

Starlings exploiting patches: the effect of recent experience on foraging decisions

INNES C. CUTHILL*†, ALEJANDRO KACELNIK‡§, JOHN R. KREBS*,
PATSY HACCOU** & YOH IWASA††

*Edward Grey Institute, Department of Zoology, South Parks Road, Oxford OX1 3PS, U.K.

‡King's College Research Centre, King's College, Cambridge CB2 1ST, U.K.

**Institute of Theoretical Biology, Groenhovenstraat 5, 2311 BT Leiden, The Netherlands

††Department of Biology, Faculty of Science, Kyushu University, Fukuoka 812, Japan

Abstract. Laboratory and field experiments have shown that, as predicted by the marginal value model, starlings, *Sturnus vulgaris*, stay longer in a food patch when the average travel time between patches is long. A laboratory analogue of a patchy environment was used to investigate how starlings respond to rapidly fluctuating changes in travel time in order to find out the length of experience over which information is integrated. When there was a progressive increase in the amount of work required to obtain successive food items in a patch (experiment 1), birds consistently took more prey after long than after short travel times; travel experience before the most recent had no effect on the number of prey taken. Such behaviour does not maximize the rate of energy intake in this environment. The possibility that this is the result of a simple constraint on crop capacity is rejected as, when successive prey were equally easy to obtain up until a stepwise depletion of the patch (experiment 2), birds took equal numbers of prey per visit after long and short travel times: the rate-maximizing behaviour. A series of models are developed to suggest the possible constraints on optimal behaviour that affect starlings in the type of environment mimicked by experiment 1.

Our aim in this study is to investigate how the time spent foraging by starlings, *Sturnus vulgaris*, in a food patch is influenced by recent experience. Optimality models show that in a patchy environment, long-term rate of energy gain is maximized by staying in each patch for an amount of time dependent upon the average patch quality and the average travel time between patches (Charnov 1976; Parker & Stewart 1976; literature review in Stephens & Krebs 1986). The average to which we refer is the 'average for the environment', so that in different environments the rate-maximizing patch time may differ. However, the classical optimality model of this problem (the marginal value theorem, Charnov 1976) makes no reference to how parameters should be estimated; that is, over what length of experience information should be integrated. Cowie (1977) suggested a 'memory window' of recent experience that would allow a forager to adjust its behaviour to current conditions, but this concept was never precisely formulated. More quantitative models of the influence of past experience on foraging decisions are common in the oper-

ant psychological literature (see review by Kacelnik et al. 1986), although the adaptive significance of such decision rules is rarely examined (but see Dow & Lea 1986; McNamara & Houston 1987).

In this paper we focus on one particular aspect of change in the environment, namely changes in travel time. We consider the optimal (rate-maximizing) strategy for each of four different classes of change in travel time. The experiments investigate how foraging decisions made by starlings in a patchy environment are influenced by recent experience of different travel times.

Previous work (Kacelnik 1984; Cuthill 1985) has shown that when starlings forage in a patchy environment, the relationship between the average travel time and the average amount of food taken from each patch can be accurately predicted by the marginal value theorem. In the field experiment of Kacelnik (1984), inspired by Tinbergen's (1981) observations of natural foraging, parent starlings were trained to collect food for their young from a mealworm dispenser and travel time was manipulated by altering the distance between the feeder and the nest. In a previous laboratory study (Cuthill 1985) travel time was manipulated by varying the number of times the bird had to fly between two perches before it could gain access to a patch

†Present address: Department of Zoology, University of Bristol, Woodland Road, Bristol BS8 1UG, U.K.

§Present address: Edward Grey Institute, Department of Zoology, South Parks Road, Oxford OX1 3PS, U.K.

consisting of a feeder that dispensed turkey starter crumbs. In both studies the birds experienced diminishing returns within each patch visit: this was achieved by making the interval between successive prey items progressively longer, either, as in Kacelnik's study, by increasing the time (a 'progressive interval schedule') or, as in Cuthill's study, by increasing the amount of work (perch-hopping) required to obtain successive prey (a 'progressive ratio schedule'). The dependent variable in both studies was the number of prey taken per visit to a patch ('prey per visit'). In the field study, prey were mostly carried to the nestlings, whilst in the laboratory study they were all consumed by the foraging bird itself.

From these studies, we conclude not only that the average number of prey taken per visit is accurately described by a model of rate maximization, but also that starlings are able to adjust the number of prey taken in relation to their past experience of travel times: when the average travel time was increased, the birds on average took more prey per visit. However, in both experiments travel time was held constant within an experimental session, so it is unclear how much of their past experience subjects used: responding solely to the last travel time would be equivalent to a decision based on the long-term average. Our aim in the present study is to investigate in more detail how recent experience of travel times influences the number of prey taken per visit. In the experiments described below we analyse how starlings respond to short-term changes in travel time within an experimental session.

Optimality Models: Rules for Rate-maximizers

If a rate-maximizing forager experiences short-term changes in travel time from one patch to the next, how should it respond? Should it track the changes by adjusting the prey taken per visit, or should it respond to the average travel time and thus ignore the fluctuations? Assume that the currency is energy rate maximization and that the time available for foraging is very long (an effectively infinite time horizon). The answer depends on how the travel time changes. Two aspects of changes are important: first, whether changes in travel time occur with respect to time, independently of the animal's actions (time-dependent changes), or independent of time, at a rate determined by the animal's actions (animal-dependent changes); and second, whether or not successive travel times are correlated or uncorrelated with one another. To

illustrate these points we consider the example of an environment in which all patches are identical and in which travel time changes between two values, long and short, and ask whether the forager should employ a fixed patch-residence strategy related to the average travel time (which we term 'averaging') or adjust its patch residence time to recently experienced long or short travel times ('tracking').

Consider first the environment in which travel time changes in an animal-dependent fashion, that is, the probability of the travel time changing from long to short, or vice versa, is fixed per patch-use cycle rather than being dependent upon time per se. Stephens & Krebs (1986) showed that for this sort of environmental change, the animal should use a fixed patch-residence strategy dependent upon the average travel time. This is because the expected rate of gain from leaving a patch (a function of average patch quality and average travel time) is independent of the animal's recent experience and current behaviour. This is true whether there are long runs of long or short travel times (correlated environment) or changes occur with probability 0.5 per patch-use cycle.

Turn now to the environment in which changes are time-dependent. To get an indication of the optimal policies in this type of environment, we assume as a first (and mathematically tractable) approximation that environmental changes occur according to a continuous time Markov chain (Appendix 1). The crucial difference between this and the sequential environment is that changes in the travel time can occur without the animal leaving a patch. If a change from long to short travel times occurs while the animal is exploiting a patch, the forager would benefit by reducing the amount of travelling it had to do in the period when travel time was long. Thus, in general, it is possible to imagine that by staying longer in a patch following a long travel time, the forager would be able to increase the likelihood that expected conditions would have changed towards the average for the environment (i.e. changed for the better) by the time it leaves the patch. Similarly, when travel times are short, by leaving a patch early, the animal reduces the probability of the expected travel time having returned to the average (i.e. changed for the worse). In this illustrative example, the travel time is assumed to have a fairly high probability of remaining the same in successive time intervals. Therefore, staying longer after a long travel time and vice versa

Table 1. Classification of environments according to whether (a) changes in travel time between patches occur relative to elapsed time ('time-dependent') or with respect to patch-use cycle ('animal-dependent') and (b) successive travel times are correlated or uncorrelated with one another

Change	Travel sequence	
	Correlated	Uncorrelated
Animal-dependent	Average	Average
Time-dependent	Track	Average

In each cell the optimal policy for a rate-maximizer is indicated. 'Average' means adjust patch time in relation to the long-term average travel time, whilst 'track' means adjust patch time in relation to short-term changes in travel time.

is the optimal policy for a rate-maximizer in an environment where changes in travel time are time-dependent and where successive travel times are strongly correlated. The weaker the correlation, the weaker the influence of recent experience and, if successive travel times are uncorrelated, the optimal policy for the rate-maximizer is to respond to the average travel time (Appendix 1).

To summarize, the two dichotomies discussed here, time-dependent versus animal-dependent and correlated versus uncorrelated, allow us to classify foraging environments into four possible sets (Table 1). In three of the four sets the optimal strategy for the rate-maximizer is to adjust to the long-term average travel time, whilst in one, it pays to track short-term changes in travel time.

In the experiments that follow, we consider one of the sets we have described, in which successive travel times are uncorrelated, and changes occur with respect to patch-use cycle, rather than being time-dependent.

GENERAL METHODS

We used captive starlings foraging in a cage similar to that used by Cuthill (1985), described in Kacelnik & Cuthill (1987). The bird travelled by flying between two perches a predetermined number of times. When it had completed this travel requirement, it foraged in a patch consisting of an operant feeder and a perch on which the bird had to hop to gain access to the food. The feeder delivered rewards according to a preprogrammed schedule. If a bird left the patch and returned to travelling,




Figure 1. Plan view of an experimental aviary. Each such cage houses one starling.

the schedule was reset to its initial value and the bird had to complete the travel requirement before it could resume feeding in the patch. Thus at any moment in the patch, the bird had the choice between continuing to work in the present patch and travelling to a 'new' patch. In experiment 1, the reinforcement schedule in the patch was a progressive ratio schedule in which each successive reward required more hops, mimicking the naturally occurring phenomenon of resource depression (gradually diminishing returns). In experiment 2, the schedule was a fixed ratio, each reward being offered after a fixed number of hops, with a maximum value beyond which no further rewards could be obtained (stepwise depletion, or in operant jargon, 'a sudden transition to extinction').

The travel requirement in both experiments was either long (60 flights) or short (five flights), the two values occurring in a random sequence ($P=0.5$). By examining the birds' persistence in the patch, measured as prey taken per visit, following specific sequences of travel requirements we were able to infer how recent experience influences patch-leaving decisions.

Birds

We used 12 wild-caught, experimentally naive, starlings (six in experiment 1 and six in experiment 2). The birds were held in large outdoor flight cages prior to testing where they were fed on a mixture of Sluis insectivorous bird food, turkey starter crumbs and mealworms.

Apparatus

The birds were individually housed in T-shaped cages (Fig. 1). Three cages were in a room measuring 4×3.5 m. The two travel perches were 1.2 m apart and 0.3 m from the floor of the cages. The feeder was 5 cm above the floor on one wall of the cage; the feeder perch was 0.5 m away from the

feeder and 15 cm off the floor. The feeder itself consisted of a gravity feed, cylindrical Perspex reservoir that could be reloaded with turkey starter crumbs from outside the cage and an access hole measuring 1 cm in diameter, covered by a metal lid. The bird could open the feeder by breaking an infrared beam above the metal lid, which caused the lid to slide to one side for a preset period of 2 s, during which the bird could collect an average of 0.08 g of food.

The cycle of events, consisting of a visit to a patch followed by a period of travel, was signalled to the birds by coloured lights as follows. When the bird was 'in a patch', a red light on the feeder perch was on, indicating that the bird could gain rewards according to the programmed schedule by hopping on the feeder perch. At the same time both travel perches had an illuminated green light, indicating to the bird that it could resume travelling at any time instead of foraging in the patch. If the bird left the patch and alighted on one of the two travel perches, both this travel perch light and the feeder perch light were extinguished, while the light on the other travel perch remained on. At this point the bird could not gain access to the feeder until completing the programmed travel requirement of either five or 60 flights. When the bird flew to the other travel perch, this light was extinguished and the light on the perch it had just left came on. Thus the bird could use the rule 'fly to the green light' to complete its travel requirement. On completion of the travel requirement, both perch lights went off and simultaneously the red feeder perch light came on. The bird then hopped on the feeder perch until, according to the schedule of reinforcement, a red light on the feeder itself, next to the access hole, came on. The bird could then gain access to the food for 2 s by approaching the feeder and breaking the infrared beam, which caused both the lid to open and the red feeder light to go out. Each successive reward was signalled by the red feeder light, whilst the feeder perch light remained on throughout the patch visit. The progressive ratio schedule for experiment 1 offered rewards following a rounded-down integer approximation of the function $h = \exp(G/1.5)$, where G is the cumulative number of rewards since arrival at the patch and h is the cumulative number of hops on the feeder perch. Thus rewards were delivered after 1, 3, 7, 14, 28, 54, 106 ... cumulative hops on the feeder perch. Because the birds varied in their rate of hopping, their gain functions in time also varied (Fig. 2a).

In experiment 2, the gain function was programmed as a fixed ratio schedule with rewards offered after every five hops. After the fifth reward the schedule was terminated, so that no more rewards could be obtained within a patch without travelling to reset the patch to the start of the schedule. As in experiment 1, the cue light in the patch remained on until the bird started travelling. Figure 2b shows the linear gain functions experienced by the six birds in experiment 2.

The cue lights, reward schedules and feeder opening were controlled by an Acorn System 4 microcomputer, which also recorded the following data: travel time (the time from the first landing at one of the travel perches to the first landing on the feeder perch), the time of each reward, number of rewards per patch visit and duration of the patch visit (time from the first hop at the feeder perch to the first landing on one of the travel perches).

Water was provided in drinking fountains placed next to the feeder and the birds were given bathing water once per week.

Protocol

The experiments were conducted in a closed economy (Hursh 1980). Thus the birds lived in the experimental cages throughout the experiment and obtained all their food from the feeders, apart from a supplement of mealworms given at the end of the day. The feeder access time and travel requirements were chosen so that the birds would have to work for most of the day in order to maintain their body weight. Each day the lights came on at 0900 and went out at 1700 hours.

Before we started to record the data, the birds were left in the cages to get used to the feeder, which was left permanently open at this point, and then gradually trained to open the feeder, hop on the feeder perch and eventually to complete a travel requirement in order to activate the feeder perch. This whole sequence usually lasted about 3 weeks and ended with the birds performing the complete experimental procedure with a travel time intermediate between the long and short values (32 flights). In each experiment birds remained in the final experimental conditions for 12–14 days. We present analyses of data from the last 6 days, when all individuals were fully familiar with the schedules in operation.

The first group of six birds was tested with the depleting (progressive ratio) schedule and the




Figure 2. Cumulative prey gain as a function of median time in the patch for each of the birds in (a) experiment 1 (depleting patches), and (b) experiment 2 (linear gain function). Ordinate and abscissa should strictly be reversed, but are plotted according to the 'marginal value theorem' convention. Δ : bird 0; \circ : bird 1; \bullet : bird 2; \diamond : bird 3; \blacksquare : bird 4; \square : bird 5.

second group of six with the linear (fixed ratio) schedule. In the second group, three birds were observed for 1 h per day over a period of 10 days. The purpose of these observations was to examine how the birds actually spent their time in the experiment, especially how they allocated time to foraging (travelling and working in the patch) as opposed to other activities such as preening, singing and resting. The observations were recorded on a computer-compatible event recorder for subsequent analysis.

Time Inhomogeneity

A potential problem with data from closed economy experiments is that the behaviour of the animals changes during the course of the day, either because of endogenous rhythms or because of

extrinsic events such as after-effects of cleaning the cages. We refer to these changes as time inhomogeneities. To take this effect into account in our analyses, the data from experiment 1 were inspected for time inhomogeneity by means of cumulative plots of the number of items taken, the patch residence times and the travel times, against the sequence number. These plots indicate whether global changes occur and, if so, if these changes are continuous or more or less abrupt. In the latter case, a change in the slope of the plot indicates the place of a change point (Haccou et al. 1983). The plots can also indicate outliers, i.e. extremely large values.

On the basis of visual inspection of these plots, outliers were removed (<1% of the data) and the records were divided into six periods: 0930–1030,

1030–1200, 1200–1300, 1300–1400, 1400–1500 and 1500–1700 hours. The choice of these periods was a compromise between the need for homogeneous data within a period and the need for sufficient numbers of data points within each period for analysis. Thus, some of the periods still contain change points, especially the first two periods, which follow routine cleaning of the cages. In the analysis of the results of experiment 2, the same periods were used. We expect these data to be less inhomogeneous, since here cleaning took place well before or after the period of the day during which data were recorded. Therefore, for these results the subdivision is probably conservative. In the analyses of both experiments, the effect of period of the day is included.

EXPERIMENT 1: UNPREDICTABLE TRAVEL TIMES AND DEPLETING PATCHES

Our aim in this experiment was to analyse the birds' response to an unpredictable sequence of travel requirements when there were gradually diminishing returns within a patch. The dependent variable used to measure the birds' response was the number of prey taken per patch visit. The experimental environment corresponded to the top right-hand cell of Table I: successive travel requirements were uncorrelated and the changes of long to short travel times occurred with probability 0.5 per patch-use cycle, independent of elapsed time.

Procedure

The travel requirement was either five or 60 flights. At the start of each period of travel, the value was chosen at random with a probability of 0.5 for both long and short. Thus on average the birds experienced an equal number of long and short travel requirements and the expected run length (number of successive travel requirements of similar length) was two.

Methods of Analysis

The data were analysed with an analysis of covariance (ANCOVA, Scheffé 1959) in which the number of prey taken per visit was the dependent variable and the following five factors were independent variables. (1) Travel history is the effect of recently experienced travel requirements. Recent experience was defined as including the last three




Figure 3. Experiment 1: number of prey taken per patch visit following short travel requirements (▨) and long travel requirements (□). Also shown, for each bird, are the values predicted by the marginal value theorem if the bird estimated travel time from the previous travel time alone (■: predicted prey per visit following short travel requirement, +: predicted following long travel requirement).

travel/patch-use cycles. This gives eight possible levels of recent experience, corresponding to all possible combinations of long and short travel requirements over the preceding three cycles. The other factors were (2) bird, testing for individual differences, (3) period, concerning variation over the periods of the day that were defined in the Methods under the heading of time inhomogeneity, (4) day, testing for daily variation, and (5) last patch visit, entered as a covariate, testing for an effect of the number of prey eaten during the preceding patch visit on the number eaten in the present visit.

Of these, travel history is the effect of prime interest and bird, day and period are essentially nuisance variables that have to be controlled for in the analysis. Bird, day and period had a significant three-way interaction, but no obvious pattern could be discerned in this variation. Whilst bird and day could be considered to be random effects, multiple mixed models are hard to interpret (see e.g. Armitage & Berry 1987) and with only six birds in the experiment we simply treated these factors as fixed effects and limit our interpretation of results to these birds alone. Analysis of the residuals revealed no significant departures from assumptions of the ANCOVA model.

Results

Travel history

Our main purpose in the experiment was to analyse the effect of preceding travel requirements

Table II. Experiment 1: results of ANCOVA on the mean number of prey taken per patch visit

Source of variation	<i>df</i>	<i>F</i> *
Last meal	1, 817	68.78
Bird	5, 817	53.30
Day	5, 817	8.56
Period	5, 817	6.33
History	7, 817	42.09
Bird*Day*Period	123, 817	2.37

Travel history is the factor of prime interest, representing the effect of the three most recent travel times on the mean number of prey taken per visit. Bird, day and period (of day), whilst of secondary interest, control for unexplained daily and individual variation. The covariate last meal is the number of prey taken per visit taken on the previous patch visit.

*All *F*-ratios significant at $P < 0.001$.

on the number of prey per visit. All birds took more prey per patch visit after a long travel requirement than after a short one (Fig. 3). The effect of travel history on prey taken per visit was highly significant, as were the effects of other factors included in the ANCOVA (Table II). There were no significant interaction effects between travel requirement and other factors, which indicates that there were no individual differences or daily or periodic variations in the way recent travel experience affected the birds' behaviour.

To analyse in more detail exactly how recent travel experience affects the number of prey taken per visit, we compared different sequences of travel histories. As we have already mentioned, the analysis considered sequences of the last three travel requirements. By pairwise comparison of the effects of these sequences it is possible to find out exactly which part of the previous experience determines the birds' behaviour. For instance, let SSL denote the sequence short-short-long (that is, the most recently experienced travel requirement is long). By comparing prey taken per visit after experiencing SSL with that after SSS, the effect of the immediately preceding travel requirement can be calculated, since the only difference between these two sequences is the last travel requirement. Similarly, effects of the penultimate travel requirement can be studied by comparing prey taken per visit after, for instance, SLS and SSS. Testing such pairwise differences by means of contrasts (Scheffé 1959) is a more sensitive analysis than a test for the averaged

effect of, say, the last travel requirement, especially when various components of past experience are not independent (Armitage & Berry 1987). The contrasts for testing effects of last, penultimate and antepenultimate travel requirements are given in Table III. The value of a contrast gives the difference in average prey taken per visit. For instance, the birds take on average 0.80 food items more after the travel sequence SSL than after SSS (Table III). These results show that the only part of the previous experience influencing the number of prey taken per visit is the most recent travel requirement. The level of patch exploitation does not appear to be influenced at all by travel experience further back in time.

These results show that the birds track the changes in travel time by staying longer in a patch after a long travel requirement than after a short one. One possible explanation of their behaviour is that the birds behave according to the marginal value model, but based only on the most recent travel time. To test this idea, we estimated the expected numbers of prey taken per visit after long and short travel requirements, based on the marginal value model, for environments with only one type of travel requirement. The calculations were based on the median travel times per bird and the estimated gain functions (Fig. 2a). The observed number of prey taken per visit changed less than expected if the birds were behaving according to the marginal value model basing the travel time only on the last trip (Fig. 3).

Last patch visit

The number of prey eaten on the preceding patch visit is positively correlated with the number eaten during the current visit. The estimated regression coefficient is 0.26 (± 0.06 , 95% confidence intervals) and the correlation coefficient 0.28.

Conclusion

The main conclusion of experiment 1 is that the number of prey taken per patch visit is greater after long than after short travel times and that travel times before the most recent have no significant effect on patch exploitation.

A simple explanation for the strong influence of previous travel requirement on rewards consumed could be a digestive constraint on the number of

Table III. Pairwise-comparisons of the mean number of prey taken per visit following different travel histories, using contrasts

Most recent travel requirement		Penultimate travel requirement		Antepenultimate travel requirement	
Classes	Contrasts	Classes	Contrasts	Classes	Contrasts
SSL-SSS	0.80*	SLS-SSS	-0.10	LSL-SSL	0.00
SLL-SLS	0.84*	SLL-SSL	-0.06	LSS-SSS	0.13
LSL-LSS	0.94*	LLS-LSS	0.09	LLL-SLL	0.18
LLL-LLS	0.67*	LLL-LSL	-0.18	LLS-SLS	0.01

* $P < 0.05$. Contrasts represent the difference in the mean number of prey taken per visit, with all other effects in the ANCOVA (Table II) controlled for. Each pair differs in only one of the three most recent travel requirements, e.g. SSL (short-short-long) versus SSS (short-short-short). All contrasts in which the most recent travel requirement differs are significant, but no others.

prey taken per visit. This is analysed further in the next experiment.

EXPERIMENT 2: PATCHES WITH A LINEAR GAIN FUNCTION

Our purpose in this experiment is to test the hypothesis that the results of experiment 1 are caused by a digestive constraint. Assume that the bird fills its crop or gut to the same level at every patch visit; the capacity to eat at the next feeding opportunity will be related to the time (and/or energy expenditure) until that opportunity occurs. Thus, following a strategy of topping up to the same level at each meal would result in larger meals after long travel requirements (long crop-emptying episodes) than after short ones.

However, the results of experiment 1 (last patch visit analysis) showed that the numbers of prey taken on successive visits are positively rather than negatively correlated, as one would have expected if there was a digestive bottleneck. Furthermore, an additional analysis showed that the time since the last patch visit had no effect on the number of prey taken per visit, again arguing against a digestive bottleneck. Nevertheless, a more direct test of the digestive constraint hypothesis is necessary and that is the aim of experiment 2.

In this experiment we used the same varying travel schedule as in experiment 1, but instead of the depleting patch schedule, we used a non-depleting ('fixed ratio') patch schedule. Up to five consecutive rewards could be obtained by performing the same number of perch hops each time, after

which the patch depleted/extinguished in a stepwise fashion. We refer to this schedule as a 'linear gain function'. The schedule was chosen such that the amount of work necessary to obtain the maximum number of rewards (25 hops) was equivalent to that needed to obtain the same number of rewards in experiment 1, where the work requirement increased exponentially. The maximum number of prey taken per visit was fixed at five as this was at the upper end of the range of the median number of prey taken per visit by birds in experiment 1. Naturally, if such a non-depleting schedule had no truncation point, digestive constraints might eventually become important.

If the results of experiment 1 are explicable in terms of a fixed crop-filling strategy, the birds should also take more prey per visit after longer travel requirements in experiment 2. If, on the other hand, the birds use a rate-maximizing criterion, then all the models discussed in the Introduction predict they should take the maximum number of prey per visit, regardless of travel requirements.

Procedure

We exposed the birds to two travel requirements, five or 60 flights, with a probability of 0.5 per patch-use cycle, as in experiment 1. The only difference was that the progressive ratio in the patch was replaced by a fixed ratio in which rewards were delivered every five hops, up to a maximum of five rewards. The work done for five rewards (25 hops) is thus close to that done for an equivalent number

of rewards in experiment 1 (28 hops). The median number of rewards taken after a short travel requirement in experiment 1 ranged from one (bird 2) to four (bird 5), with no upper quartile exceeding four rewards (in 14 hops).

As birds worked under closed economy conditions, they naturally performed other, non-foraging, activities during the experiment. In an attempt to quantify these activities three birds were directly observed over the course of the experiment. Using an event recorder, any two birds were watched for 1 h each day, between 1000 and 1200 hours, in a rota balanced for order. Location of each bird was recorded (travel perch, feeder, cage floor, etc.), as well as eating, drinking, preening, singing, travelling, bill wiping, feather ruffling, probing (where the bill is opened to prise open, say, the cage-lining material) and 'exploring' (where the bird walks around the bottom of the cage, apparently investigating its surroundings).

Analysis

The records were first divided into the six time periods per day used for the analysis of experiment 1. To study the effect of the previous travel requirement on the number of rewards taken, we considered the proportion of patch visits in which the maximum possible amount (five items) was taken. This, rather than the mean number of items taken, was the preferable statistic as we were interested in detecting any deviations from taking five items, and a test based on these proportions is more sensitive to such deviations.

Let i be the period number, j the bird number, k the previous travel requirement and p_{ijk} the chance that five rewards were taken under conditions i , j and k . We assume that these chances do not depend on the day of observation and that the number of items taken at different patch visits are independent. Under these conditions, the number of times that bird j takes five items in period i after travel requirement k is binomially distributed with parameters N_{ijk} and p_{ijk} . Here, N_{ijk} is the total number of patch visits under conditions i , j and k . (All tests were performed conditionally on the observed N_{ijk} .)

We want to infer whether the chances p_{ijk} depend on the preceding travel requirement and whether they deviate much from 1. The analysis consists of a sequence of likelihood ratio tests (see e.g. Cox & Hinkley 1974).

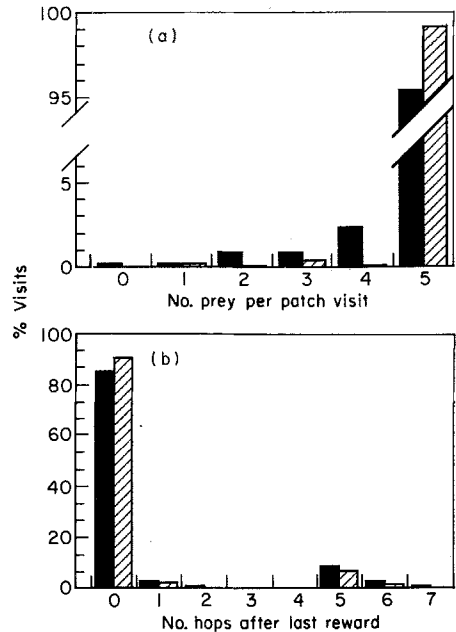


Figure 4. Experiment 2: frequency distributions, following short (■) and long (▨) travel requirements, of (a) the number of prey taken per patch visit, and (b) the number of hops on the feeder perch after the last reward was taken (visits of five rewards duration only). Pooled data from all birds (total number of patch visits following short travel requirements = 887, after long travel requirements = 819).

Results

Prey taken per visit

The main result is that there is no effect of preceding travel time on the probability of taking five items (Fig. 4a), although there are significant differences between periods and/or birds (Table IV). To distinguish these effects we then tested separately whether there are differences between periods and between birds. Both tests gave significant results (Table IV).

Persistence in the patch

Analysing those visits in which birds took the maximum five rewards, it is clear that birds usually ceased hopping immediately after the last reward (Fig. 4b), even though patch extinction was not cued directly. This indicates sensitivity to the schedule being used, coupled with the fact that when birds did continue hopping (for a non-existent sixth reward) the modal figure was five times. Using likelihood ratio tests, the proportion of visits in which

Table IV. Experiment 2: likelihood ratio tests on the proportion of visits that were of the maximum possible duration (five rewards)

Test	Null hypothesis	Alternative	Statistic*	df	P
1	$p_{ijk} = p_{ij}$	p_{ijk} differ	46.41	36	0.11
2	$p_{ijk} = p$	$p_{ijk} = p_{ij}$	87.11	35	$< 10^{-5}$
3	$p_{ijk} = p_j$	$p_{ijk} = p_{ij}$	65.23	30	$< 10^{-5}$
4	$p_{ijk} = p_i$	$p_{ijk} = p_{ij}$	61.54	30	$< 10^{-3}$

p_{ijk} is the probability that five rewards were taken in period of day i , by bird j , following travel requirement k (long or short). There is no effect of previous travel requirement on the probability of taking five rewards (test 1), despite variation between individuals and/or periods (test 2). Between-period (test 3) and individual differences (test 4) are both significant.

*The test statistic has an approximately chi-squared distribution.

birds hopped at least once after the last reward, was greater following short than long travel requirements (likelihood ratio $\chi^2 = 8.94$, $df = 1$, $P = 0.003$).

Time budgets

Direct observation revealed that on average 30–50% of the birds' time was spent on activities that were not directly related to eating and drinking (as food was dry, birds drank regularly after rewards). If we lump these 'non-foraging' activities together, birds 0 and 2 performed more of them following short than long travel requirements (Table V). For both, a major contributor to this difference was more preening, feather ruffling and bill wiping following short travel requirements, and in the case of bird 0 (a male) singing tended to follow short travel requirements rather than long ones. Bird 1 showed no obvious difference in the allocation of time to other behaviour patterns following long versus short travel requirements.

Conclusion

The results do not support the 'gut constraint hypothesis', since travel requirement had no effect on the number of prey taken per visit. Although the numbers vary between birds and periods, the birds usually took the maximum number of prey and left immediately ($> 85\%$ of visits; Fig. 4b). If perch hopping after the last reward (Fig. 4b) is interpreted as high motivation to feed, then it is paradoxical that birds were more persistent after short than long travel requirements. However, if such non-reinforced hopping is seen as a counting error (i.e. a failure to detect the extinction of the patch by

Table V. Time budgets in experiment 2: mean times (s) spent on activities during feeding periods by birds 0, 1 and 2

After*:	Bird					
	0		1		2	
	S	L	S	L	S	L
Eat + Drink	28	33	18	19	17	20
Explore + Bathe	30	55	4	4	15	14
Preen + Ruffle						
+ Bill wipe	120	11	4	24	35	4
Sing	142	14	0	4	0	1

*Previous travel requirement either short (S) or long (L).

inaccurate counting of the number of rewards already obtained) then it is perhaps unsurprising that more errors occur when the cost of making such errors is lower, in other words when travel time is short.

GENERAL DISCUSSION

Our main result is that starlings track short-term changes in travel time by taking fewer prey per patch visit following shorter travel requirements (experiment 1, Fig. 3). Whilst this is true when patches deplete gradually, when patches yield food at a constant rate up to a maximum, the birds do not track changes in travel time (experiment 2, Fig. 4a). In experiment 1, there is no evidence that any travel requirements prior to the most recent have any influence on the number of prey taken per visit (Table III).

Viewed from the perspective of optimality models, our results present a paradox. In experiment 2, the birds behaved as rate-maximizers by taking the maximum number of prey per visit regardless of travel requirement, whilst in experiment 1 they did not maximize rate (Table I): they tracked changes in travel requirement instead of responding to the average.

In the following discussion, we consider four hypotheses to account for this apparent discrepancy between the two experiments. (1) Starlings in experiment 1 might have used a strategy that is adaptive in their natural environment, but inappropriate in the context of experiment 1. (2) Starlings may respond to interruptions to their normal foraging and behave as if the 'time horizon' is short. (3) The birds may be using only very recent past experience to estimate travel times and hence adjust patch time. (4) Rate maximization may be an inappropriate currency that only fits by chance in experiment 2.

Hypothesis (1): Non-adaptive Behaviour

Experiment 2 and the studies cited in the Introduction show that starlings are able to behave as rate-maximizers, but perhaps they are unable to cope with the environment presented in experiment 1. Suppose the birds are designed to respond as though successive inter-patch travel times are strongly positively correlated, and changes in travel time occur in a time-dependent fashion. As shown in the Introduction and in Table I, adjusting the number of prey taken per visit to changes in travel time would then be the rate-maximizing strategy for experiment 1, and taking the maximum number of prey on each visit would maximize the rate in experiment 2.

This is a plausible hypothesis, but when stated in such a general way it would be difficult to test, other than by showing that, in the starlings' natural environment, changes in travel time always occur in a correlated, time-dependent way. However, as explained below, hypotheses (2) and (3) also invoke the notion that the birds may have behaved in a non-adaptive way in experiment 1.

Hypothesis (2): Interruptions to Foraging

In our experiments the period per day available for foraging was fixed, but the birds spent much of their time on activities not directly related to foraging, so any one foraging bout could be considered

to have an effectively random, finite time horizon. The effect of finite time horizons is examined in Appendix 2 by invoking random interruptions of the foraging period (e.g. the end of the day, or disturbance), after which travel time is unpredictable. Note that mathematically this is equivalent to a model discounting future foraging gain (cf. Kagel et al. 1986; Mazur & Vaughan 1987). The main conclusions from the model (Appendix 2) are that averaging is still optimal when changes in travel time occur with probability 0.5 (as in experiment 1), or when the rate at which the foraging period ends is small. However, when there is positive sequential correlation of travel times and the probability of interruption is relatively large and therefore the environment is unlikely to change before a foraging bout ends, tracking is optimal.

Thus, for finite time horizons to explain the strong dependence of the number of prey taken per visit on previous travel time in experiment 1, we would have to return to a more specific version of hypothesis (1) and postulate that the birds are adapted to respond as though the environment is different from the one experienced in experiment 1. The birds would have to be adapted to environments with correlated travel times with relatively large disturbance rates or at least, in our experiments, perceive travel times to be correlated and disturbances during the day, such as loud noises, as interruptions to their foraging periods.

Hypothesis (3): Use of Recent Experience

The starting point for our experiments was the question of how foragers estimate the parameters of their environment on the basis of previous experience. The results of both experiments 1 and 2 could be explained by suggesting that starlings are rate-maximizers that use an estimate of travel time heavily weighted towards the most recently experienced travel time. In experiment 1 this would lead to tracking the changes in travel time, whilst in experiment 2 it would not, since the rate-maximizing policy is to take the maximum number of prey per visit regardless of travel time.

Although our results show no effect of travel requirements further back than the most recent one, a weak dependence upon experience before the most recent travel requirement cannot be ruled out. The expected run length of short or long travel requirements was only two, so that 'past experience' would tend to average out over very few patch-use

cycles. If the birds used a rule in which a weighted average of the last and earlier travel times was used to predict the future, the 'earlier travel time' component would be a constant, so that the only observable effect would be that of the most recent travel time.

Is a memory rule in which very recent experience is heavily weighted, and which therefore results in tracking of changes in travel time, adaptive? Tracking changes does not result in rate maximization in experiment 1, but it does quite well as a rate maximizing strategy over a wide range of environmental conditions. We investigated this numerically for the model in Appendix 1 with several, widely differing, values of the transition rates between long and short travel times. This showed that tracking the environment to a certain extent (as evidenced in experiment 1, see Fig. 3), although it costs somewhat in efficiency in some cases, does globally better than averaging. Thus, if a forager does not know the rate of transition between travel times, and it is difficult to estimate them accurately, it may be better to track.

In summary, the results of both experiments could be explained by the hypothesis that starlings use a memory strongly weighted towards the most recent travel time. This may be an unavoidable constraint, but it may also be a good overall policy for rate maximization, which happens to do badly in experiment 1.

Hypothesis (4): Alternative Currencies

In the models considered above, rate maximization is used as a currency, but perhaps an alternative currency is more appropriate. For instance, variance in foraging success may also affect decision making (risk-sensitive foraging: Real 1980; Stephens & Charnov 1982). However, Appendix 3 shows that, for random time horizons, variance minimization would also lead to the policy 'average' in the present context.

Another possibility would be a compound currency of which rate maximization is only one component: namely there is a trade-off with another behaviour. This can explain our results only if we assume that for some reason the cost of performing other activities after short travel requirements is smaller than after long ones in experiment 1, but not in experiment 2. The results of experiment 2 indicate that two of the three observed birds perform more other activities after a patch visit following a short travel requirement, perhaps indicating

that their cost is lower, but in this experiment the number of prey taken per visit does not differ between visits following long and short travel requirements. Hence our evidence does not support the idea of a trade-off.

Proximate Versus Functional Models

Our discussion of finite time horizons and of memory for recent travel history parallels other discussions of models of behavioural control. Future discounting (e.g. Kagel et al. 1986; Mazur & Vaughan 1987) is mathematically equivalent to our functional model of random time horizons and would account for the birds' behaviour in the same conditions. Similarly, models in which memory is represented as a linear combination of recent and past experience (linear operator or LINOP: literature review in Kacelnik et al. 1986), would be consistent with observed patch-use if recent experience is given heavy weighting.

However, the aim of our discussion has not simply been to find the best-fitting model, but also to consider why a particular decision rule should be favoured in a particular environment. As we have seen, future discounting would be an appropriate strategy for a rate-maximizer only if successive travel times tend to be similar and interruptions to foraging are frequent. On functional grounds, a LINOP-like rule with heavy weighting to recent experience, which could account for the results of both experiments, might be a robust strategy for a rate-maximizer over the wide range of environmental conditions represented in Table 1.

Conclusions

To summarize, the failure of starlings to maximize gain rate in experiment 1 could be explained in at least two ways. The birds may (1) treat interruptions to foraging (due to switches to other behaviour patterns or to external events such as loud noises) as producing random, finite time horizons or (2) use only very recent experience of travel time in adjusting their patch residence time. This could be an adaptive strategy in many environments but it is not in experiment 1.

Future experiments could investigate these hypotheses either by manipulating the frequency of interruptions to foraging (as e.g. Lucas 1987) or by supplying the birds with alternative information (e.g. by light cues) of future travel times.

APPENDIX 1

Time-dependent Changes

Suppose that the environment changes between two states, L and S, according to a continuous time Markov process (see e.g. Feller 1968). If travelling starts when the state is L, the travel time is T_L , if it starts during state S, the travel time is T_S ($T_L > T_S$). The forager cannot observe directly which state it is in at any moment. It can only infer the state from its travel experiences. Because of the Markovian nature of the environment, all the relevant information about the current state is contained in the length of the most recently experienced travel time. Therefore, the optimal strategy only depends on the last travel time. Let x_L denote the patch residence time after travel time T_L and x_S that after travel time T_S . We now calculate the asymptotic rate of energy intake as a function of x_L and x_S .

Let α be the chance per unit of time of a change to S if the current state is L and let λ be the transition rate of S to L. If the current state is given, the chances on the different states after a finite period of length T can be readily calculated. Thus, if we denote the probability of a change from state i at time t to state j at time $t + T$ as $P_{ij}(T)$, we can note the following

$$\begin{aligned} P_{LL}(T) &= (\lambda + \alpha \exp[-(\alpha + \lambda)T]) / (\alpha + \lambda) \\ P_{LS}(T) &= (1 - \exp[-(\alpha + \lambda)T])\alpha / (\alpha + \lambda) \\ P_{SL}(T) &= (1 - \exp[-(\alpha + \lambda)T])\lambda / (\alpha + \lambda) \\ P_{SS}(T) &= (\alpha + \lambda \exp[-(\alpha + \lambda)T]) / (\alpha + \lambda) \end{aligned}$$

Thus, if the state at the start of travelling is L, the chance that after completing the travel time T_L and the succeeding patch residence time x_L the state of the environment remains L, is equal to

$$P_{LL}(T_L + x_L)$$

Since the state transitions at different moments are independent, the number of travel times in a run of long (or short) travel times has a geometric distribution

$$P[\text{number of successive long travel times} = k] = [P_{LL}(T_L + x_L)]^{k-1} P_{LS}(T_L + x_L)$$

where $k = 1, 2, 3, \dots$. From this it can be inferred that the mean run length of long travel times is

$$1/P_{LS}(T_L + x_L)$$

Similarly, the average run length of short travel times is

$$1/P_{SL}(T_S + x_S)$$

Let $G(x_L)$ and $G(x_S)$ be the expected net energy gain at patch residence times x_L and x_S , respectively. Then the long-term average foraging rate is (by the renewal theorem, see e.g. Feller 1968)

$$r(x_L, x_S) = \frac{G(x_L)/P_{LS}(T_L + x_L) + G(x_S)/P_{SL}(T_S + x_S)}{(T_L + x_L)/P_{LS}(T_L + x_L) + (T_S + x_S)/P_{SL}(T_S + x_S)}$$

The optimal values of x_L and x_S are such that $r(x_L, x_S)$ is maximized. We consider two contrasting cases.

Case 1: slowly changing environment

Suppose $(\alpha + \lambda)T_L \ll 1$. Since $1 - e^{-x} \simeq x$ for small x , in this case

$$P_{LS}(T_L + x_L) \simeq \alpha(T_L + x_L) \text{ and } P_{SL}(T_S + x_S) \simeq \lambda(T_S + x_S)$$

Therefore

$$r \simeq [(1/\alpha)G(x_L)/(T_L + x_L) + (1/\lambda)G(x_S)/(T_S + x_S)] / (1/\lambda + 1/\alpha)$$

Thus, r will be maximized by x_L maximizing $G(x_L)/(T_L + x_L)$ and x_S maximizing $G(x_S)/(T_S + x_S)$, since $G(x)$ is monotonically increasing and concave. Thus, in this case the optimal policy is to track.

Case 2: quickly changing environment

Suppose $(\alpha + \lambda)T_s \gg 1$. Since $e^{-x} \approx 0$ for large x , then $P_{Ls}(T_L + x_L) \approx \alpha/(\alpha + \lambda)$ and $P_{SL}(T_S + x_S) \approx \alpha/(\alpha + \lambda)$. Thus

$$r(x_L, x_S) \approx [G(x_L)/\alpha + G(x_S)/\lambda] / [(T_L + x_L)/\alpha + (T_S + x_S)/\lambda]$$

Thus r is maximized if $G'(x_L) = G'(x_S)$ (where $G'(x)$ is the derivative function), which indicates that at the maximum $x_L = x_S$. Therefore, the optimal policy is now to average.

APPENDIX 2

Animal-dependent Changes and Finite Time Horizon

Assume that the sequence of long and short travel times is generated by a discrete two-state Markov chain (as opposed to the continuous process of Appendix 1) with transition probabilities p and q . A long travel time is followed by another long travel time with probability $1 - p$ and by a short travel time with probability p , while a short travel time is followed by another short travel time with probability $1 - q$ and by a long travel time with probability q . (The situation in the experiments described in this paper corresponds to such an environment with $p = q = 1/2$.) We consider the case in which a foraging period is randomly terminated with constant rate μ , i.e. the period's length is exponentially distributed with expectation $1/\mu$.

First, note that, because of the Markovian environmental change, the future environment for the forager depends only on the length of the most recently experienced travel time, and any additional information of past trajectory does not improve the prediction. Therefore the optimal patch residence time should depend only on whether the most recent travel time is long or short. Let x_L and x_S be the patch times after a long and a short travel time, respectively. Let w_L be the total foraging success until the end of the current foraging period for a forager who just finished a long travel time. Let w_S be a similar quantity for a forager who has just finished a short travel time. We assume that if the period ends during a stay in a patch the forager is allowed to finish its patch visit.

Consider a forager who just finished a long travel time. The expected foraging success w_L is

$$w_L = G(x_L) + (1 - p)\exp[-\mu(x_L + T_L)]w_L + p \exp[-\mu(x_L + T_S)]w_S \tag{2.1}$$

where the first term is the expected gain of foraging in the current patch (compare e.g. Iwasa et al. 1984). The second term is the probability that the next travel time is also a long one, multiplied by the probability that the foraging period does not end before the beginning of the next patch use, and multiplied by the expected foraging success after the beginning of the next patch. The third term accounts for the event that the next travel time is short. The corresponding equation for w_S is

$$w_S = G(x_S) + q \exp[-\mu(x_S + T_L)]w_L + (1 - q)\exp[-\mu(x_S + T_S)]w_S \tag{2.2}$$

Both w_L and w_S are maximized by the same optimal pair of x_L and x_S . This pair of values can be calculated by putting the derivatives of (2.1) and (2.2) with respect to x_S and x_L equal to zero. This gives

$$G'(x_L) = \mu \exp[-\mu x_L] \{ (1 - p)\exp[-\mu T_L]w_L + p \exp[-\mu T_S]w_S \} \tag{2.3a}$$

$$G'(x_S) = \mu \exp[-\mu x_S] \{ (1 - q)\exp[-\mu T_S]w_S + q \exp[-\mu T_L]w_L \} \tag{2.3b}$$

If $p = q = 1/2$, equations (2.3a) and (2.3b) give $G'(x_L)\exp[\mu x_L] = G'(x_S)\exp[\mu x_S]$. Hence $x_L = x_S$, i.e. the forager should use the same patch time after long and short travel times in this case. On the other hand, the pair of linear equations (2.1) and (2.2) can be solved to obtain

$$w_L = \{ (1 - (1 - q)\exp[-\mu(x_S + T_S)])G(x_L) + (p \exp[-\mu(x_L + T_S)])G(x_S) \} / D$$

where $D = \{ (1 - (1 - p)\exp[-\mu(x_L + T_L)])(1 - (1 - q)\exp[-\mu(x_S + T_S)]) - pq \exp[-\mu(x_L + x_S + T_L + T_S)] \}$ (2.4)

The corresponding expression for w_s can be obtained by exchanging p and q and the suffixes S and L. By examining the pairs of x_L and x_s that maximize w_L (and w_s), the following can be proved.

(1) If both p and q are smaller than $1/2$, and if μ is relatively large, then the optimal forager should stay longer after a long travel time ($x_L > x_s$).

(2) When both $[p/(x_L + T_L)]$ and $[q/(x_s + T_s)]$ are smaller than μ (i.e. the sequential correlation is so large that environmental change is unlikely to occur in a foraging period), and when μ is also small (a foraging period includes many patch visits), then the optimal strategy is to use the patch times predicted by the marginal value theorem, based on the separate travel times (i.e. $x_L > x_s$).

(3) If both p and q are $1/2$ (no sequential correlation) and if μ is relatively small (many visits in a single period), then $x_L = x_s$, the forager should use the same patch time, based on the average travel time.

(4) If μ is much smaller than $[p/(x_L + T_L)]$ and $[q/(x_s + T_s)]$ (i.e. many transitions occur during a single foraging period), then the difference between x_L and x_s becomes negligibly small.

APPENDIX 3

Animal-dependent Changes: Variance of Daily Foraging Gain

In the context of risk-sensitive foraging, it is usually assumed that an optimal forager maximizes

$$E[\text{TFS}] - a \text{Var}[\text{TFS}] \tag{3.1}$$

where TFS denotes the total foraging success until the end of the day and a is a constant (Real 1980). We consider the same currency here.

Consider the same model as in Appendix 2, i.e. assume that the environment changes according to a discrete Markov chain and that the length of a foraging period is distributed as a negative exponential with mean $1/\mu$. Here, the foraging period corresponds to the total foraging time in a day, irrespective of interruptions. Since many patch visits occur in a day, the effective foraging period is long. Thus μ is quite small.

Let v_L be the variance of total foraging success (TFS) for a forager who has just finished a long travel time and let v_s be the corresponding quantity for short travel times. First note the following general relation for conditional variance

$$\text{Var}[\text{TFS}] = \sum_i \text{Var}[\text{TFS}|A_i]P(A_i) + \sum_i E[\text{TFS}|A_i]^2 P(A_i) - E[\text{TFS}]^2 \tag{3.2}$$

where A_1 , A_2 and A_3 are exclusive events, $P(A_i)$ is the probability of event A_i , and $E[\text{TFS}|A_i]$ and $\text{Var}[\text{TFS}|A_i]$ are the conditional mean and variance of TFS, respectively. For a forager who has just finished a long travel time, one of the following three events will occur.

$A_1 = \{\text{the next travel time is long, and the period does not end before the next patch encounter}\}$

$A_2 = \{\text{the next travel time is short, and the period does not end before the next patch encounter}\}$

$A_3 = \{\text{the period ends before the next patch encounter}\}$

For these events, we can calculate the probability of each type of event and the conditional mean and variance of total foraging success. Thus

$$P(A_1) = (1-p)\exp[-\mu(x_L + T_L)], \quad E[\text{TFS}|A_1] = G(x_L) + w_L, \quad \text{Var}[\text{TFS}|A_1] = V_G(x_L) + v_L$$

$$P(A_2) = p \exp[-\mu(x_L + T_s)], \quad E[\text{TFS}|A_2] = G(x'_L) + w_s, \quad \text{Var}[\text{TFS}|A_2] = V_G(x_L) + v_s$$

$$P(A_3) = 1 - P(A_1) - P(A_2), \quad E[\text{TFS}|A_3] = G(x_L), \quad \text{Var}[\text{TFS}|A_3] = V_G(x_L)$$

$V_G(x_L)$ denotes the variance in net energy gain during one patch visit of length x_L . Equation (3.2) gives

$$\begin{aligned} v_L = & V_G(x_L) + G(x_L)^2 - G(x_L) \\ & + (1-p)\exp[-\mu(x_L + T_L)]\{v_L + w_L^2 + [2G(x_L) - 1]w_L\} \\ & + p \exp[-\mu(x_L + T_s)]\{v_L + w_s^2 + [2G(x_L) - 1]w_s\} \end{aligned} \tag{3.3}$$

The equation for v_s can be derived similarly. Expressions for V_L and V_S can be derived from these equations. Since $1/\mu$ is large, we can pick up leading terms only. By noting that w_L and w_S given by (2.4) are also large (of the order of $1/\mu$), it can be derived that

$$v_L = c(w_L^2 + w_S^2)/\mu + O(1/\mu^2) \quad (3.4)$$

where $O(1/\mu^2)$ indicates the terms of the order of $1/\mu^2$ and c is a constant. Therefore, maximizing $w_L - av_L$, is in fact the same as maximizing $w_L - aw_L^2$, which corresponds to maximizing w_L when $w_L < 1/(2a)$. In this case, we expect the relevance of variance, and hence a , to be small. Therefore, the maximization of (3.1) gives the same result as the maximization of the average foraging success if μ is large.

ACKNOWLEDGMENTS

The work was supported by a grant from NERC to J.R.K. (GR4503'A'). Further support was provided by Brasenose College (I.C.C.), King's College Research Centre (A.K.), the Royal Society (P.H.), the Sasakawa fund (Y.I.) and the British Council (P.H. and Y.I.). We thank Julian Howe and Alan Woodington for technical assistance and Daniela Brunner, Evert Meelis, Alan Grafen, Alasdair Houston, Rudi Drent and John McNamara, for discussion. This manuscript was greatly improved by comments from Michel Treisman, Evert Meelis and an anonymous referee.

REFERENCES

- Armitage, P. & Berry, G. 1987. *Statistical Methods in Medical Research*. 2nd edn. Oxford: Blackwell Scientific Publications.
- Charnov, E. L. 1976. Optimal foraging, the marginal value theorem. *Theor. Popul. Biol.*, **9**, 129–136.
- Cowie, R. J. 1977. Optimal foraging in great tits (*Parus major*). *Nature, Lond.*, **268**, 137–139.
- Cox, D. R. & Hinkley, D. V. 1974. *Theoretical Statistics*. London: Chapman & Hall.
- Cuthill, I. C. 1985. Experimental studies of optimal foraging theory. D. Phil. thesis, Oxford University.
- Dow, S. M. & Lea, S. E. G. 1986. Foraging in a changing environment: simulations in the operant laboratory. In: *Quantitative Analyses of Behavior. Vol. 6: Foraging* (Ed. by M. L. Commons, A. Kacelnik & S. J. Shettleworth), pp. 89–113. Hillsdale, New Jersey: Lawrence Erlbaum.
- Feller, W. 1968. *An Introduction to Probability Theory and its Applications. Vol. 1*. 3rd edn. New York: John Wiley.
- Haccou, P., Dienske, H. & Meelis, E. 1983. Analysis of time-inhomogeneity in Markov chains applied to mother–infant interactions of rhesus monkeys. *Anim. Behav.*, **31**, 927–945.
- Hursh, S. R. 1980. Economic concepts for the analysis of behavior. *J. exp. Anal. Behav.*, **34**, 219–238.
- Iwasa, Y., Suzuki, Y. & Matsuda, H. 1984. Theory of oviposition strategy of parasitoids. 1. Effect of mortality and limited egg number. *Theor. Popul. Biol.*, **26**, 205–227.
- Kacelnik, A. 1984. Central place foraging in starlings (*Sturnus vulgaris*). I. Patch residence time. *J. Anim. Ecol.*, **53**, 283–299.
- Kacelnik, A. & Cuthill, I. C. 1987. Starlings and optimal foraging theory: modelling in a fractal world. In: *Foraging Behavior* (Ed. by A. C. Kamil, J. R. Krebs & H. R. Pulliam), pp. 303–333. New York: Plenum Press.
- Kacelnik, A., Krebs, J. R. & Ens, B. 1986. Foraging in a changing environment: an experiment with starlings (*Sturnus vulgaris*). In: *Quantitative Analyses of Behavior. Vol. 6: Foraging* (Ed. by M. L. Commons, A. Kacelnik & S. J. Shettleworth), pp. 63–87. Hillsdale, New Jersey: Lawrence Erlbaum.
- Kagel, J. H., Green, L. & Caraco, T. 1986. When foragers discount the future: constraint or adaptation? *Anim. Behav.*, **34**, 271–283.
- Lucas, J. R. 1987. Foraging time constraints and diet choice. In: *Foraging Behavior* (Ed. by A. C. Kamil, J. R. Krebs & H. R. Pulliam), pp. 239–269. New York: Plenum Press.
- McNamara, J. M. & Houston, A. I. 1987. Memory and the efficient use of information. *J. theor. Biol.*, **125**, 385–396.
- Mazur, J. E. & Vaughan, W. 1987. Molar optimization versus delayed reinforcement as explanations of choice between fixed-ratio and progressive-ratio schedules. *J. exp. Anal. Behav.*, **48**, 251–261.
- Parker, G. A. & Stewart, R. A. 1976. Animal behavior as a strategy optimizer: evolution of resource assessment strategies and optimal emigration thresholds. *Am. Nat.*, **110**, 1055–1076.
- Real, L. A. 1980. Fitness, uncertainty, and the role of diversification in evolution and behavior. *Am. Nat.*, **115**, 623–638.
- Scheffé, H. 1959. *The Analysis of Variance*. New York: John Wiley.
- Stephens, D. W. & Charnov, E. L. 1982. Optimal foraging: some simple stochastic models. *Behav. Ecol. Sociobiol.*, **10**, 251–263.
- Stephens, D. W. & Krebs, J. R. 1986. *Foraging Theory*. Princeton: Princeton University Press.
- Tinbergen, J. M. 1981. Foraging decisions in starlings (*Sturnus vulgaris* L.). *Ardea*, **69**, 1–69.

(Received 29 June 1987; initial acceptance

14 September 1987; final acceptance 19 August 1989;

MS. number: 3047)