



A state-free optimization model for sequences of behaviour

ALAN GRAFEN

Department of Zoology, University of Oxford

(Received 23 February 2001; initial acceptance 11 May 2001;
final acceptance 11 September 2001; MS. number: 6858)

The view of animals as optimizers is extremely important in many parts of biology, and is represented in a variety of theoretical models. Current developments in animal welfare create a need for a new theoretical model, which treats the animal as acting optimally in its choice of a sequence of behaviours, and places as few restrictions as possible on what quantity is being maximized. The intention is to facilitate the study of what the animal chooses, which will often be relevant in considering animal welfare. Currently, economic models of the consumer are borrowed in this situation, but they are not well suited for the increasingly exact analyses being performed. A new alternative model is presented here analytically in general, and a special case is studied numerically. The model does not attribute internal states to the individual, which is a simplification with advantages and disadvantages.

© 2002 The Association for the Study of Animal Behaviour

The use of economic analogies has been a major technique in animal welfare since it was introduced 20 years ago by Dawkins (1981), and recently the applications have become more and more quantitative (for example Cooper & Mason 2000; Mason et al. 2001). The central ideas are that individuals reveal their preferences by observed choices, including experimentally presented choices, and that individuals to at least some extent 'know what is good for them'. Thus, although there is no question of accepting automatically that whatever an animal chooses is best for it in welfare terms, it will often be useful to study what an animal chooses. The conceptual framework that has been used in this work has been the economic model of the consumer, as outlined in standard economics textbooks such as Mas-Colell et al. (1995). In applied work, this provides a family of techniques for investigating animal welfare. In more purely scientific terms it provides the challenge of understanding sequences of behaviour as the result of an optimization process.

The need for a new model arises first because the model of the economic consumer represents rather poorly the situation of an animal in animal welfare experiments, and this point is explained more fully below. The second reason is that existing biological models of sequences of behaviour (McFarland & Houston 1981; Houston & McNamara 1999) are constructed to allow an adaptationist analysis of behaviour. The mathematical structures are consistent with a Darwinian interpretation of the maximand, and they enforce a kind of Darwinian

rationality on the animal. All models represent a compromise between simplicity and complexity, and these models usually assume what the animal is meant to be maximizing, and are designed to enquire whether we have understood how it is that the animal's behaviour succeeds in maximizing the unquestioned maximand.

In contrast, the economic model of the consumer leaves the utility function (almost completely) unspecified, with the aim of describing any reasonable behaviour within the optimization framework. This is the spirit of the animal welfare work, and this is why a model designed for that area should aim to allow any 'reasonable' behaviour on the part of the animal. We wish to be able to infer from an experiment that play, say, is extremely important to an animal, without having to have any idea about why the preference for play is adaptive in a Darwinian sense.

In this paper, we aim to present as simple a model as possible that allows as wide a range of reasonable behaviour as possible in the experimental arrangement used by Cooper & Mason (2000), and allows us to pose questions of substitutability, elasticities and demands. One use for this model is to explore whether the inexactness of the economic model is important in its quantitative applications so far. It could also in principle be used to develop a parallel theory to the economic model of the consumer, providing the same facilities but with basic assumptions that are more closely met by the situation of the experimental animals. Another use is to apply the logic of the optimization approach to features of the behaviour that are elided in the economic analogy, particularly those to do with sequencing.

Correspondence: A. Grafen, St John's College, Oxford OX1 3JP, U.K.
(email: alan.grafen@sjc.ox.ac.uk).

The interpretation of sequences of behaviour as the outcome of an optimization process does arise in fields other than animal welfare, and it may be that the model developed here has correspondingly wider applications. For example, real human consumers also make choices of sequences of behaviour.

THE ECONOMIC MODEL OF THE CONSUMER

My aim in this section is to explain which aspects of the basis of the economic model of the consumer are likely not to be met in applications in animal welfare.

It will suffice to consider the model with two commodities, quantities of which will be represented by x and y . The utility of the consumer is represented by the function $u(x, y)$. The prices of the two goods will be denoted p and q , and the consumer's budget will be b . The theory states that the consumer chooses x and y , subject to the constraints $x \geq 0$, $y \geq 0$ and $px + qy \leq b$, to maximize the value of $u(x, y)$. The hope is that any reasonable behaviour on the part of the consumer can be represented by a suitable choice of the function u . (Quite a lot of effort in economics has been spent considering much more precisely what 'reasonable' might and should mean, and in what conditions the hope is justified. These refinements are not needed here.)

In one application, Cooper & Mason (2000) studied mink, *Mustela vison*, in a closed economy arena based on commercial practice but with some added activity areas reached by wire corridors. They varied the cost of an activity by changing the weight on the door through which access was gained. Cooper & Mason (2000) discussed thoroughly the fact that this is not a cost per unit of activity, but resembles more closely an entry fee.

We now explore what would have to be true for the simple consumer model to apply, with the analogy of door weighting to price. One possibility is to identify x and y as the time spent in the activities. One constraint on x and y is the time constraint, that there is only a finite amount of time in the day. For this constraint, the prices must be equal, as a minute on one activity reduces the time budget by the same amount as a minute on the other. But then there is no role in the model for door weighting as a price. We could instead ignore the time constraint, and analogize p and q to the door weightings. This constraint on x and y does not work well, because we multiply the cost per visit by the total duration of all visits, making a nonsensical product. This point led Cooper & Mason (2000) to a second way of constructing the analogy.

This other possibility is to identify x and y as the numbers of visits to the locations for the two activities. This is an improvement so far as the constraint is concerned, as the door-weight per visit is multiplied by the number of visits. We would have to make strong assumptions of additivity of door-pushes: why not take the logarithm or the square root of the door-weight as the price? Let us note that assumption and carry on. There is a problem now with the utility function, which has to be a function of the number of visits, and not of the duration, as the variables must have the same

interpretation in the utility function as in the constraint. If visit duration varies a lot during the experiment as the door weightings change, then we would expect this problem to present a serious concern. There remain two difficulties with the constraint in this interpretation. First, the time constraint is again missing from the model. Second, is the budget really constrained? It is reasonable to expect the mink to indulge less frequently in activities with a heavier door weighting, but if the constraint is well analogized to a budget, then the mink should act as though constrained by the sum of the door weightings they endure. In terms of the economic notation, we must in that case expect $px + qy$ to remain constant as p is varied. It would be useful if reports on experiments using the analogy considered this point, and provided the relevant information. There seems no reason in general to expect $px + qy$ to remain constant as the door weightings are changed.

How much do these potential failures in the analogy matter? Elasticities can be calculated as an indication of sensitivity of demand to environmental changes, and may serve this purpose well even though further technical properties of elasticities that hold within the consumer model (such as reconstructing utilities from them) would not hold. For other analyses, however, the constancy of $px + qy$ does matter. The calculation of 'consumer surplus' as reported by Mason et al. (2001) is intended as a measure of the difference to utility made by a change in price; but the analogy on which this calculation depends is brought into doubt if $px + qy$ is not constant.

The reason for the relevance of variation in $px + qy$ to consumer surplus involves a small excursion into economic theory. Mas-Colell et al. (1995) derived and discussed the calculation of consumer surplus in their Section 3.I and illustrated it in their Figure 3.I.3, of which a version is reproduced here as Fig. 1. They distinguished two kinds of demand curves. Uncompensated (or Walrasian or market) demand curves show how the quantity consumed (or 'demand') for one commodity changes as its price changes, assuming there is no simultaneous change in other prices or in the consumer's budget. Compensated (or Hicksian) demand curves show how demand for one commodity changes as its price changes, and no other prices change, but the budget is changed in parallel with the price, and by exactly the amount necessary to keep the consumer at the same level of utility. Mathematical arguments prove that the area to the left of the compensated demand curve between two prices can be interpreted as the difference that would be made to the utility of the consumer by that change in price if it were not compensated. The hedging phrase 'can be interpreted as' is necessary, because the area does not correspond to the difference in the consumer's actual utility function (which is usually considered unobservable in economics), but only to a money-metric equivalent to that utility function. The area to the left of the uncompensated demand curve is sometimes used by economists as a surrogate for the area to the left of the compensated demand curve, which is much harder to observe. This use is an approximation that assumes the

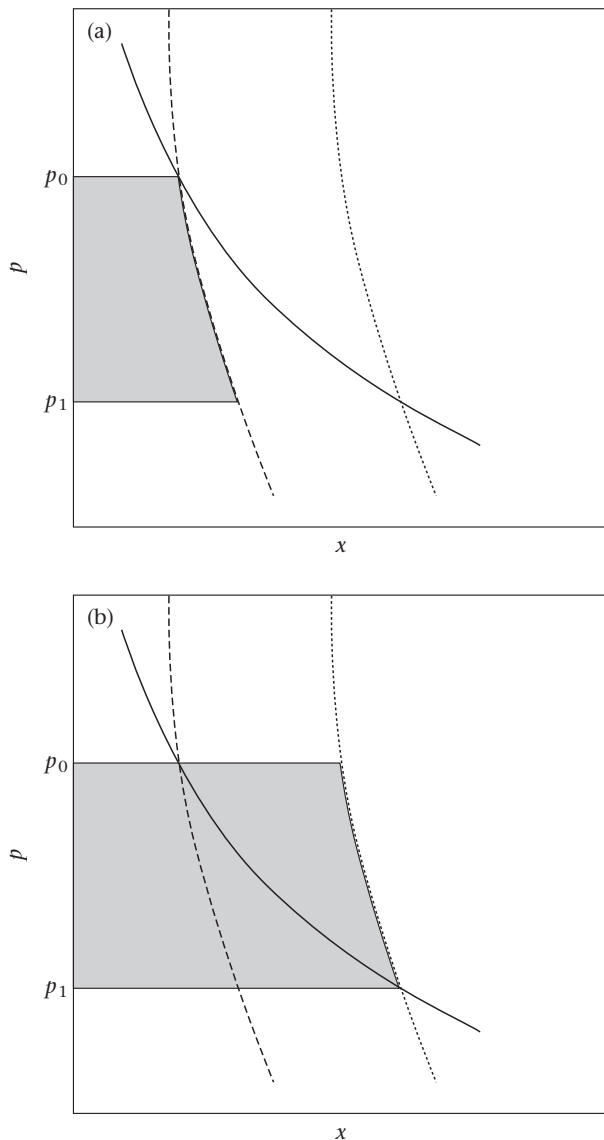


Figure 1. The figure shows the effect of reducing the price of one good (p) from p_0 to p_1 on the quantity of that good purchased by an individual (x). The dashed and dotted lines are the compensated demand curve when the price is p_0 and p_1 , respectively. The solid line is the uncompensated demand curve. (a) The shaded area is the equivalent variation (how much money would need to be given to the individual when the price is p_0 to obtain the same level of utility as would be brought about by reducing the price to p_1). (b) The shaded area is the compensating variation (how much money would need to be given to the individual when the price had been reduced to p_1 to restore the same level of utility that had been attained when the price was p_0). Only if the compensated curves are close together is the area to the left of the uncompensated demand curve a measure of the change in utility.

wealth effects of the change in price are small enough to ignore. However, it seems quite likely that when the analogy is made between price and door weightings, $px+qy$ varies as p varies. Variation in $px+qy$ between different points on a demand curve creates additional wealth effects. If $px+qy$ varies little, then the demand curve will approximate an uncompensated demand

curve; if it varies so as to maintain utility, then it will approximate a compensated demand curve. Variation in $px+qy$ could cause the slope of the observed demand curve to be much greater or much less than the slopes of the compensated and uncompensated demand curves. This could lead to a substantial under- or overestimate of the utility change.

Even if one study found that $px+qy$ remained roughly constant, this would still leave unsettled whether the same would hold good in future studies. The economic model is clearly designed for a situation in which the structure of the problem forces $px+qy$ to remain constant. In economics, you cannot spend money you do not have. A robust model for the welfare case would not contain this biologically unnatural requirement.

Thus, the importance of how tightly the analogy with economics can be drawn increases as more sophisticated parts of the economic methodology are borrowed for application in biology. The motivation for the model, presented below, is that, while the application of such sophisticated techniques is an important advance, they will require modification and rejustification in a model constructed more specifically for the biological situation. The drive towards experimental precision and conceptual clarity requires, and deserves, the support of a custom theoretical apparatus.

THE MODEL

The main discrepancy between the economic consumer and the captive animal is in the dynamic nature of the choices facing the animal. A sequence of behaviours is produced, each with a duration. Furthermore, the timing of the behaviours is important. Eating for half an hour each day is not the same as eating for 2 years and then fasting for 46 years. Furthermore, some manipulations affect each bout of a behaviour in the same way, no matter what its duration. Half an hour spent eating might incur the penalty of opening one heavy door if undertaken in a single session, but the penalty of opening a heavy door three times if undertaken in 10-min bouts.

The link between Darwinism and optimization by animals can operate at many different levels. One currently conventional view is that animals are designed by natural selection to be efficient optimizers in the environment in which their ancestors lived for many generations, but that, while their behaviour in artificial environments can also be interpreted in terms of optimization, the quantity optimized may show little resemblance to Darwinian fitness. The discrepancy is very acute in many applications of animal welfare, in which reproduction is strictly controlled or prevented by the farmer. A further complication is that we do not wish to give an artificial priority to activities whose contribution to fitness is currently understood (e.g. feeding) over activities that are less well understood (e.g. swimming, grooming, playing). The approach adopted in the rest of the paper will therefore be the same as that adopted by economists studying humans. We assume that the animal acts according to an optimization principle, but we aim to leave open what the maximized quantity is. In principle,

we would discover the nature of the maximand from observing behaviour. We are studying what the animal chooses to do, and placing as few constraints as possible on what those choices might be. This method has been discussed and described in a biological context as the 'inverse optimality approach' (McFarland & Houston 1981).

A notable feature of the model presented here is that it does not use the concept of the state of an animal, and this omission requires justification. Other models of sequences of behaviour such as those of McFarland & Houston (1981) and Houston & McNamara (1999) do use states. In the case of hunger, the state can be clearly defined (at least in retrospect over short periods) as the size of the deficit, roughly measurable as how much the animal would eat now if offered food ad libitum. It is natural to link the strength of hunger and thirst to disutility. But for play and sleep, it is much harder to define the state. Trade-offs between different kinds of activities are exactly at issue in welfare debates, and this model treats the activities as equal a priori. The omission of states means that activities with easily defined states are not privileged over activities without, and this is one advantage. This omission can be justified on two further grounds. A practical issue is that information about states, even where those states are in principle definable, will often not be available in an experiment; or would be expensive to collect. Furthermore, a major purpose of the model is to represent more formally ideas about what different experiments on sequences of behaviour tell us about the animal. When such ideas do not involve internal states, there is a role for a theory that simplifies by abstracting from them.

The model can be outlined as follows. The animal chooses a sequence of behaviours that is considered to be repeated indefinitely. There are n activities, which will be indexed by a . Each activity a is considered to give rise to a 'benefit' that depends on the periods for which it is and is not performed. The benefit of performing activity a for duration $t > 0$ will be denoted $f_a(t)$, and the benefit of not performing activity a for duration $t > 0$ will be denoted $g_a(t)$. The benefit is the sum of the f s for all the bouts of the activity and of the g s for all the gaps between bouts. An activity like feeding would be considered to produce a positive benefit while being performed, but a negative benefit (through use of the food, emptying of the stomach) while not being performed. Each activity produces a 'benefit rate' which is the rate through time at which benefit is accrued. Finally, the utility rate of the animal is a function of all the benefit rates. The animal is presumed to choose its behaviour to maximize the utility rate.

More formally, a strategy is a finite sequence of activities $S = (S_1, S_2, S_3 \dots S_N)$ and a corresponding sequence of durations $T = (T_1, T_2, T_3 \dots T_N)$. It will sometimes be helpful to avoid certain kinds of degenerate sequences, in two ways. First, no two succeeding activities can be the same. That is, $S_i \neq S_{i+1}$ and, because the sequence is thought of as repeating indefinitely, we also require $S_1 \neq S_N$. Second, none of the durations can be zero, i.e. $T_i > 0$. In any given application, some further criteria may be relevant in

limiting the set of possible sequences we permit. For example, if an animal is placed in an experimental apparatus that enforces certain orders on activities, such as the need to open a door before feeding, then sequences that broke this requirement would simply be excluded from the set of permitted strategies.

It is helpful to introduce for each bout the duration afterwards for which that activity is not performed, say U_i . For the last bout of a given activity, the duration is, in accordance with the idea of indefinite repetition of the whole sequence, 'wrapped round' to the beginning of the first bout. A strategy has nonunique representations, as it can be chosen to start on any one of its elements. Furthermore, a strategy with N bouts can be repeated k times to form a strategy with kN bouts, and these are also alternative representations of the same strategy.

Let $supp(S)$ denote the set of activities contained in S . Before defining the utility rate, we need to consider activities that are never performed. Although it may be difficult in practice to identify this potentially infinite set, it will be important to include in one experimental regime activities that are performed in another. In a period of time t , such an activity produces a benefit of $g'_a(t)$. As t grows larger, the rate of accrual of benefit is the limit of $g'_a(t)$, which we shall assume exists (but may be minus infinity) and notate as $g'_a(\infty)$.

The benefit rate over the sequence for activity a can be written as

$$b_a(S, T) = \begin{cases} \frac{\sum_{i:a=S_i} (f_a(T_i) + g_a(U_i))}{\sum_j T_j} & a \in supp(S) \\ g'_a(\infty) & a \notin supp(S) \end{cases}$$

and the utility rate can be written as

$$r(S, T) = u(b_1(S, T), b_2(S, T) \dots b_n(S, T))$$

We can write a formal optimization program as

$$(S, T) \max r(S, T), \\ (S, T) \in C$$

where C is a constraint set limiting the permissible values of (S, T) . In words, we consider the choice of (S, T) to maximize $r(S, T)$ over the set C of permitted sequences.

The function u plays the same role as the utility function in economic theory, and can represent substitutability between benefits (and so indirectly of their activities). The functions b_a , through their dependence on the f_a and g_a , capture the effects of timing of activities.

Theoretical Results

This section, for theoretical readers, develops a notation that allows the model to be manipulated analytically. For a given strategy (S, T) , first-order conditions based on optimality of T can be derived. Let $II(k)$ denote the set of i that U_i is increased if T_k is increased: that is, the set of gaps that include the k -th bout. Let U_k^-

denote the gap for S_k preceding bout k . We also introduce the notation $u_{(a)}$ to mean the partial derivative of u with respect to b_a , and so $u_{(S_k)}$ is the partial derivative of u with respect to the benefit rate of whichever activity is undertaken in the k -th bout of the sequence. In taking these partial derivatives, the other T_i are held constant, and so the total duration of the cycle of activities changes in line with the changes in T_k . Then differentiating the expression for u with respect to T_k yields

$$\left(\sum_i T_i\right) \frac{\partial u}{\partial T_k} = u_{(S_k)} f'_{S_k}(T_k) + \sum_{i \in \Pi(k)} u_{(S_i)} g'_{S_i}(U_i) - \sum_{a \in \text{supp}(S)} b_a u_{(a)}$$

These all equal zero at an internal equilibrium, from which it follows that

$$\sum_{a \in \text{supp}(S)} b_a u_{(a)} = u_{(S_k)} f'_{S_k}(T_k) + \sum_{i \in \Pi(k)} u_{(S_i)} g'_{S_i}(U_i)$$

independently of k . This provides a way to search for an optimum. Values of k for which the RHS is less than (more than) the LHS should have T_k increased (decreased).

By subtracting the values for k and $k+1$ we obtain

$$(f'_{S_k}(T_k) - g'_{S_k}(U_k))u_{(S_k)} = (f'_{S_{k+1}}(T_{k+1}) - g'_{S_{k+1}}(U_{k+1}))u_{(S_{k+1})}$$

This final set of first-order conditions has lost completeness, as it does not reflect on the possibility of changing the total duration of the cycle. However, it has an appealing intuitive interpretation in terms of altering the durations of two neighbouring activities without altering the sum of those durations.

A Simple Example

We construct here a simple example with benefit functions and a utility function, with the aim of displaying a few features of the model in a simplified theoretical version of the experiment of [Cooper & Mason \(2000\)](#). The more realistic situation of inferring the nature of these functions from data is beyond the scope of the present paper. The example has three activities, feeding, drinking and resting, for which we use the subscripts 1, 2 and 3. Once the behaviour of this very simple case has been explored a little in this section, the next will look at adding doors between the compartments. The doors will be allowed to vary in the force required to open them, and the mink will be assumed not to like opening weighty doors.

Feeding

We must define f, g and a fragment of a utility function relating to feeding. After feeding for t_1^{\min} , an animal is assumed to take in food at a rate with diminishing returns and ultimately with no further intake. The amount of food taken in for $t \geq t_1^{\min}$ will be assumed to be

$$K_1(1 - \exp(-\lambda_1(t - t_1^{\min})))$$

where K_1 is the maximum amount that can be taken in one bout, and λ_1 represents how fast the asymptote is reached. If we assume that the need for food grows at a constant rate μ_1 while feeding or not, then we can define

$$f_1(t) = \begin{cases} -\mu_1 t & t \leq t_1^{\min} \\ K_1(1 - \exp(-\lambda_1(t - t_1^{\min}))) - \mu_1 t & t \geq t_1^{\min} \end{cases}$$

$$g_1(t) = -\mu_1 t$$

The utility fragment must also be chosen. We will shortly define the overall utility as the product of the separate fragments, so that if the feeding fragment is zero, then the whole utility is also zero. To have a continuing deficit of food is lethal, and so will be allocated a utility of zero. To have too large a surplus is not advantageous. It would be possible to insist on an exact balance, but it is analytically more convenient, and probably more realistic, to assume that animals like to run a small surplus. This will allow them not to be continually monitoring food levels, and the potential energy surplus can be burnt off by slight metabolic inefficiency, or responsive digestive inefficiency could ensure it was never acquired. Recall that b_1 is the net rate at which food value is accrued, adding up over the cycle the f_1 value for each feeding bout and the g_1 value for nonfeeding periods, and dividing by the cycle length. The following function makes the optimal value of b_1 equal to a parameter ω_1 . Negative values of b_1 are effectively outlawed by assigning a utility of zero to them, in conjunction with a convention that utilities are always non-negative. The parameter κ_1 controls how strong the preference is that b_1 should be close to ω_1 . A high κ_1 means that even a small deviation from ω_1 results in a large reduction in utility. The function is

$$u_1(b_1) = \begin{cases} 0 & b_1 \leq 0 \\ \left(\frac{b_1}{\omega_1}\right)^{\omega_1 \kappa_1} \exp(-\kappa_1(b_1 - \omega_1)) & b_1 \geq 0 \end{cases}$$

In the absence of time constraints, therefore, the animal would choose its feeding bouts to ensure $b_1 = \omega_1$. The maximum value of u_1 is attained in that case, and is equal to 1.

Drinking

The structures for drinking will be assumed to be identical to those for feeding, except that all the parameters may be different. Thus,

$$f_2(t) = \begin{cases} -\mu_2 t & t \leq t_2^{\min} \\ K_2(1 - \exp(-\lambda_2(t - t_2^{\min}))) - \mu_2 t & t \geq t_2^{\min} \end{cases}$$

$$g_2(t) = -\mu_2 t$$

$$u_2(b_2) = \begin{cases} 0 & b_2 \leq 0 \\ \left(\frac{b_2}{\omega_2}\right)^{\omega_2 \kappa_2} \exp(-\kappa_2(b_2 - \omega_2)) & b_2 \geq 0 \end{cases}$$

Resting

Animals do not usually spend all their time feeding and drinking. 'Resting' is now introduced, as a 'default' activity. Assumptions about default activities may have important implications for how other activities are patterned, as the marginal value of a unit of time is affected by them. In this example, we assume that the animal has a preference for spending half of its time resting. We achieve that by the following definitions:

$$\begin{aligned}f_3(t) &= t \\g_3(t) &= 0 \\u_3(b_3) &= 4b_3(1 - b_3)\end{aligned}$$

We have again ensured that if the animal chose without constraints, the value of its utility fragment would be 1.

Joining three activities together

The final formal structure is the choice of the function u , which combines the utility fragments together into a single overall utility function. The simple assumption to be adopted for the sake of the example is that the total utility is the product of the utility fragments for each activity, formally

$$u(b_1, b_2, b_3) = u_1(b_1)u_2(b_2)u_3(b_3)$$

The fragments are all non-negative by definition, and the formula implies that if any one fragment equals zero, then the overall utility equals zero. Thus we expect all three activities to be present in the optimal sequence if a positive utility is feasible.

The utility also has a maximum possible value of 1, inherited from each of its components.

This choice of utility function ensures a particular pattern of preferences across the three activities, as the effect on utility of trading off between two of them is quite unaffected by the value of the third. This enforces 'quasilinear preferences', and in particular there are no substitution or complement effects. This is of course quite unrealistic, as eating and drinking in many animals are tightly linked physiologically. The numerical simulation is nevertheless of interest, as it helps us to uncover what kinds of behaviour can be explained without recourse to

ideas of substitution and complementarity. These could be further explored by adopting a different function for u .

Numerical simulations

A complete analysis of the model is not attempted here. A numerical investigation is presented that centres on the following parameter values, which should be assumed except where otherwise stated:

$$\begin{aligned}\lambda_1 &= \lambda_2 = 1 \\ \mu_1 &= \mu_2 = 0.01 \\ t_1^{min} &= t_2^{min} = 0.2 \\ \omega_1 &= \omega_2 = 0.1 \\ \kappa_1 &= \kappa_2 = 10 \\ \kappa_1 &= \kappa_2 = 1\end{aligned}$$

Both feeding and drinking have an optimal net rate of 0.1, with quite a sharp decrease in utility for rates above and below. Both feeding and drinking have a requirement that grows constantly at rate 0.01.

Figure 2 shows a schedule of optimal responses as the ω_i are increased, corresponding to increasing the optimal net intake rate. Before $\omega_i = 0.13$, the mink has no marginal value for time. It can attain the optimal benefit rates for feeding and drinking and resting, and must obtain food and water in a time-inefficient manner. Otherwise, it would gain either too much food or water over time, or would have to rest for more than the optimal fraction of time. The mink has time on its hands, as life is very easy. This is reflected in the attained utility of 1.0. Note that feeding is suboptimal so far as feeding efficiency is concerned, but optimal from the point of view of utility.

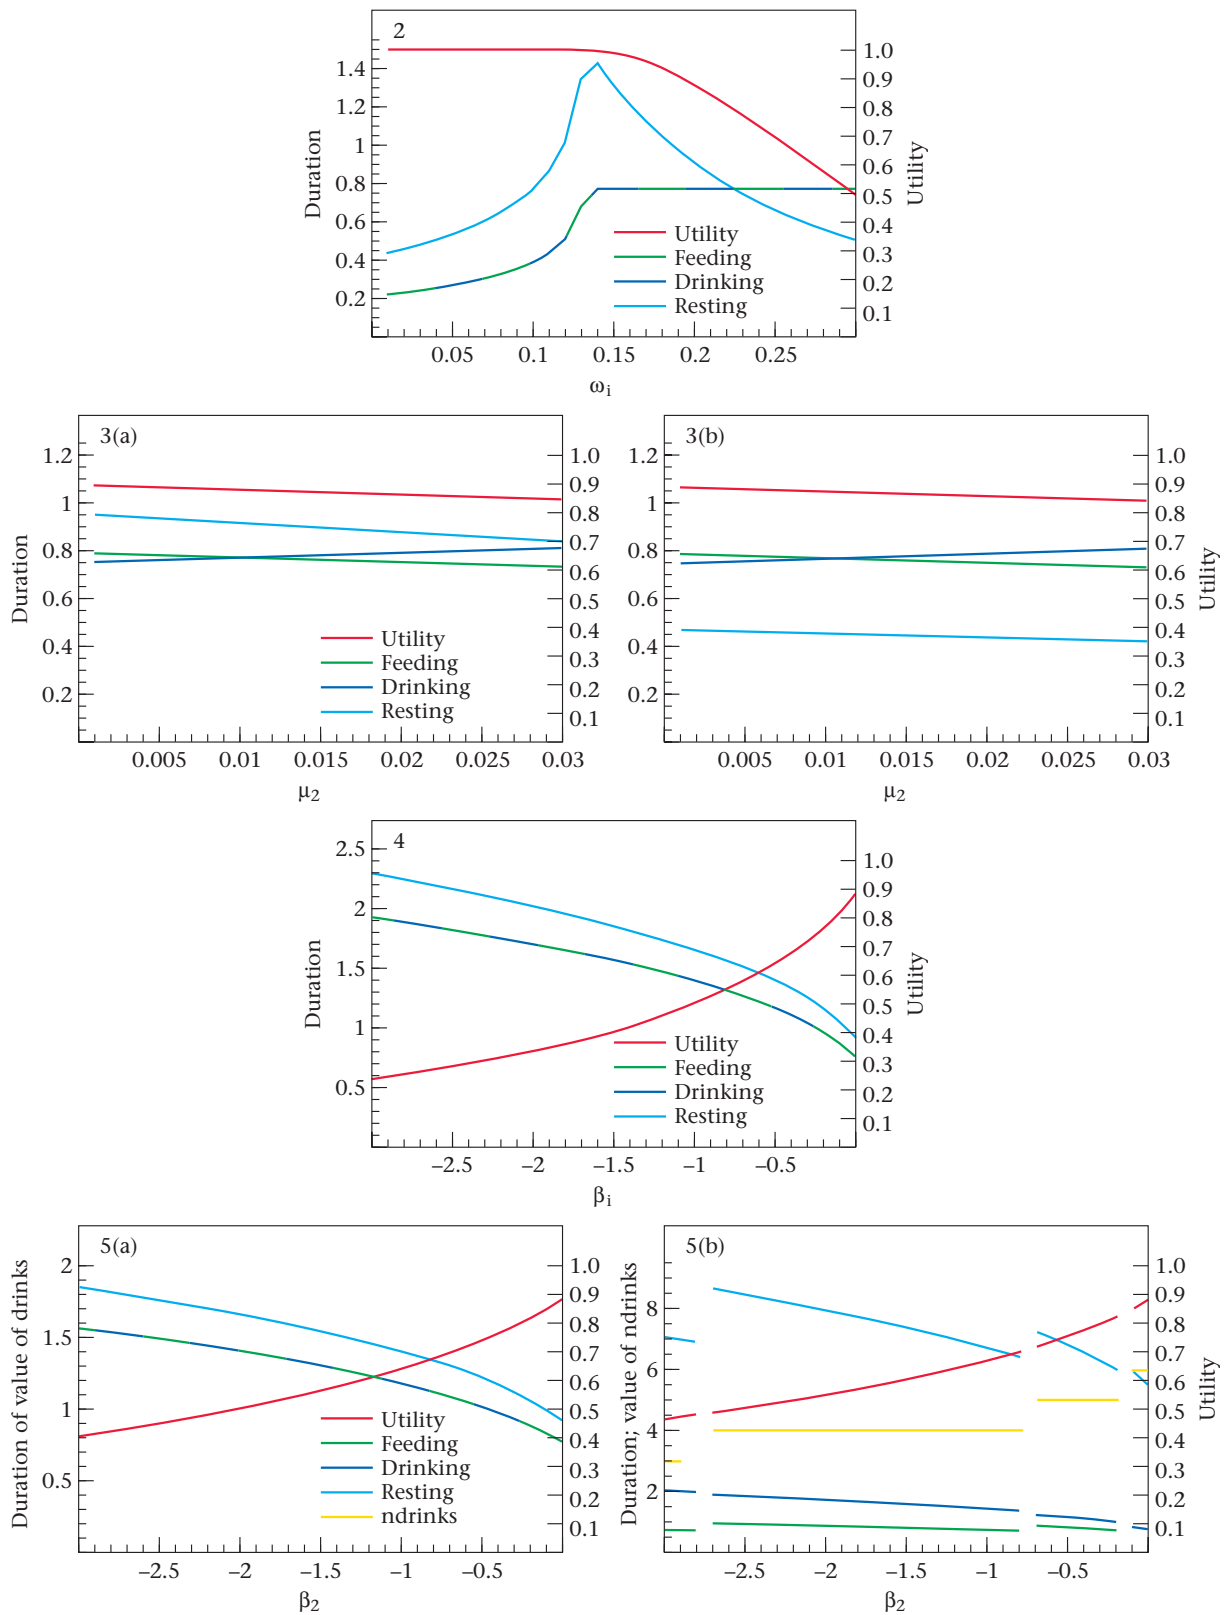
After $\omega_i = 0.13$, the situation for the mink changes. It cannot optimize separately for all three activities. Rest drops from the isolatedly optimal fraction of one half as the food and drink requirements increase. Feeding and drinking now have to be done efficiently, as time has a marginal value. Utility drops below 1.0 once simultaneous unconstrained optimization of all activities becomes impossible.

Figure 2. Effect of varying both resource requirements (ω_i) on optimal behaviour. Resting occupies half of the whole sequence while the ω_i are low enough, but less than half once the requirements have increased beyond $\omega_i = 0.13$. The utility starts to fall at this same point.

Figure 3. Effect of varying rate of water use (μ_2) on optimal behaviour. (a) As the rate of water use increases, drinking increases at the expense of the other two activities. In (b) there are two rest periods. The utilities and durations of eating and drinking bouts are the same as in (a), and the duration of each rest bout is half that of the single rest bout in (a).

Figure 4. Effect of varying both door weightings (β_i) on optimal behaviour. The bout lengths increase, so as to reduce the frequency with which doors must be opened.

Figure 5. (a) Effect of varying door weighting for drinking (β_2) on optimal bout durations within a fixed cycle of activities. The sequence of activities is constrained to have equal numbers of the three activities. The durations all increase, and feeding and drinking durations remain equal. (b) Effect of varying door weighting for drinking (β_2) on optimal behaviour with 12 bouts per sequence available for feeding or drinking. 'ndrinks' is the number out of the 12 bouts that are used for drinking. The curves have all been interrupted at values where ndrinks has changed.



Figures 2-5.

Choosing the values $\omega_i=0.2$, at which time has a marginal value, the value of μ_2 is now varied. As μ_2 increases, the rate of water use increases, and so the requirement for drinking increases. We expect to see the mink reallocate time from both feeding and resting to drinking as μ_2 increases, and Fig. 3 shows this happening. Utility declines as μ_2 increases. Note the expected symmetry between feeding and drinking at $\mu_2=\mu_1=0.01$.

The mink could choose to rest after both feeding and drinking, rather than just once in the sequence. The optimal times within such a sequence have been calculated for the same parameter values as in Fig. 3a. As expected from the linearity of the effect of rest, the drinking and feeding times are unaffected, and the sum of the two resting times equals the length of the single original resting times, over the whole schedule of values of μ_2 . These results are shown in Fig. 3b. However, this linearity of the effect of rest is a very special case, and we would not expect the same result in general. The maximizations were performed using the FindMinimum command in Mathematica 4.0 (Wolfram 1999).

These results may give some feeling for how the model works. The next stage is to add doors of varying weight.

Doors of Varying Weight

At this stage we introduce into the model a representation of the door weighting that Cooper & Mason (2000) studied experimentally. Assume that the feeding and drinking compartments each have a door through which they must be entered, and that by loading these doors, it can be made more or less aversive for the mink to enter the compartments. To make the results of the model easier to interpret, we shall assume that it takes no time to open a door, but that the necessity to open doors affects the utility function. Suppose we measure the cost of opening the feeding and drinking doors as 'negative benefits' $\beta_1 \leq 0$ and $\beta_2 \leq 0$, respectively, and the total cost of opening doors over a behaviour sequence is the sum of the costs. So the 'benefit rate' for door opening, b_D , would be the sum of all the β values over the sequence divided by the duration of the sequence. Then the utility fragment u_D for door opening will be

$$u_D(b_D) = \exp(b_D)$$

which multiplies with the others to form the whole utility function.

To complete the formalities according to which this new model fits into the optimization program given earlier, we introduce two new activities numbered 4 and 5 to represent opening the feeding and drinking door, respectively; we define $f_4(t)=\beta_1$, and $f_5(t)=\beta_2$, and $g_4(t)=g_5(t)=0$; we choose the constraint set C so that activities 4 and 5 have a duration of zero; and that activities must be chosen by sampling repeatedly from the sequences (4, 1), (5, 2) and (3). In further discussion, we will not mention the door openings as activities, but it is important that all of these models fall within the same formal optimization program.

Figure 4 shows the effect on the mink's optimal sequence of having the same value of β for both doors, for a range of values of β . The heavier the door, the longer the duration of all bouts. This is because the doors must be opened once each per cycle, and extending the duration of the cycle reduces the frequency with which they need to be opened. Note that utility decreases as β_i decreases.

Figure 5a has $\beta_1=0$ but varies β_2 , and shows the effect on the lengths of bouts in the fixed sequence (feeding, drinking, resting). Note that durations of feeding and drinking bouts are the same here. Although the drinking door is weighted, it must be opened exactly once in each sequence and so does not produce a difference between feeding and drinking. To reduce the rate of disutility from opening the door, the durations of all three activities increase as the door weighting increases. Utility, unsurprisingly, falls as the door becomes heavier. One technical point is reassuring. The durations for $\beta_1=0$, $\beta_2=-2.8$ shown in Fig. 5a are the same as for $\beta_1=\beta_2=-1.4$ in Fig. 4. The sum of β s over the cycle is the same, and as they combine linearly the effect on durations should be, and is, the same.

However, the fixed sequence (feeding, drinking, resting) does not remain optimal as β_2 changes. As the cost increases, there is a tendency for the optimal sequence to contain relatively fewer but longer bouts of drinking. Figure 5b shows the properties of the optimal sequences out of those with exactly 12 bouts shared between feeding and drinking, for a range of β_2 values. The balance is six of each activity with equal costs ($\beta_1=\beta_2=0$), but a decreasing number of longer drinking bouts as β_2 increases. In theory, there would be a continuous shift in the fraction of bouts that were drinking bouts, and a continuous increase in their duration, and Fig. 5b shows an approximation to that continuous response. During a sequence of β_2 values for which the number of drinking bouts remains constant, all the durations increase, to dilute the fixed penalty of opening doors over a longer cycle. When the number of drinking bouts decreases at a particular value of β_2 , this produces a reduction in that penalty, and the duration of resting bouts springs back. The durations of feeding and drinking bouts change at these boundaries for that reason, but also in response to the change in the number of feeding and drinking bouts in the cycle.

DISCUSSION

The model presented in this paper has a very wide range of application in principle, but has so far been applied in only one, very special, hypothetical case. One general conclusion that can be drawn so far is the importance of assumptions made about the default activity and the possibility of 'deliberate inefficiency' when the animal is easily attaining the maximum possible utility.

The results shown in the figures formalize some of the intuitions of Cooper & Mason (2000), which are themselves borne out by the data they present, about the effect of altering door weights, and hold out the prospect of quantitative predictions of experimental results. The model is much more appropriate in its assumptions to

experiments of this kind than the standard economic model of the consumer.

Adapting the model to analyse experimental results could involve estimating parameters of the model. Thus, the values of the utility decrement from opening a door of varying weights could be discovered by testing different weights in turn, and inferring a value of β for each weight, from the resulting behaviour. The most useful level of behaviour to quantify in this case might be the frequency of bouts of each type of behaviour, and the average duration of each type of bout. This very simple example suggests that the general approach could have a role in designing, analysing and interpreting experiments on sequences of behaviour.

The main limitation of the model is that it ignores the state of the animal, and that may at some stage require the construction of a more complex model. However, there does seem to be scope for representing current ideas about behaviour sequences in this simpler framework. For example, the question of whether deficits build up for certain kinds of behaviour should be solvable.

A neighbouring literature to that on animal welfare is the psychological literature on reinforcement schedules, and I am grateful to Professor Alex Kacelnik for suggesting the possibility of applying this kind of model to those experiments.

In general terms, even human economic consumers make decisions about sequences of behaviour, and the micro-economic model of the consumer is probably inadequate to account for all consumer behaviour, just as it is inadequate to account for mink behaviour.

Acknowledgments

I am very grateful to Dr Georgia Mason for helpful discussions and for useful comments on the manuscript. Professor Richard Dawkins made helpful remarks about the links between different kinds of optimization in animals, and Professor Marian Dawkins encouraged a clearer statement of the link between the animal's utility function and the animal's welfare. Dr Ezequiel Di Paolo made some very acute and sympathetic comments on the manuscript.

References

- Cooper, J. J. & Mason, G. J. 2000. Increasing costs of access to resources cause re-scheduling of behaviour in American mink *Mustela vison*: implications for the assessment of behavioural priorities. *Applied Animal Behaviour Science*, **66**, 135–151.
- Dawkins, M. S. 1981. Battery hens name their price: consumer demand theory and the measurement of ethological needs. *Animal Behaviour*, **31**, 1192–1205.
- Houston, A. I. & McNamara, J. 1999. *Models of Adaptive Behaviour: An Approach Based on State*. Cambridge: Cambridge University Press.
- McFarland, D. J. & Houston, A. I. 1981. *Quantitative Ethology: The State Space Approach*. London: Pitman.
- Mas-Colell, A., Whinston, M. D. & Green, J. R. 1995. *Micro-economic Theory*. Oxford: Oxford University Press.
- Mason, G. J., Cooper, J. & Clareborough, C. 2001. Frustrations of fur-farmed mink. *Nature*, **410**, 35–36.
- Wolfram, S. 1999. *The Mathematica Book*. Cambridge: Cambridge University Press.