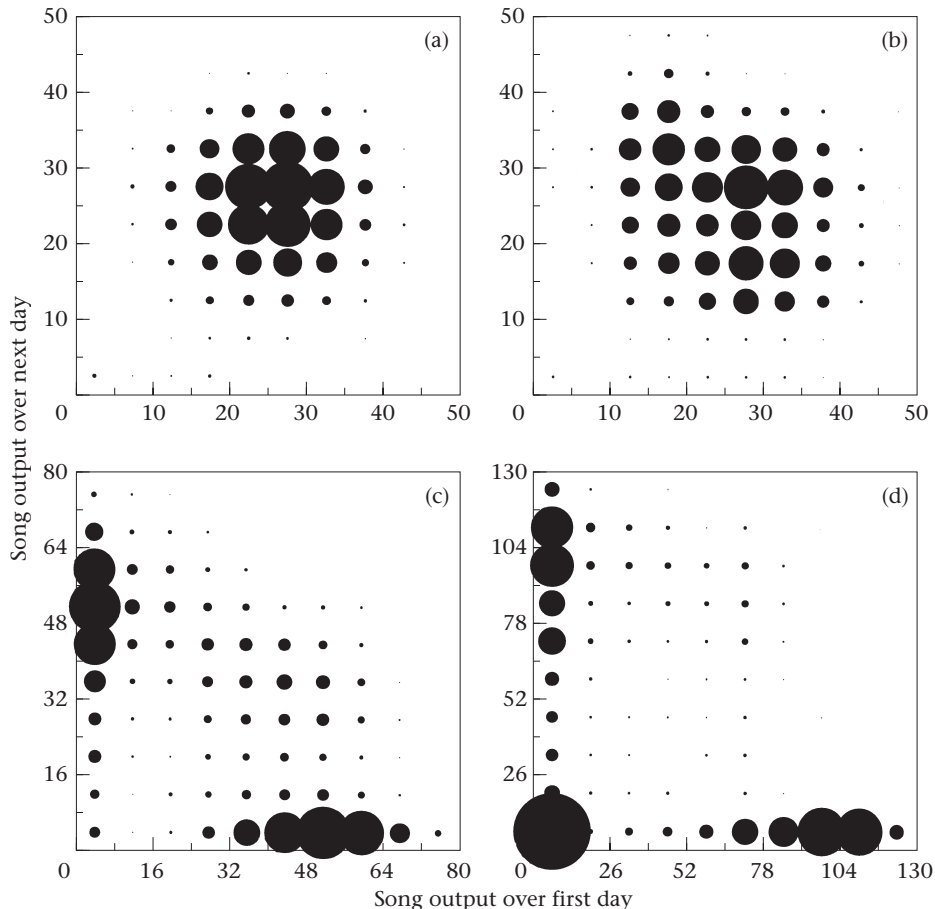
In the rest of this section we discuss correlations between state at one time and state at another, but the same arguments apply to correlations between behaviour at one time and behaviour at another. There are other possibilities, for instance to relate state at one time with subsequent behaviour (e.g. dusk reserves with the subsequent dawn chorus). In SDP models behaviour is a function of state, and if the same behaviour is optimal in different states, the function loses information. Consequently, correlations in behaviour may be weaker than correlations in state. To recover extra information from behaviour that is categorical one can average over successive time intervals (e.g. calculate the proportion of early-morning time intervals spent singing).

In many SDP models the correlation between state at one time and state at a later time rapidly disappears as the interval between observations increases. One reason for this lack of correlation is evident in the original version of the singing routine model. There is only one state variable (energy reserves) and the policy is to sing when reserves are above a switchpoint. So if a bird has high reserves, it sings and its reserves fall to the switchpoint; if it has low reserves, it forages until they rise to the switchpoint. Consequently all birds have similar reserve levels, and there is no correlation between reserve levels more than a few time steps apart. However, it is not always the case that state variables are independent from one time to another, and we now consider a series of factors favouring a correlation.

### *When does state at one time correlate with state at another?*

Correlations may persist if the state variable changes slowly. In birds recent evidence from increasing clutch sizes and removing second clutches suggests the existence of some state variable that decreases in response to the effort required to feed the young, and that still affects performance the next breeding season (Gustafsson & Pärt 1990; Verhulst 1998). Energy reserves change too rapidly, but immunocompetence or plumage condition are more plausible mechanisms (Gustafsson et al. 1994; Nilsson & Svensson 1996). Although the mechanism is uncertain, it has still been informative to develop models with a 'condition' state variable having slow dynamics (McNamara et al. 1998; Houston & McNamara 1999). For instance, Houston & McNamara (1999) considered a model of the timing of reproduction in which the state variables were energy reserves and condition. When food was supplemented well before the breeding season, the major effect was on the number of individuals that bred, whereas when the supplementation was just before the breeding season, the major effect was on the timing of breeding. This is a novel and testable prediction.

Some state variables change very slowly or not at all, for instance an individual's adult size, its quality, or the richness of its environment. Such factors may generate a positive correlation between state at different times if each pair of measurements is derived from a different individual. This may occur even if the state variable



**Figure 4.** Correlation in total song output over successive days (measured in time intervals spent singing; 150 time intervals=24 h), from the model of [Hutchinson et al. \(1993\)](#). The area of each dot is proportional to the number of occurrences in  $10^4$  independent simulations. For (a) the daily singing routine is similar to [Figure 1](#), the probability of pairing does not depend on past song output, and there is no correlation between song output on successive days. In (b)–(d), the probability of pairing depends on a sliding average of recent song output. As the importance of singing in long bouts increases, in (b) a slight negative correlation in song output on successive days is apparent, in (c) birds mostly alternate days of high song output with days of silence, and in (d) birds take 2 days to build up reserves after a day of very high song output.

measured is not the one that is unchanging. For instance, birds in territories with a high perceived predation risk may maintain comparatively low body weights throughout the day or season. Correlations generated in this way are hardly evidence for a particular SDP model, so the experimental design should rule out such explanations. Manipulations offer one solution; another is to measure the correlation on repeated observations of the same individual (replicating over several individuals).

Even if state variables change quickly, one can still get correlations. In some versions of the singing routine example, given their reserves, almost all population members do best to forage maximally throughout the afternoon. This consistency in behaviour preserves their relative levels of reserves. Another case is if the range of actions available cannot affect the rate of change of the state variable much. Energy reserves in a small bird will fall considerably overnight, but because the bird is unable to take action to alter the rate of fall by much (e.g. it is too dark to forage), we predict a strong correlation between reserves at dusk and at dawn.

A related argument explains why correlations can occur in models in which the effort spent foraging can be continuously adjusted, but the higher the level of foraging the higher the risk of predation. Then the optimal policy is often not to forage at the maximum intensity, and consequently reserve levels may converge only slowly ([Houston et al. 1993](#)).

Negative correlations can occur when an available action has a very large effect on the state variable, and the action must be all or nothing. Such an action can cause the value of the state variable to overshoot the switchpoint. A possible example may be those albatrosses that breed only in alternate years ([Prince et al. 1981](#)). Any breeding attempt may so reduce condition that birds do best to spend the year after a breeding attempt regaining condition ([McNamara & Houston 1996](#)).

Correlations are also liable to persist when the policy depends on two or more state variables. We saw in the singing routine model that with only one state variable different states would normally converge to the switchpoint; plotting the switchpoint against time gives a critical



line. With two state variables, state converges to a critical surface instead. Once in the neighbourhood of this surface there need be no further tendency for state to converge, so individuals starting in different states can continue to differ in state, and consequently their future behaviour can differ. For instance, consider the version of the singing routine example when a second state variable remembers how intensively a bird has been singing recently. If at a particular time two individuals do happen to have identical reserves, they need not be identical in respect of the memory state variable. In that case the bird with the higher memory variable can lie above the switchpoint surface and will continue to sing, whereas the other lies below and will forage. Consequently their reserves will diverge.

If there is a large stochastic influence on the change in state, correlations over time will obviously be smaller; but the effect of stochasticity is one of degree rather than kind.

### USING MANIPULATIONS: ADVANTAGES AND PITFALLS

There are several reasons why an experimental manipulation of state may be informative, some already mentioned. Another is that some state variables are easier to manipulate than to measure. For instance, rather than measuring food availability in different bird territories it is easier to give supplementary food in some of them (e.g. Thomas 1999a, and references therein). This also avoids food availability being confounded with male quality, which is likely if birds have fought for the best territories; after a manipulation we know what varies. Manipulations enable us to investigate conditions that rarely occur naturally; in SDP models rare dangerous events often dominate explanations of the optimal policy. Manipulations may also be useful in making some state variables and the environment more uniform. For instance, if we manipulate the stochasticity in the food regime, it makes sense to test all individuals before and after the manipulation on the same sort of nonstochastic regime, both to detect how policies have changed and to facilitate comparison between individuals (e.g. Ekman & Hake 1990).

An attraction of SDP models is that experiments chosen because of their tractability in the field can be mirrored by experimental manipulations on the computer, either by a sudden change of state or by a change in parameter values (e.g. Houston & McNamara 1989). Not only does this mirroring generate the most appropriate predictions of the manipulation's effects, but trying out the manipulation first on the computer can suggest which empirical manipulations are likely to have the most detectable effects, and on which aspects of behaviour and state. For instance, on the basis of results from our SDP models of the avian singing routines (McNamara et al. 1987; Hutchinson et al. 1993), Thomas (1999b) decided that his manipulation of food stochasticity might have its clearest effect on the relative strengths of the dusk and dawn choruses. This choice enabled birds to act as their own controls, thus factoring out variation in overall song effort between birds. Also, the prediction was not made by other hypotheses for the dawn chorus.

However, there are potential pitfalls with interpreting the effects of a manipulation, which can also apply when interpreting the consequences of natural environmental variation. The problems arise from not knowing whether an organism has interpreted an environmental change as short term, long term, or something intermediate. The optimal change of behaviour may be in opposite directions depending on the permanence of the environmental change. McNamara & Houston (1994) demonstrated this with an analytical model of animals optimally choosing between feeding options that differed in food reward and predation risk. A general increase in the food rewards would cause foraging to be more intensive if the manipulation were interpreted as a short-term change, but less intensive if interpreted as a permanent change.

Consider how in an SDP model we would calculate the effect of changing an environmental parameter, say foraging success. We might leave the policy unchanged and just recalculate the sequences of states and actions after the change. In contrast we might first recalculate the policy so that it is optimal in the new conditions. The latter is equivalent to including the parameter as a state variable of fixed value to which the organism can respond.

Modellers must carefully consider which of these approaches is appropriate to each particular empirical system and empiricists should check which approach was used to generate predictions. Predictions based on no change of policy are more likely to be appropriate to short-term manipulations when the subject has had little chance to learn about the change. Such predictions are also applicable if the manipulated feature naturally fluctuates rapidly and unpredictably, so that the manipulation is unlikely to be interpreted as permanent. Predictions based on a changed policy are most appropriate when the subject could have learnt over a long period to adjust to local or current conditions and to be confident that they will persist. A change of policy could also come about through evolution, so such predictions would be appropriate for manipulations lasting several generations and for nonexperimental comparisons of species and populations. With characters that do not usually change within the lifetime of an individual (e.g. the host on which a parasitoid is feeding), predictions based on a change of policy might be appropriate for a nonexperimental comparison between individuals, but it is uncertain whether an individual could adjust its policy if the character were artificially manipulated. For wandering animals the size and scale of the spatial variation affects how much the policy should adjust to conditions at the present site; it also matters whether temporal differences are consistent over adjacent sites (Weber et al. 1999).

Judging whether an organism will adjust its policies in the course of an experiment is not always straightforward, and a continuum of intermediate cases can be envisaged. Suppose that  $r$  measures the natural correlation of environmental conditions at one time with those at a later time. If  $r=0$ , current conditions predict nothing about future conditions, so the policy should not change if conditions are experimentally altered. If  $r=1$ , the change of conditions must be permanent, so the policy should change. In intermediate cases, the organism can

**Table 1.** The predicted effects of a cold night on the avian singing routine, depending on the assumptions underlying the policy

|     | Assumptions of policy                      |   | After several normal nights |                        |               | Then after a cold night    |                        |               | Then after a normal night  |                        |               |
|-----|--|---|-----------------------------|------------------------|---------------|----------------------------|------------------------|---------------|----------------------------|------------------------|---------------|
|     | Overall probability of cold (or hot) night | Probability next night cold, given cold tonight | Percentage singing at dawn  | When switchpoint rises | Dusk reserves | Percentage singing at dawn | When switchpoint rises | Dusk reserves | Percentage singing at dawn | When switchpoint rises | Dusk reserves |
| (1) | 0  | 0   | 59                          | 26                     | 368           | 7                          | 26                     | 368           | 59                         | 26                     | 368           |
| (2) | —  | 1   | 59                          | 26                     | 368           | 7                          | 22                     | 405           | 97                         | 26                     | 368           |
| (3) | 0.1  | 0.1   | 86                          | 26                     | 387           | 27                         | 26                     | 387           | 86                         | 26                     | 387           |
| (4) | 0.1  | 0.325   | 84                          | 26                     | 385           | 24                         | 24                     | 396           | 93                         | 26                     | 385           |
| (5) | 0.1  | 0.775   | 75                          | 26                     | 378           | 15                         | 23                     | 403           | 96                         | 26                     | 378           |

In all cases we envisage a sequence of nights of normal temperature, and study the effects of a single cold night. For each day we give the percentage of the population singing in the first time interval after dawn, the number of time intervals after dawn at which the switchpoint first rises (which corresponds to the end of the morning plateau in song output, see Fig. 1) and the level of dusk reserves. The percentage singing at dawn is affected by the level of reserves at the previous dusk and the overnight temperature, whereas the latter two measures depend only on the policy used that day. In line (1) the policy remains constant under the assumption that overnight temperature will always be normal in the future. In (2) the policy changes, under the assumption that all subsequent nights will be as cold as the last. In (3)–(5) the policies are adapted to low probabilities ( $P=0.1$ ) of a warm and a cold night. These policies differ in their assumptions of how likely temperature is to remain the same on successive nights; the more likely, the closer the routines are to (2). Parameters and other model details are as in Figure 1.

get only some indication of future conditions from current conditions. Its policy should change, but probably not by as much as when  $r=1$ .

As an example, consider a simplified version of the singing routine model in which now only one, rather than three, night-time temperature can occur (but there is still considerable variation in energy use overnight). Assume in the first instance that the bird is hardwired to assume this particular temperature. An experimental one-night decrease of temperature will affect the routine the next day (e.g. less song at dawn), but the policy will not change, and so reserves by dusk will be as before (Table 1, line 1). Consequently, if the next night's temperature is normal again, the next day's routine will revert to the original pattern.

A second possibility is that the bird is hardwired to assume that the temperature does not vary, but that it learns this temperature through experience. Let us suppose (unrealistically) that a single colder night would change its policy. The routine on the day after the manipulation will differ from that when the policy was unchanging (Table 1, compare lines 1 and 2). Furthermore, the new low-temperature policy leads to greater reserves at dusk, so that a reversion to normal temperatures the next night means a larger surplus of reserves remaining at dawn, which extends the dawn chorus.

A more realistic possibility is that the policy is adapted to an environment in which nights often vary in their temperature. For simplicity we assume that there are only three possibilities for overnight temperature, and that the animal is hardwired to know these temperatures and the frequency with which each occurs. We now additionally assume a correlation in temperature between successive nights (as in Fig. 2b). The optimal policy is then a rule that specifies actions on the basis not only of current reserves but also of the previous night's temperature. An experimental decrease in one night's temperature will, in a technical sense, not affect this policy, because now part

of the policy is to take account of the previous night's temperature, and this rule is fixed. But a decrease in the preceding night's temperature will cause a change in the switchpoint values of reserves at which singing replaces foraging (e.g. Table 1, lines 4 and 5). Thus there is a change in the rule relating behaviour to reserves, but no change in the metarule, about the updating of this rule according to previous temperature. How much the switchpoints change will depend on what this metarule assumes about the correlations and frequencies of overnight temperatures (Table 1, cf. lines 3–5).

Researchers have manipulated temperature and measured birds' reserves. Evening levels of reserves do increase in direct response to several days of cold (Wansink & Tinbergen 1994; Rogers 1995; Hake 1995, the latter using a natural change of temperature rather than a manipulation), but one study found no such response to a single cold night (Lilliendahl et al. 1996). There is a suggestion also that the response to single cold days increased as the birds learnt that the between-night variation in temperature had increased (Bednekoff et al. 1994). But in no system do we have a full understanding of what cues and rules organisms use to update their policies. So when one predicts the effect of manipulations we advocate making two sets of predictions corresponding to whether the policy has, or has not, adjusted (e.g. Weber et al. 1999). Another approach is to design the experiment so that adjustment of the policy is either very unlikely (by making the manipulations short and minor in extent), or very likely to be complete (by a long-term manipulation). To answer how long the policy takes to adjust, one might continue until the behaviour showed no significant trend over consecutive days.

Unfortunately these are not watertight approaches. For example, Thomas (1999a) supplemented the food of robins, *Erithacus rubecula*, for a one-off manipulation lasting 4–7 h during a single day, intending that this should just affect the state variable but not the policy.

However, song output increased significantly not only at dusk but also the following dawn. The easiest way for our SDP model to predict this effect on the dawn chorus is to suppose that the robins interpreted the food supplementation as indicative of the ready availability of food continuing the next day.

Conversely, even though the experimenter knows that a manipulation is long term, the experimental organism itself cannot be certain that conditions will not change back. Thus many vertebrates in captivity seek to sample inferior food sources, mark territories and keep a look out for predators, however long they have been kept isolated in a safe environment with an unvarying food supply (McNamara & Houston 1980; Houston 1987). Risk of predation is not an easy quantity for organisms to estimate, so it is not surprising that they are hardwired always to assume some risk. A similar problem is that organisms may be hardwired to respond to indirect environmental cues rather than to whatever has the direct effect on fitness (e.g. clear skies might predict cold nights better than does afternoon temperature). If the experimenter does not happen to include these unknown cues in the manipulation, the response would not be optimal in the laboratory, even though it could be appropriate to natural variation in the wild.

## DISCUSSION

When statistically analysing empirical data, biologists know that one should attempt to factor out as much variation as possible so as to minimize the unexplained noise or error. It would be considered a weakness in a statistical analysis to average over individuals and conditions if the information was available. Many of our suggestions about what to predict and test can also be viewed as avoiding unnecessary averaging. For instance, we have advocated measuring the multivariate distribution of state, and keeping track of individuals rather than calculating an average routine.

We hope that we have shown that SDP models can make a wide variety of predictions, but some readers may now feel an overabundance of choice. We have advocated testing several predictions from each model, but testing too many predictions may be wasteful because if some are true others will necessarily also hold. In practice we expect that choices will be constrained by which behaviours, states and parameters can be most easily and accurately measured without disturbing the system, and by what can be manipulated. In some systems state is more difficult to measure than behaviour (e.g. if state includes the animal's perception of the environment), whereas in others behaviour is the harder (e.g. if we were modelling the allocation of resources between reproductive and storage organs).

When there is a choice, the SDP model itself may be used to predict the empirical approaches most likely to yield significant results. For instance we can estimate the best time to perform a particular manipulation, or the sample size sufficient to achieve significance given the stochasticity in the system. The model may incorporate some sources of stochasticity that cause differences

between individuals, and others that cause differences between days, in which case the model can suggest the most efficient way to replicate observations over individuals and days. A sensitivity analysis should check which predictions are robust to minor variations in the model's assumptions, and we might also model rival or null hypotheses to investigate which approach will best distinguish their predictions from those of the main model.

## Acknowledgments

We thank Innes Cuthill, Rob Thomas, Simon Verhulst, Thomas Weber and a referee for their comments on the manuscript. J.M.C.H. was supported by a grant from NERC to J.M.M. and I. C. Cuthill.

## References

- Bednekoff, P. A. 1996. Translating mass dependent flight performance into predation risk: an extension of Metcalfe & Ure. *Proceedings of the Royal Society of London, Series B*, **263**, 887–889.
- Bednekoff, P. A., Biebach, H. & Krebs, J. 1994. Great tit fat reserves under unpredictable temperatures. *Journal of Avian Biology*, **25**, 156–160.
- Brunner, D., Kacelnik, A. & Gibbon, J. 1992. Optimal foraging and timing processes in the starling, *Sturnus vulgaris*: effect of inter-capture interval. *Animal Behaviour*, **44**, 597–613.
- Bull, C. D., Metcalfe, N. B. & Mangel, M. 1996. Seasonal matching of foraging to anticipated energy requirements in anorexic juvenile salmon. *Proceedings of the Royal Society of London, Series B*, **263**, 13–18.
- Clark, C. W. & Ydenberg, R. C. 1990. The risks of parenthood II. Parent-offspring conflict. *Evolutionary Ecology*, **4**, 312–325.
- Clutton-Brock, T. H., Stevenson, I. R., Marrow, P., MacColl, A. D., Houston, A. I. & McNamara, J. M. 1996. Population fluctuations, reproductive costs and life-history tactics in female Soay sheep. *Journal of Animal Ecology*, **65**, 675–689.
- Ekman, J. B. & Hake, M. K. 1990. Monitoring starvation risk: adjustments of body reserves in greenfinches (*Carduelis chloris* L.) during periods of unpredictable foraging success. *Behavioral Ecology*, **1**, 62–67.
- Enquist, M. & Leimar, O. 1987. Evolution of fighting behaviour: the effect of variation in resource value. *Journal of Theoretical Biology*, **127**, 187–205.
- Galvani, A. & Johnstone, R. 1998. Sperm allocation in an uncertain world. *Behavioral Ecology and Sociobiology*, **44**, 161–168.
- Gibbon, J. & Allan, L. G. (Eds) 1984. *Timing and Time Perception*. New York: New York Academy of Sciences.
- Gladstein, D. S., Carlin, N. F. & Austad, S. N. 1991. The need for sensitivity analyses of dynamic optimization models. *Oikos*, **60**, 121–126.
- Gustafsson, L. & Pärt, T. 1990. Acceleration of senescence in the collared flycatcher *Ficedula albicollis* by reproductive costs. *Nature*, **347**, 279–281.
- Gustafsson, L., Nordling, D., Andersson, M. S., Sheldon, B. C. & Qvarnström, A. 1994. Infectious diseases, reproductive effort and the cost of reproduction in birds. *Philosophical Transactions of the Royal Society of London B*, **346**, 323–331.
- Hake, M. 1995. Adaptive control of starvation risk in greenfinches (*Carduelis chloris*). Ph.D. thesis, University of Göteborg.
- Heimpel, G. E., Mangel, M. & Rosenheim, J. A. 1998. Effects of time limitation and egg limitation on lifetime reproductive success of a parasitoid in the field. *American Naturalist*, **152**, 273–289.



- Houston, A. I. 1987. The control of foraging decisions. In: *Quantitative Analyses of Behavior. Vol. VI. Foraging* (Ed. by M. L. Commons, A. Kacelnik & S. J. Shettleworth), pp. 41–61. Hillsdale, New Jersey: L. Erlbaum.
- Houston, A. I. & McNamara, J. M. 1987. Singing to attract a mate: a stochastic dynamic game. *Journal of Theoretical Biology*, **129**, 57–68.
- Houston, A. I. & McNamara, J. M. 1988. Fighting for food: a dynamic version of the Hawk-Dove game. *Evolutionary Ecology*, **2**, 51–64.
- Houston, A. I. & McNamara, J. M. 1989. The value of food: effects of open and closed economies. *Animal Behaviour*, **37**, 546–562.
- Houston, A. I. & McNamara, J. M. 1999. *Models of Adaptive Behaviour: an Approach Based on State*. Cambridge: Cambridge University Press.
- Houston, A. I., McNamara, J. M. & Thompson, W. A. 1992. On the need for a sensitive analysis of optimization models, or, 'This simulation is not as the former'. *Oikos*, **63**, 513–517.
- Houston, A. I., McNamara, J. M. & Hutchinson, J. M. C. 1993. General results concerning the trade-off between gaining energy and avoiding predation. *Philosophical Transactions of the Royal Society of London, Series B*, **341**, 375–397.
- Hurly, T. A. 1992. Energetic reserves of marsh tits (*Parus palustris*): food and fat storage in response to variable food supply. *Behavioral Ecology*, **3**, 181–188.
- Hutchinson, J. M. C. 1999. Bet-hedging when targets may disappear: optimal mate-seeking or prey-catching trajectories and the stability of leks and herds. *Journal of Theoretical Biology*, **196**, 33–49.
- Hutchinson, J. M. C., McNamara, J. M. & Cuthill, I. C. 1993. Song, sexual selection, starvation and strategic handicaps. *Animal Behaviour*, **45**, 1153–1177.
- Hutchinson, J. M. C., McNamara, J. M., Houston, A. I. & Vollrath, F. 1997. Dyar's Rule and the Investment Principle: optimal moulting strategies if feeding rate is size-dependent and growth is discontinuous. *Philosophical Transactions of the Royal Society of London, Series B*, **352**, 113–138.
- Iwasa, Y. 1991. Sex change evolution and cost of reproduction. *Behavioral Ecology*, **2**, 56–68.
- Lilliendahl, K., Carlson, A., Welander, J. & Ekman, J. B. 1996. Behavioural control of daily fattening in great tits (*Parus major*). *Canadian Journal of Zoology*, **74**, 1612–1616.
- Lucas, J. R. & Howard, R. D. 1995. On alternative reproductive tactics in anurans: dynamic games with density and frequency dependence. *American Naturalist*, **146**, 365–397.
- Lucas, J. R. & Walter, L. R. 1991. When should chickadees hoard food? Theory and experimental results. *Animal Behaviour*, **41**, 579–601.
- Mace, R. 1993. Nomadic pastoralists adopt subsistence strategies that maximise long-term household survival. *Behavioral Ecology and Sociobiology*, **33**, 329–334.
- McNamara, J. M. 1990. The starvation-predation trade-off and some behavioural and ecological consequences. In: *Behavioural Mechanisms of Food Selection* (Ed. by R. N. Hughes), pp. 39–59. New York: Springer-Verlag.
- McNamara, J. & Houston, A. 1980. The application of statistical decision theory to animal behaviour. *Journal of Theoretical Biology*, **85**, 673–690.
- McNamara, J. M. & Houston, A. I. 1986. The common currency for behavioral decisions. *American Naturalist*, **127**, 358–378.
- McNamara, J. M. & Houston, A. I. 1987. Starvation and predation as factors limiting population size. *Ecology*, **68**, 1515–1519.
- McNamara, J. M. & Houston, A. I. 1992. Risk-sensitive foraging: a review of the theory. *Bulletin of Mathematical Biology*, **54**, 355–378.
- McNamara, J. M. & Houston, A. I. 1994. The effect of a change in foraging options on intake rate and predation rate. *American Naturalist*, **144**, 978–1000.
- McNamara, J. M. & Houston, A. I. 1996. State-dependent life histories. *Nature*, **380**, 215–221.
- McNamara, J. M., Mace, R. H. & Houston, A. I. 1987. Optimal daily routines of singing and foraging in a bird singing to attract a mate. *Behavioral Ecology and Sociobiology*, **20**, 399–405.
- McNamara, J. M., Welham, R. K. & Houston, A. I. 1998. The timing of migration within the context of an annual routine. *Journal of Avian Biology*, **29**, 416–423.
- Mangel, M. & Clark, C. W. 1986. Towards a unified foraging theory. *Ecology*, **67**, 1127–1138.
- Mangel, M. & Clark, C. W. 1988. *Dynamic Modeling in Behavioral Ecology*. Princeton, New Jersey: Princeton University Press.
- Monaghan, P., Nager, R. G. & Houston, D. C. 1998. The price of eggs: increased investment in egg production reduces the offspring rearing capacity of parents. *Proceedings of the Royal Society of London, Series B*, **265**, 1731–1735.
- Newman, J. A., Parsons, A. J., Thornley, J. H. M., Penning, P. D. & Krebs, J. R. 1995. Optimal diet selection by a generalist grazing herbivore. *Functional Ecology*, **9**, 255–268.
- Nilsson, J.-Å. & Svensson, E. 1996. The cost of reproduction: a new link between current reproductive effort and future reproductive success. *Proceedings of the Royal Society of London, Series B*, **263**, 711–714.
- Parker, G. A. & Maynard Smith, J. 1990. Optimality theory in evolutionary biology. *Nature*, **348**, 27–33.
- Prince, P. A., Ricketts, C. & Thomas, G. 1981. Weight loss in incubating albatrosses and its implications for their energy and food requirements. *Condor*, **83**, 238–242.
- Reinhardt, U. G. & Healey, M. C. 1999. Season- and size-dependent risk taking in juvenile coho salmon: experimental evaluation of asset protection. *Animal Behaviour*, **57**, 923–933.
- Rogers, C. M. 1995. Experimental evidence for temperature-dependent winter lipid storage in the dark-eyed junco (*Junco hyemalis oreganus*) and song sparrow (*Melospiza melodia morphna*). *Physiological Zoology*, **68**, 277–289.
- Rowe, L. & Ludwig, D. 1991. Size and timing of metamorphosis in complex life cycles: time constraints and variation. *Ecology*, **72**, 413–427.
- Stearns, S. C. & Schmid-Hempel, P. 1987. Evolutionary insights should not be wasted. *Oikos*, **49**, 118–125.
- Thomas, R. J. 1999a. Two tests of a stochastic dynamic programming model of daily singing routines in birds. *Animal Behaviour*, **57**, 277–284.
- Thomas, R. J. 1999b. The effect of variability in the food supply on the daily singing routines of European robins: a test of a stochastic dynamic programming model. *Animal Behaviour*, **57**, 365–369.
- Verhulst, S. 1998. Multiple breeding in the great tit, II. The costs of rearing a second clutch. *Functional Ecology*, **12**, 132–140.
- Wansink, D. & Tinbergen, J. M. 1994. The influence of ambient temperature on diet in the great tit. *Journal of Avian Biology*, **25**, 261–267.
- Weber, T. P., Ens, B. J. & Houston, A. I. 1998. Optimal avian migration: a dynamic model of fuel stores and site use. *Evolutionary Ecology*, **12**, 377–401.
- Weber, T. P., Fransson, T. & Houston, A. I. 1999. Should I stay or should I go? Testing optimality models of stopover decisions in migrating birds. *Behavioral Ecology and Sociobiology*, **46**, 280–286.
- Witter, M. S., Cuthill, I. C. & Bonser, R. H. 1994. Experimental investigations of mass-dependent predation risk in the European starling, *Sturnus vulgaris*. *Animal Behaviour*, **48**, 201–222.
- Witter, M. S., Swaddle, J. P. & Cuthill, I. C. 1995. Periodic food availability and strategic regulation of body mass in the European starling, *Sturnus vulgaris*. *Functional Ecology*, **9**, 568–574.