

Rate currencies and the foraging starling: the fallacy of the averages revisited

Melissa Bateson and Alex Kacelnik

Department of Zoology, University of Oxford, South Parks Road, Oxford OX1 3PS, UK

In classical optimal foraging models long-term rate of energy intake (the ratio of expected amount of food over expected time) is assumed to be the maximized currency, because doing this is consistent with minimizing the loss of alternative opportunities. Here this possibility and various alternatives are examined quantitatively using European starlings (*Sturnus vulgaris*) in the laboratory. The birds chose between two cues. One signaled an option that led to either a fixed delay to food ("single standard," Experiment 1) or to one of two equally probable delays to food ("double standard," Experiments 2 and 3). The other cue signaled an "adjusting option" consisting of a single delay to food. This option adjusted according to the previous choices made by the birds, improving when the standard had been preferred and worsening when the adjusting option had been preferred. Adjustments were made either by changing the delay to food or the amount of food. The rationale underlying this procedure was that the parameter values at which the adjusting option stabilizes should reflect the subjective value of the standard. This was validated in Experiment 1. In Experiments 2 and 3 the adjusting option fluctuated around parameter values that are interpreted as yielding subjective equivalents of the double standard. The results contradict the predictions based on minimizing the lost opportunity. First, the birds did not include all the time intervals in their assignment of value to the two options, and second, the birds used the expected ratio of amount over time rather than the ratio of expected amount over expected time as their rate currency. *Key words*: currency, European starling, optimal foraging, rate maximization. [*Behav Ecol* 7:341–352 (1996)]

Optimality models are built on the premise that the pay-offs of different behavioral decisions can be expressed in a common "currency." Although the ultimate currency of all decisions must be Darwinian fitness, it is often assumed that organisms may use a proximate surrogate for fitness to guide their behavior in the short-term. This assumption is necessary to generate predictions that are testable using short-term behavioral observations. In contrast to the descriptive models often used in psychology and economics, in which observed behavior is used to specify a currency (or utility) *a posteriori*, in evolutionarily inspired foraging models the currency is postulated on the basis of reasoning about the fitness consequences of foraging. The currency assumed in the majority of existing foraging models is average net long-term rate of energy intake (for a review of these models see Stephens and Krebs, 1986). This hypothesis has been justified on the grounds that it combines three dimensions of importance (gains, losses, and time) in a sensible and simple manner. In this article we test this currency hypothesis against potential alternatives using European starlings (*Sturnus vulgaris*) foraging in the laboratory.

The average net long-term rate of energy intake for a foraging option is best estimated by the ratio of the expected net energy gained (G) over the expected time spent foraging (T). Thus

$$\text{long-term rate, or ratio of expectations (RoE)} = \frac{\sum_{i=1}^n G_i}{\sum_{i=1}^n T_i} \quad (1)$$

where G_i is the energy gained from the i th food item, T_i is the time taken to acquire this item, and n is the total number of food items available in the option. The choice of long-term

rate (or RoE) is justified on the grounds that, all else being equal, natural selection should favor animals that on average have a higher net energy gain during the time dedicated to foraging. Both the survival and reproduction components of fitness have an obvious relation to energy and time: the more food energy, the more energy available for reproduction, and the less time spent feeding, the more time available for other activities, such as predator avoidance and reproductive behavior. This is encapsulated in the principle of lost opportunity (Stephens and Krebs, 1986) that identifies optimal choices by comparing the gains accrued from each possible choice with those that would accrue if the same time were to be devoted to an alternative.

Despite these persuasive arguments for long-term rate, an alternative rate currency for combining G and T has periodically been advocated, starting with an article entitled, "The fallacy of the averages in ecological optimization theory" (Templeton and Lawlor, 1981). The currency in question is the expected ratio of G over T , where, using the same definitions as above,

$$\text{expectation of Ratios (or EoR)} = \frac{\sum_{i=1}^n \left(\frac{G_i}{T_i} \right)}{n} \quad (2)$$

The expectation of the ratios (EoR) has been referred to by a number of names including few-prey-item maximization (Turelli et al., 1982), per-patch rate (Stephens and Krebs, 1986), and short-term rate (Real, 1991). Since each of these names may be questioned on different grounds, we prefer the more descriptive EoR. The crucial difference between long-term rate and EoR is that in the latter the rates at which individual prey items are obtained are not weighted according to the time over which this rate of intake is experienced. Thus, if there is no variability in T , then long-term rate and EoR give identical values. However, when T is variable, for any set of values of G and T , it will be true that EoR is not equal to long-term rate, and that for two options of equal long-term rate the one with the greater variance in T will have the great-

Received 23 January 1995; revised 3 July 1995; accepted 30 October 1995.

1045-2249/96/\$5.00 © 1996 International Society for Behavioral Ecology

er EoR (Caraco et al., 1992). Hence the so-called "fallacy of the averages," whose proponents point out that the average of a function of random variables is not equal to the function of the averages, and that foragers ought to maximize EoR rather than long-term rate as assumed in classical OFT (Templeton and Lawlor, 1981). This proposal for the use of EoR was originally rejected on theoretical grounds because of the lack of an obvious link between maximization of this currency and fitness ("the fallacy of the fallacy of the averages" (Gilliam et al., 1982; Possingham, et al., 1990; Turelli et al., 1982). In essence, EoR violates the principle of lost opportunity.

In the psychological literature, however, currencies similar or identical to EoR have received support from a number of empirical studies (e.g., Mazur, 1984, 1986; McDiarmid and Rilling, 1965). For example, Mazur (1984, 1986) has found that pigeons choosing between a fixed and a variable delay to reward behave as would be predicted if they attribute value to the options according to the following equation:

$$\text{Value} = \frac{\sum_{i=1}^n \left(\frac{A_i}{1 + KD_i} \right)}{n} \quad (3)$$

where A is the immediate value of the food reward and D is the delay to receive it in seconds. There are three differences between the above expression and EoR. The first is the constant 1 added to the delay. This prevents the value becoming infinite when the delay is zero, although in practice this is unlikely to be a problem because a food item can never really take no time to obtain. The second is the constant K , a free parameter with units s^{-1} that varies between individuals, but was found by Mazur to be close to unity in pigeons. The final difference is that the D in Mazur's rule is not exactly equivalent to the T in EoR because it includes only the delay between the decision and reward and not other time intervals in the foraging cycle such as the intertrial interval.

In the foraging literature theoretical debate related to the use EoR has been revitalized by the emergence of experimental data on starlings (Bateson and Kacelnik, 1995b; Cuthill et al., 1990), pigeons (*Columba livia*; Todd and Kacelnik, 1993), and bumble bees (*Bombus pennsylvanicus*, Harder and Real, 1987; Real et al., 1990) that are compatible with maximization of EoR. The evidence comes from foraging experiments involving variance in T , in which the predictions of long-term rate and EoR maximization are different. In some such experiments the discrepancy between the observed behavior and the predictions of long-term rate maximizing can be explained if EoR is in fact the currency under maximization. For example, when given a choice of five units of food after a 20 s delay with certainty (the fixed option) versus a delay that was either 2.5 s or 60.5 s with equal probability (the variable option), starlings strongly preferred the variable option, despite the fact that it gave the lower long-term rate of energy intake (Bateson and Kacelnik, 1995b). We argued that this result could be compatible with maximization of EoR, since when valued in this currency the variable option offers a greater subjective value than the fixed option. However, this study does not represent a strong test of the EoR currency because other theories predict the same qualitative preference. For example, both risk-sensitive foraging theory under negative energy budgets and time discounting due to unpredictable interruptions to foraging also predict a preference for the variable option (McNamara and Houston, 1987), as does the Scalar Expectancy Theory (Bateson and Kacelnik, 1995b; Gibbon et al., 1988; Reboresda and Kacelnik, 1991).

In this article we use a design for which quantitative predictions for a variety of rate currencies can be made. Our rationale is based on the experiment described above (Bate-

son and Kacelnik, 1995b), in which a variable delay to food was preferred over a fixed delay. We reasoned that if the subjective value of the fixed option were increased, a point would be reached at which the two options would become subjectively equivalent in value. The fixed option could be improved either by increasing the amount of food delivered or by reducing the delay to receive food. The predicted points of subjective equality depend on the currency in which the options are valued, and they can therefore be used to determine the currency used by the birds.

Points of subjective equality between the two options can be determined experimentally by establishing when subjects treat a fixed option as being equal in value to a variable option. We employ a procedure in which the fixed option is periodically modified in response to the birds' choices until it is chosen in an average of 50% of trials. This procedure involves presenting an animal with the two foraging options on response keys of different colors. The schedule programmed for one option, the "standard option," does not alter. The other, the "adjusting option," has a schedule that is adjusted according to the previous preference of the animal: if the animal prefers the standard then the value of one parameter of the adjusting option (the "adjusting parameter," either the delay to reward or the amount of reward) is modified so that the subjective value of the adjusting option is increased, and if the animal prefers the adjusting option the opposite adjustment takes place. An increase in subjective value is achieved by a shorter delay to reward or a larger reward, and a decrease means the opposite in each case. After a few sessions the adjusting parameter is generally found to fluctuate around a stable value, which is interpreted as the value of this parameter that makes the adjusting option equivalent in subjective value to the standard.

GENERAL METHODS

Subjects

The subjects were six wild-caught, first-year European starlings (*Sturnus vulgaris*), housed individually in cages measuring 120 cm × 50 cm × 60 cm. Two months previously they had served in the choice experiment described by Bateson and Kacelnik (1995). The birds were visually but not acoustically isolated. Temperature in the laboratory ranged between 7°C and 13°C, and the lights were on between 0600 h and 1800 h. The birds were food deprived from 1700 h until the start of each session at 0800 h the following morning. During the experimental sessions the birds were rewarded with turkey starter crumbs, and after the session the birds were given four mealworms and *ad libitum* turkey crumbs until 1700 h. This regime resulted in the birds being maintained at approximately 90% of their *ad libitum* feeding weights.

Apparatus

Each cage had an operant panel in the center of the back wall with two response keys (3.5 cm in diameter) at 6 cm on either side of a central food hopper. The food hopper was connected to a pellet dispenser (Campden Instruments) filled with turkey crumbs sieved to an even size. One unit of food averaged 0.012 g of turkey crumbs and took 1 s to deliver. A BBC Master microcomputer running SPIDER experimental control language (Paul Fray, Ltd.) controlled the stimulus events and response contingencies, as well as the recording of data. Since the birds were already familiar with the apparatus and had been trained to peck illuminated keys to obtain food, no further training was required.

Table 1
Schematic diagram showing the delay (indicated by the length of the horizontal line) and the amount of food (indicated by the diameter of the circle) available in the standard and adjusting options in each comparison made (not to scale)

Experiment	Treatment (=adjusted parameter)	Standard option	Adjusting option
1	Amount	—20—(5)	—20—(7)↑
1	Delay	—20—(5)	←7—(5)
2	Amount	P=0.5 { —2.5—(5) 60.5—(5) }	—20—(7)↑
2	Delay	P=0.5 { —2.5—(5) 60.5—(5) }	←7—(5)
3	Amount	P=0.5 { —5—(2) 20—(2) }	—20—(7)↑

The arrows indicate which parameter was titrated in the adjusting option. The question marks indicate the point subjective equality sought in each comparison. In the double standards (Experiments 2 and 3) the two delays to reward occurred with equal probability (indicated by "P=0.5").

Adjusting procedure

Colored lights on the pecking keys were used as the discriminative stimuli indicating the two options. For each bird, either the standard was cued by a green light and the adjusting option by a red light or vice versa, such that the assignment of colors to options was balanced across birds. Depending on the experiment, the standard was either a "single standard" (Experiment 1), that consisted of a fixed delay to obtain a fixed amount of food, or a "double standard" (Experiments 2 and 3) that consisted of two different delays programmed to occur with equal probability to obtain a fixed amount of food. The values of the delays and amounts in the two options of each experiment are shown schematically in Table 1.

For all of the experiments we used the same basic procedure consisting of discrete trials with a fixed intertrial interval (ITI) of 40 s. There were two types of trials, forced trials and choice trials. The forced trials allowed the birds to experience the two options and learn about their characteristics, while the choice trials tested the birds' preferences. A forced trial (see Figure 1) began with one of the key lights flashing (on for 0.7 s and off for 0.3 s). When the bird pecked the key the light changed from flashing to being continuously illuminated and a delay to receive the associated reward began. The first peck after this delay had timed out extinguished the key light and caused the delivery of food to the hopper. The ITI started to time when food delivery was complete. The choice trials were identical to the forced trials with the exception that a choice trial began with both keys flashing, one in each color, and as soon as the bird pecked one of the keys it was illuminated continuously and the other was extinguished (see Figure 1). Thus a bird committed itself to one of the options with its first peck.

An experimental session was composed of 12 blocks of 10 trials. Each block consisted of eight forced trials followed by two choice trials. The forced trials comprised four of each option given in a different randomly chosen order for each

block, with half of the trials of each type appearing on the left key and the other half on the right (see Figure 2). This procedure was designed to ensure that the birds had experienced both options an equal number of times before having to choose between them, and that the adjusting option was experienced as invariable in the short term. The birds were given one session per day. When all of the birds had completed the session, generally at around noon, they were given *ad libitum* food until 1700 h when they were deprived for the night.

In each experiment the adjusting parameter started at some extreme value (see methods of individual experiments for details) and the birds were trained until they completed one session with at least a 90% preference in the choice trials for one of the options. When a bird reached this criterion titration began. After each pair of choice trials the schedule for the subsequent block was altered as follows. If the bird chose the standard twice then the value of the adjusting option became one step better (either 1 s shorter delay or 1 unit more food depending on the treatment), if a bird chose the adjusting option twice its value became one step worse (either 1 s longer delay or 1 unit less food depending on the treatment). If a bird chose each option once no change followed. Each new session began with the adjusting option set at the value at which it finished in the previous session. The birds were run 7 days a week, and titration continued for a minimum of 22 sessions (264 blocks).

EXPERIMENT 1

Rationale and predictions

This experiment was designed to test whether the adjusting procedure produced valid subjective equivalents in starlings. The strategy we adopted was to perform two titrations, one with amount and one with delay as the adjusting parameter, in which the point of subjective equality could be unequivocally predicted. In both treatments we titrated against a single standard consisting of a delay of 20 s to obtain five units of food. In the amount treatment the adjusting option consisted of a 20 s delay to obtain an adjusting quantity of food, and in the delay treatment the adjusting option consisted of an adjusting delay to obtain five units of food (see Table 1). Thus, if the process of adjustment did not bias choice, the subjective equivalents derived from the adjusting parameter should be equal to the value of the same parameter in the standard option. In the amount treatment the mean value of the adjusting amount should be five units, and in the delay treatment the mean value of the adjusting delay should be equal to 20 s.

Methods

All six birds did both treatments; Birds 0, 1, and 2 did the titration with amount followed by the titration with delay, and Birds 3, 4 and 5 did the treatments in the reverse order. The adjusting parameter started out as 15 units of food in the amount treatment and as 5 s delay in the delay treatment. All other details are as described for a titration against a single standard in the General Methods section above.

Results and discussion

Statistics on the value of the adjusting parameter in the last 100 blocks of the amount and delay treatments are shown in Figures 3 and 4 respectively. In both treatments the value of the adjusting parameter fluctuated symmetrically (i.e., the




Figure 1
The events in a forced trial (left) and a choice trial (right). The color of the lines indicates the state of the pecking keys, with white indicating that there is no light illuminated on the pecking key.

mean and median of the data were similar) around the value of the same parameter in the standard. In the amount treatment the mean value of the adjusting amount across birds is 5.22 units, and the mean values obtained from each bird are not significantly different from the predicted value of five units (Wilcoxon one-sample test against a median of 5, $T = 13.0$, $p = .675$, two-tailed). Similarly, in the delay treatment the mean value of the adjusting delay across birds is 21.77 s, and the mean values obtained from each bird are not significantly different from the predicted delay of 20 s (Wilcoxon one-sample test against a median of 20, $T = 14.0$, $p = .529$, two-tailed). These results suggest that the adjusting procedure does not introduce a significant bias and thus it is a valid method for establishing points of subjective equality.

EXPERIMENT 2

Rationale and predictions

Having established that the adjusting procedure itself did not introduce any strong bias, we used it to find the value of the adjusting option subjectively equivalent to a double standard. Again there were two treatments. In both treatments the double standard consisted of a delay to obtain five units of food that was delayed either 2.5 s or 60.5 s with equal probability. In the amount treatment the adjusting option consisted of a 20 s delay to obtain an adjusting amount of food, while in the delay treatment it consisted of an adjusting delay to obtain five units of food. In both treatments the points of subjective equality predicted by long-term rate and EoR were calculated as described in the Appendix and are shown in Table 2. As well as the basic predictions, we show the predictions that are made if the birds do not include all of the time intervals in

the trial in their rate calculations, because although according to the logic of lost opportunity all time intervals of a foraging cycle (searching, chasing, handling etc.) have equal impact on future intake, they may not be equal in their subjective impact. We shall return to this issue in the discussion.

Methods

All six birds did both treatments; Birds 0–2 did the delay titration followed by the amount titration, and Birds 3–5 did the titrations in the reverse order. In the amount treatment the adjusting amount began at five units of food and in the delay treatment the adjusting delay began at 20 s. In all other respects the methods are as described in the General Methods section for a titration against a double standard.

When the titration in the delay treatment had been completed, the adjusting delay for each bird was set at the mean of the values that it had taken in the last 100 blocks of the titration (see Figure 5), and the birds were given six sessions consisting entirely of choice trials. The value of the adjusting delay was not altered during these sessions. Thus, data were collected on 360 choices for each bird with the exception of Bird 2 that failed to complete the last 13 trials of the second session.

Results and discussion

Statistics on the value of the adjusting delay in the last 100 blocks of the delay treatment are shown in Figure 5. The value of the adjusting delay dropped from its initial value of 20 s to fluctuate symmetrically around a mean value of 5.61 s. The mean values for each bird are not significantly different from

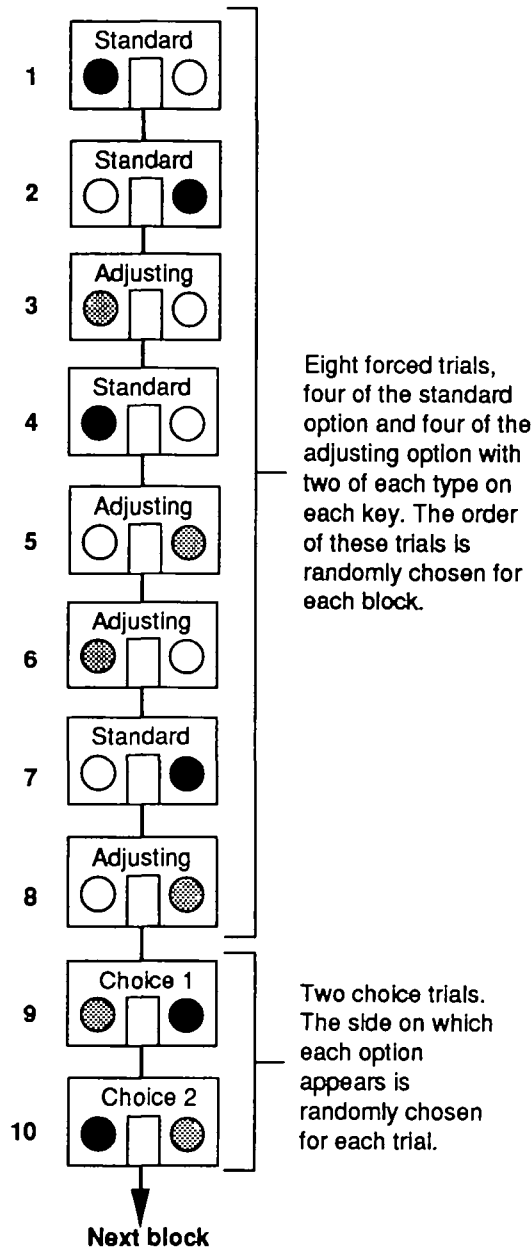


Figure 2
The sequence of trials in a block. The icon represents the panel with two pecking keys and a central food hopper. The color of the pecking keys indicates the type of trial. White pecking keys are not illuminated and the other two colors represent the cues for the standard and adjusting options.

the prediction made by EoR calculated with the delay time only (Wilcoxon one-sample test against a median of 4.80, $T = 18.0$, $p = 0.142$, two-tailed), but are significantly different from the predictions made by all of the other EoR and long-term rate currencies under consideration (Wilcoxon one-sample test against a median of 8.46 or greater, $T = 0$, $p = .036$, two-tailed).

In the final six sessions of choice trials the birds showed no consistent preference for either option. Birds 0–3 showed no significant preference, Bird 4 had a significant preference for the adjusting option, and Bird 5 had a significant preference for the double standard (binomial tests with a null hypothesis that the two options have an equal probability of being cho-




Figure 3
Results from the amount treatment in Experiment 1. The bars show the mean + 1 standard deviation of the value of the adjusting delay in the last 100 blocks of the treatment for each bird; the stars show the median of the data. The dashed line shows the amount in the standard option.

sen). The proportion of times each bird chose the double standard are as follows: Bird 0: 0.52; Bird 1: 0.48; Bird 2: 0.47; Bird 3: 0.53; Bird 4: 0.42; and Bird 5: 0.58. Thus, even the two birds that had significant preferences were not far from the random expectation of 0.5. These results support our inter-




Figure 4
Results from the delay treatment in Experiment 1. The bars show the mean + 1 standard deviation of the value of the adjusting delay in the last 100 blocks of the treatment for each bird; the stars show the median of the data. The dashed line shows the delay in the standard option.

Table 2
Points of subjective equality predicted in the amount and delay treatments of Experiment 2

Currency	Time included	Value of adjusting amount in amount treatment (units of food)	Value of adjusting delay in delay treatment (s)
Long-term rate (RoE)	All time	4.21	31.5
	Delay and feeding time only	3.17	31.5
	Delay only	3.17	31.5
EoR	All time	4.95	20.65
	Delay and feeding time only	11.82	8.46
	Delay only	20.83	4.80 (harmonic mean)

Predictions for a range of different currencies are shown. For the calculation of these predictions see the Appendix.

pretation that the adjusting procedure produces valid points of subjective equality between a fixed delay to food and a double standard. To summarize, the results from the delay treatment suggest that (all else being equal) starlings value foraging options using EoR rather than long-term rate, and that they do not include all of the time in the trial in their rate calculations.

In the amount treatment, four out of six birds continued to prefer the double standard independently of the amount of food available in the adjusting option, with the consequence that the value of the adjusting amount continued to increase. Eventually we removed these birds from the experiment before they had completed the full 264 blocks. No analyses were performed on the data from the amount treatment. There are various possible explanations for the failure of the

titrations to stabilize. A theoretical possibility is that the starlings learned to "cheat" the procedure by continuing to choose the less attractive option in the choice trials to increase the amount of food they received in the forced trials. However, this seems unlikely because this requires sacrificing immediate payoff in choice trials in exchange for delayed gains later on, an unlikely behavior given what is known about animal choice. Furthermore, in the delay treatment and in Experiment 1 the same birds did not learn to do this. An alternative interpretation takes into account the expected value at equilibrium. For the currency suggested by the results from the delay treatment, the predicted subjective equivalent in this treatment is reached when the adjusting option delivers approximately 21 units of food (see Table 2). The failure of the amount titrations to stabilize may have occurred because this amount of food may be too large for the starlings to eat at once, and therefore it may be systematically underestimated. If this interpretation is correct the effect should disappear when the predicted equilibrium value is smaller. We examined this possibility in the following experiment.

EXPERIMENT 3

Rationale and predictions

In this experiment the values for the delays and amounts in the double standard were such that the predicted points of subjective equality for all of the currencies lay within the bounds of what we knew the starlings could measure (i.e., about five units of food or less). The double standard consisted of a delay to obtain two units of food that was either 5 s or 20 s with equal probability. The adjusting option consisted of a 20 s delay to obtain an adjusting amount of food (see Table 1). The predicted points of subjective equality derived from the different currencies are given in Table 3. Note that the currency supported by the results of Experiment 2 (EoR with delay only) predicts that the point of subjective equality should occur when the value of the adjusting amount is exactly five units of food.

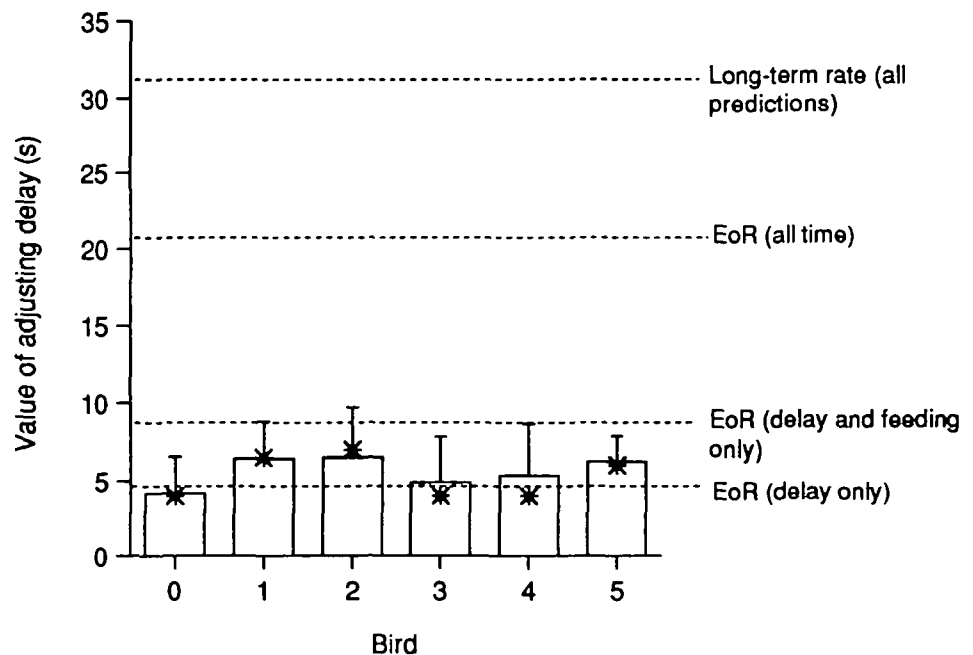


Figure 5
Results from the delay treatment in Experiment 2. The bars show the mean + 1 standard deviation of the value of the adjusting delay in the last 100 blocks of the treatment for each bird; the stars show the median of the data. The dashed lines show the predictions of the various currencies under consideration. The data are significantly different from all of these predictions with the exception of expectation of ratios (EoR) calculated without the ITI or latency.

Table 3
Predicted points of subjective equality for the different currencies under consideration in Experiment 3

Currency	Time included	Value of adjusting amount (units of food)
Long term rate (RoE)	All time	2.28
	Delay and feeding time only	3.20
	Delay only	3.20
EoR	All time	2.32
	Delay and feeding time only	4.64
	Delay only	5.00

For the calculation of these predictions see the Appendix.

Methods

The adjusting option started with two units of food. In all other respects the methods are as described in the general methods for a titration against a double standard.

Results and discussion

The amount of food in the adjusting option now stabilized for all birds. Statistics on the value of the adjusting amount in the last 100 blocks are shown in Figure 6. The value of the adjusting amount rose from its initial value of two units to fluctuate symmetrically around a mean value of 4.25 units. The points of subjective equality are not significantly different from the predictions of EoR with delay (Wilcoxon one-sample test against a median of 5.00, $T = 1.00$, $p = .059$, two-tailed) or with delay and feeding time (Wilcoxon one-sample test against a median of 4.64, $T = 4.00$, $p = .208$, two-tailed), but are significantly different from the other predictions of EoR and long-term rate (Wilcoxon one-sample test against a median of 2.32 or less, $T = 21.0$, $p = .036$, two-tailed). This result supports the view that starlings value foraging options using

EoR rather than long-term rate, and that they do not include all of the time in the trial in their rate calculations.

GENERAL DISCUSSION

Our aim in this article was to elucidate the rate currency that starlings use to choose between feeding options that differ in the time taken to acquire food items and their energetic content. We discussed two alternative ways of computing a rate: long-term rate (the ratio of the expected amount over expected foraging time, or RoE), which is the most frequently used currency in optimal foraging models, and the expected ratio of amount over foraging time (or EoR), a currency without *a priori* evolutionary rationale, because it violates the principle of lost opportunity, but consistent with the foraging decisions made by animals in other experimental situations. These currencies make different predictions about the value a forager should assign to an option if there is variability in the time taken to acquire and consume food items. Results from experiments designed to test these predictions led to two main conclusions, both of which are at odds with the principle of lost opportunity: first, the starlings used the expected ratio of amount over time (EoR) rather than long-term rate (RoE), and second, they did not include all of the time intervals in the trial in their rate calculations.

Is the adjusting procedure a valid method?

We shall begin by discussing the validity of the procedure we used to obtain the above results. It is clear from Figures 3 to 6 that in all of the experiments the adjusting parameter fluctuated considerably around its mean value. We have analyzed the causes of this variation elsewhere (Bateson and Kacelnik, 1995a). However, a question relevant to our claims in this article is whether, and if so how, this variation affects the theoretical predictions we have made for the different currencies. Our predictions are based on the assumption that the adjusting option has the same value as a food source yielding a fixed delay to obtain a fixed amount of food (see the Appendix); but given that the adjusting parameter shows considerable variation is this justified? We have three lines of argument that

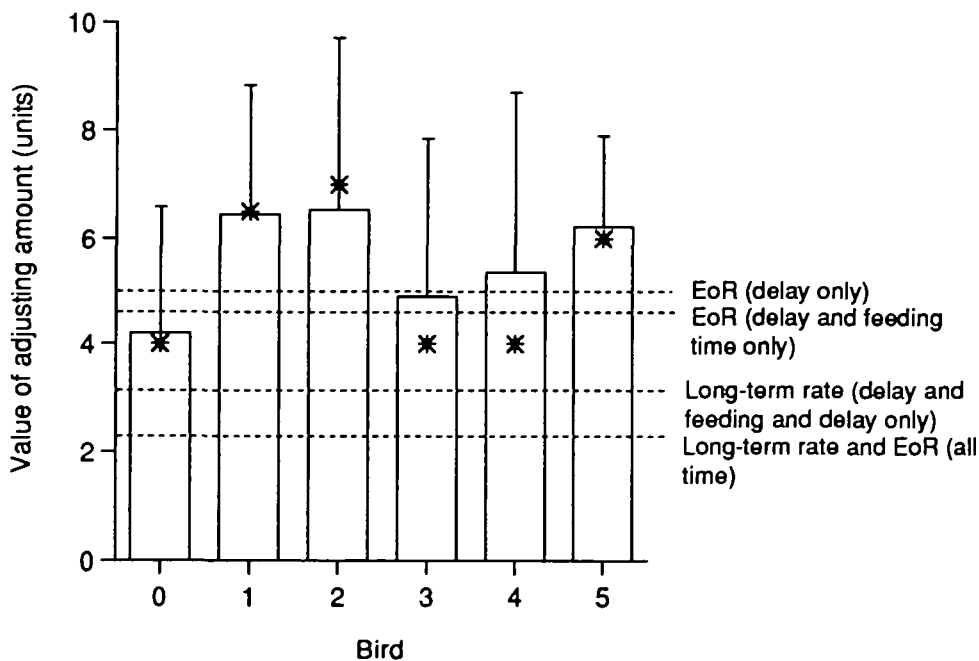


Figure 6
 Results from Experiment 3. The bars show the mean + 1 standard deviation of the value of the adjusting amount in the last 100 blocks of the treatment for each bird; the stars show the median of the data. The dashed lines show the predictions of the various currencies under consideration. The observed data are significantly different from all of these predictions with the exception of EoR calculated without the ITI or latency.

support our interpretation. First, changes in the adjusting option occurred in small steps, with each choice trial being preceded by a minimum of four experiences of the new value of the adjusting parameter. In another experiment, the behavior of starlings that experienced a sudden large transition in the length of the travel time to a patch reached a new asymptote in behavior after about six experiences of the new conditions (Cuthill et al., 1994). If this result is generalized to our experiment, then at the time of each choice trial a bird's estimate of the value of the adjusting parameter should be close to the current value of this parameter. Second, the results from Experiment 1 show that the adjusting procedure produces the equilibrium parameters that would be expected if the birds perceive the adjusting option as fixed. Finally, at the end of the delay treatment of Experiment 2 we showed that the birds remained indifferent between the two options when the adjusting parameter was fixed at the point of subjective equality derived from the adjusting procedure. Thus, we believe there is sufficient evidence that the adjusting procedure produces valid points of subjective equality between a fixed and a variable option.

Why don't the birds include all of the time intervals in their calculations?

Our result that the currency that the starling's behavior maximizes may not include the ITI or latency to peck is contrary to the predictions of the principle of lost opportunity. When a rate-maximizing forager commits itself to pursuing or consuming a given food item it trades off the average time that will be devoted to that food item against the value of the alternative foraging opportunities lost during this period. A measure of the opportunity lost when a particular foraging option is pursued is obtained by multiplying all of the time dedicated to the option by the average rate of gain in the environment as a whole. Therefore, a rate-maximizing forager should include the ITI in its calculations of the value of an option because this interval must be endured before another decision can be made. The lack of effect of the ITI in our experiments may seem surprising because in tests of patch models (such as the Marginal Value Theorem) the time interval between encountering patches affects the giving up time in a patch, as is predicted by the principle of lost opportunity (Kacelnik, 1984; Kacelnik and Todd, 1992). However, from a psychological perspective our experiment is critically different from these latter experiments. In the patch problem the time to encounter a new patch (the travel) comes between the decision to leave the current patch and future reward, with the forager starting a travel interval by deciding to abandon the current patch. In contrast, the psychological impact of the ITI in our experiment may be less than that of a travel interval, because instead of occurring between the decision and food it comes after the food has been obtained in the chosen option, with the consequence that at no point does the forager choose to enter the ITI. In support of this argument, other studies in the psychological literature have demonstrated that the impact of a time interval on choice is indeed dependent on whether it comes before or after food (e.g., Green et al., 1981; Snyderman, 1983). A possible explanation for this apparently suboptimal behavior is that freely foraging animals never encounter the situation of being forced to wait after food has been obtained before the next decision can be made. If this is the case, then animals may have evolved only to pay attention to time intervals that intervene between a decision and food. This remains an untested proposition.

Why do the birds use EoR rather than long-term rate?

The most important finding of our experiments is that starlings behaved as if they were maximizing EoR rather than long-term rate. As well as fitting with previous results from starlings and bumblebees, this is consistent with Mazur's finding that pigeons maximize EoR calculated without the ITI, but it challenges the predictions derived from overall rate maximization.

Before discussing the interpretation of the above result we shall consider whether our data shed any light on two other questions of currency currently under consideration in the foraging literature. These are, first, whether rate (energy over time) or alternatively "efficiency" (energy gained over energy spent) more accurately describes the currency maximized, and second, if it is rate that is calculated, whether it is gross or net rate of energy intake (e.g., Kacelnik, 1984, Schmid-Hempel et al., 1985; Ydenberg et al., 1994). In fact, our experiments do not address these issues because in both cases discriminating the hypotheses requires that the birds have different metabolic rates during the different periods of the foraging cycle [see Kacelnik and Houston (1984) and the Appendix of this article for proofs of this statement regarding rate and efficiency, respectively]. In our experiments it is likely that there were no substantial differences in metabolic rate during the different periods of each trial, since the birds were not forced to fly during the ITIs as they are in other paradigms in the field (Kacelnik, 1984) and the laboratory (Cuthill, et al., 1994). However, we focus our discussion on net rate maximizing because in situations where there are differential metabolic costs the data support this currency over either gross rate maximizing or efficiency.

Our results agree qualitatively but disagree quantitatively with some explanations for why animals prefer a variable delay to reward over a fixed delay equal to the arithmetic mean of the variable mixture. The first such explanation we shall consider is Scalar Expectancy Theory (SET) (Brunner et al., 1992; Gibbon et al., 1988). By combining assumptions about the processing of information relating to time intervals (and possibly also amounts of food) with functional considerations, SET accounts well for the behavior seen in a number of different foraging experiments [see Bateson and Kacelnik (in press) and Brunner et al. (1992) for examples in starlings]. SET normally assumes that time intervals are remembered with constant relative error, that the memory for each foraging option is formed as the aggregate of total experience of it, and that decisions are taken by taking a single sample from the memory for each option. It follows from these assumptions that subjective equality between a two-valued variable delay (in our case the double standard) and a fixed delay (in our case the adjusting option) should occur when the delay in the fixed option is equal to the geometric mean (in our case $\sqrt{2.5 \times 60.5} = 12.30$ s) of the two delays in the variable option [see Bateson and Kacelnik (1995b) for proof], whereas in this experiment subjective equality was found at the harmonic mean, $(\frac{1}{2} (\frac{1}{2.5} + \frac{1}{60.5}))^{-1} = 4.80$ s, which is equivalent to the EoR prediction. We do not believe that this quantitative discrepancy is fatal to the SET framework. Rather, we believe that modifications will be needed in the assumptions made about either the representation of delays in memory or their subsequent recall in order for a model in the SET framework to explain this result (e.g., Brunner et al., 1994).

A second set of models qualitatively compatible with our results are those based on the risks of interruption during foraging. According to this view, a delayed reward has a lower expected value than a more immediate one because it has a greater chance of being lost during the extra delay (e.g., Kacelnik et al., 1986; McNamara and Houston, 1987). A variety of

models have been proposed that predict different shaped discounting functions, but all of which predict that a variable delay to food should be valued more highly than a fixed delay with the same arithmetic mean value. Original predictions based on a constant probability of loss per unit time lead to exponential discounting, which is not compatible with EoR. However, newer post-hoc theoretical models predict hyperbolic discounting, which is compatible with EoR. Although these latter versions of discounting could accommodate our EoR result, there is a different reason for rejecting explanations based on the probability of interruption. Our birds were kept under continuous reinforcement for many days during the experiments, experiencing thousands of trials that always ended in reward. Under these conditions it would be curious if the birds failed to learn that the probability of interruption was zero. In order to salvage the interruption view, it could be argued that animals may be preprogrammed to behave as if interruptions are always possible because this is the case in the natural environment of the starling, but this is contradicted by evidence showing that birds can learn that reinforcement is probabilistic when this is the case (e.g., Catania and Reynolds, 1968). Thus, time discounting due to probability of reward loss does not provide a parsimonious explanation for our results or, indeed, for the many observations of hyperbolic discounting in continuous reinforcement laboratory conditions.

This section of the discussion would not be complete without mentioning risk-sensitive foraging theory. It has been shown theoretically that a bird on a negative energy budget minimizing its probability of overnight starvation should prefer an option that is variable in the predicted delay to obtain food over one that offers a fixed delay of the same average length (McNamara and Houston, 1992; Zabludoff et al., 1988). A preference for variability in our titration would lead to equilibrium, with the fixed alternative shorter than the arithmetic mean of the variable one, as we obtain. However, we cannot offer a precise test of this hypothesis because a quantitative prediction of the level of the preference (and thus of the point of equality) depends on the value of parameters that are impossible to estimate, and which (as with the interruption analysis) would not apply to the laboratory conditions. However, the risk-sensitive approach predicts that if some of these conditions are modified, the reverse preference should be observed, and to date there have been no demonstrations that the probability of preferring a variable or fixed delay (as opposed to amount) is influenced by energy budget as predicted by the risk-sensitive models (Ha, 1991; Ha et al., 1990).

In an attempt to present some plausible explanations for EoR maximization we shall begin by clarifying the difference between EoR and long-term rate. Although we initially presented long-term rate and EoR as discrete currencies, they could be seen as two extreme points on a plane of possibilities described by two parameters: the total number of prey items remembered by the forager, n , and the size of the frame over which the forager computes the rate, f . Under this view, both long-term rate and EoR are specific cases of a general expression for rate where

$$\text{Rate} = \frac{\sum_{j=1}^{n/f} \left(\frac{\sum_{i=1}^f G_i}{\sum_{i=1}^f T_i} \right)}{n/f} \quad (4)$$

When n is large and $f = n$ this converges on long-term rate, and when n is large and $f = 1$ it converges on EoR. If we use

our current knowledge about starlings from a variety of different paradigms it would appear that the currency that best describes their choices is Equation 4 with $n \approx 6$ (Cuthill et al., 1994) and $f = 1$ (e.g., this article). Note that the discrepancy between long-term rate and EoR lies in the frame size, f , rather than in the total amount of information on which the forager bases its assessment of value, n (see Table 4 for a numerical illustration of this). In fact, Equation 4 does not cover all of the plausible forms of averaging foraging experience because it implies that the last n events have equal weight in the forager's decision. In reality animals are more likely to use some type of moving weighted average in which more recently experienced events have a larger impact on the rate computed.

We shall discuss two categories of explanations for EoR, some based on optimizing under information processing constraints and others based on unconstrained optimality. The first constraint-based explanation is the neuroeconomic argument made for bumblebees by Real (1991). The idea is that information about rate of intake from each flower visited will be more efficiently stored as a single value, rather than as amounts and times to be processed separately. This assumes that the amount gained and time spent acquiring each item are perceived, or at least stored, directly as their ratio, G/T . Thus, each food item gives one estimate of rate of gain in the environment, rather than one estimate of gain per item and one of time taken per item. If foragers were forced to do this because of a neural constraint, computing EoR would be the only available approximation of the long-term rate because independent information about expected times and expected amounts per item would not be available. Evidence against this explanation in starlings comes from the patterns of key pecking during the delays to food. Peaks of pecking rate at the appropriate times show that they possess information about the possible delays to food in each of the options independent of the rate of intake provided by the option (Bateson and Kacelnik, 1995b). If EoR maximization is to be explained by constraints on the amount of memory available, it seems inconsistent that animals do remember this information but do not seem to use it to make foraging choices. However, a different neuroeconomic argument might still be made by arguing that it is possible that the constraint is not the amount of storage space available, but the relative ease of computing EoR versus long-term rate.

A second constraint-based explanation for EoR maximization arises from the fact that as the value of n approaches one in Equation 4 the two rate currencies calculated with $f = n$ and $f = 1$ converge on the same value, providing a poor instantaneous estimate of either currency because of being based on a small sample. Over a relatively long foraging period, the average of the instantaneous estimates will be equivalent to Equation 4 with a large n and $f = 1$, namely to the EoR experienced during the period. Thus if animals are constrained to use very short memory windows then this could explain why they appear to maximize EoR. There is evidence that bees may use a memory window of only one or two experiences, which Real et al. (1990) have used to explain why they maximize EoR. Given that the use of very short memory windows might give mechanistic support for the EoR result, we shall now examine if our starlings may have been using a one-event memory.

Let us consider what a one-event memory would generate in the delay treatment of Experiment 2. In the forced trials the subject experiences D_{adj} in the adjusting option and either D_{int} or D_{brg} (see nomenclature in Table 4) in the double standard. We can predict equilibrium in the titration by considering cases when D_{adj} takes different values. An animal remembering only its last experience of each option would base

Table 4
Numerical illustration of the different average rates that can be computed using Equation 4

Item	1	2	3	4	5	6	7	8	9	10	11	12		
Frame, f	Gain, G	7.09	9.59	4.41	3.12	6.09	9.53	4.47	0.82	0.46	1.83	3.15	2.21	Mean ($\Sigma G/\Sigma T$)
	Time, T	2.31	8.38	2.13	6.29	0.30	4.97	8.10	6.55	1.67	0.21	9.11	8.39	
12	ΣG	52.77											0.90 = RoE	
	ΣT	58.39												
	$\Sigma G/\Sigma T$	0.90												
6	ΣG	39.83					12.94					1.01		
	ΣT	24.37					34.01							
	$\Sigma G/\Sigma T$	1.63					0.38							
4	ΣG	24.21			20.91			7.65			0.90			
	ΣT	19.11			19.91			19.37						
	$\Sigma G/\Sigma T$	1.27			1.05			0.40						
3	ΣG	21.09		18.74		5.75		7.19		1.01				
	ΣT	12.82		11.55		16.31		17.70						
	$\Sigma G/\Sigma T$	1.64		1.62		0.35		0.41						
2	ΣG	16.68		7.53		15.62		5.29		2.29		5.36		1.22
	ΣT	10.69		8.42		5.26		14.65		1.88		17.49		
	$\Sigma G/\Sigma T$	1.56		0.89		2.97		0.36		1.22		0.31		
1	ΣG	7.09	9.59	4.41	3.12	6.09	9.53	4.47	0.82	0.46	1.83	3.15	2.21	3.29 = EoR
	ΣT	2.31	8.38	2.13	6.29	0.30	4.97	8.10	6.55	1.67	0.21	9.11	8.39	
	$\Sigma G/\Sigma T$	3.07	1.14	2.07	0.50	20.45	1.92	0.55	0.13	0.28	8.74	0.35	0.26	

All of the rates are based on the same 12 prey items ($n = 12$). Each item has a gain (G) and time (T) associated with it (here these are random numbers). The average rate experienced for the 12 items depends on the frame (f) over which gains and times are summed before a rate is computed. The table shows the calculations for six different frame sizes. Each box of the table represents a frame; its boundaries indicate which prey items it spans. The three numbers in each box are the sum of G for the frame (top), the sum of T for the frame (middle), and the ratio of these two sums (bottom). The average rate obtained for each frame size is displayed in the right-hand column of the table. Note that if $f = 12$ (i.e., $f = n$) then the resulting average rate is equal to the long-term rate (RoE), and if $f = 1$ then the resulting average rate is equal to the expectation of the ratios (EoR). This example serves to demonstrate that the discrepancy between long-term rate and EoR is due to occasional small values of T (as in the case of item 5), which have a particularly large impact on the average rate when $f = 1$.

the first of the two choices in each block on the most recent forced trials, whereas the second of the two choices will also be influenced by what is experienced as a result of the first choice. When $D_{adj} < D_{shr}$, regardless of whether the last forced trial of the double standard was D_{shr} or D_{lng} , both choices will be for the adjusting option. Due to this biased choice, in the next block, D_{adj} will be longer, but it will continue to be chosen until $D_{adj} > D_{shr}$. Contingencies in the range $D_{shr} < D_{adj} < D_{lng}$ are more complicated: in one-quarter of the pairs of choices the standard will be preferred twice (when the last forced trial gave D_{shr} and the first choice yielded D_{shr}), in another quarter the first choice will be for the standard and the second for the adjusting option (when the last forced trial gave D_{shr} and the first choice yielded D_{lng}), and in the remaining half the adjusting option will be chosen twice (when the last forced trial gives D_{lng} and the first choice thus yields D_{adj}). Thus on average D_{adj} increases twice as often as it decreases, until it reaches D_{lng} . When $D_{adj} > D_{lng}$ all choices will be for the standard and D_{adj} will therefore decrease. In summary, if choices were based on a memory for the last event in each option the value of the adjusting option would fluctuate around D_{lng} (60.5 s in our experiment), which is very far from the empirical result we obtained (mean of 5.61 s). This means that our subjects were definitely not using a one-event memory.

Although our data are incompatible with a memory of a single event, it is worth emphasizing that there are conditions, different from our experimental situation, under which using such a short memory would actually be the long-term rate maximizing policy. This happens when successive instances of a mixture do not occur independently because of serial correlation in the environment. If, for instance, short and long

delays came in bouts, then the subjects could experience a rate higher than that expected from the arithmetic mean of the two delays, because they could use the last event to predict the outcome of the next one, preferring the variable option when it is more likely to give a short than a long delay (Cuthill et al., 1990; Real, 1991). This idea must be considered here for completeness, because if the natural environment was serially correlated and the birds had acquired a rigid mechanism of using only the last event, they may follow this strategy in situations where there is no serial correlation, such as ours. This is a sound argument for the possibility of an unconstrained advantage for single-event memory in some special situations, but it can be dismissed in our experiments because the previous analysis has shown that the birds were not using the last event only.

Given that there are currently no plausible arguments for why maximizing EoR should result in higher fitness than maximization of long-term rate, we shall now examine the possibility that under natural circumstances there is no advantage of computing one over the other. In our experiment we chose a double standard with one very small and one very large value of T to exaggerate the discrepancy in the predictions of the long-term rate and EoR. The example in Table 4 makes it clear that small values of T have a larger impact on EoR than on long-term rate. However, long-term rate and EoR are identical if there is no variation in T , and in fact, as long as the distribution of T does not contain very small values, the discrepancy in energetic intake between EoR and long-term rate will be small. In this situation natural selection would be neutral with respect to whether rates are calculated using EoR or long-term rate if the only fitness consequences of the curencies are measured in terms of energetic intake. This idea

Table 5
Key to the symbols used in the Appendix

Symbol	Meaning
I	Intertrial interval (always 40 s)
L	Latency (always assigned an arbitrary value of 1 s)
A_{adj}	Amount of food in adjusting option (in units)
A_{stn}	Amount of food in standard option (in units)
D_{adj}	Delay in adjusting option (in seconds)
D_{sh}	Short delay in standard option (in seconds)
D_{lng}	Long delay in standard option (in seconds)
$F_{adj} = A_{adj}$	Feeding time in adjusting option (in seconds)
$F_{stn} = A_{stn}$	Feeding time in standard (in seconds)

Note that the feeding time, F , is always equal to the amount of food, A , because one unit of food took exactly 1 s to deliver.

could be tested by measuring the distribution of T for foragers in their natural environment and calculating whether or not there is a significant discrepancy between long-term rate and EoR (see Bateson and Whitehead, in press).

In conclusion, we have shown laboratory evidence that foraging choices in the starling are based on a currency (EoR), for which we see no obvious functional advantage. What are the implications of this for optimal foraging theory? It is possible that foraging theory, while remaining the most quantitatively precise field for the study for animal decision making, may be too narrow for our level of analysis. Functional modeling of foraging decision making may have to incorporate mechanistic factors such as the computational demands of implementing alternative algorithms.

APPENDIX

This appendix describes how the predictions given in Tables 2 and 3 were calculated. The definitions of the terms used in the equations are given in Table 5

If the starlings are using long-term rate (RoE) computed with all of the time intervals to value the two options, then the adjusting option will be equivalent to the double standard when the long-term rate in the adjusting option is equal to the long-term rate in the standard, that is,

$$\frac{A_{adj}}{I + L + D_{adj} + F_{adj}} = \frac{2A_{stn}}{2I + 2L + D_{sh} + D_{lng} + 2F_{stn}} \quad (A1)$$

Whereas, if the starlings are using the expectation of the ratios (EoR) computed with all of the time intervals to value the options, then the adjusting option will be equivalent to the double standard when the EoR in the adjusting option is equal to the EoR in the double standard, that is,

$$\frac{A_{adj}}{I + L + D_{adj} + F_{adj}} = \frac{1}{2} \left(\frac{A_{stn}}{I + L + D_{sh} + F_{stn}} + \frac{A_{stn}}{I + L + D_{lng} + F_{stn}} \right) \quad (A2)$$

Efficiency is an alternative currency to rate defined as the ratio of energy gained to energy spent. Assuming that metabolic rate is constant during foraging, efficiency can be expressed as follows:

$$\text{efficiency} = \frac{A - mT}{mT} = \frac{A}{mT} - 1 \quad (A3)$$

where A is the gross amount of food gained, T is the time

spent foraging, and m is the metabolic rate. Just as for rate, efficiency can be calculated either in the long-term or as an EoR. If the starlings use long-term efficiency to value the options, then the adjusting option will be equivalent to the double standard when the long-term efficiency in the adjusting option is equal to the long-term efficiency in the standard, that is,

$$\frac{A_{adj}}{m(I + L + D_{adj} + F_{adj})} - 1 = \frac{2A_{stn}}{m(2I + 2L + D_{sh} + D_{lng} + 2F_{stn})} - 1 \quad (A4)$$

The m terms cancel, leaving the above equation identical to Equation A1. Similarly, for the EoR form of efficiency the adjusting option will be equivalent to the double standard when

$$\frac{A_{adj}}{m(I + L + D_{adj} + F_{adj})} - 1 = \frac{1}{2} \left(\frac{A_{stn}}{m(I + L + D_{sh} + F_{stn})} - 1 + \frac{A_{stn}}{m(I + L + D_{lng} + F_{stn})} - 1 \right) \quad (A5)$$

Again the m s cancel leaving the above equation identical to Equation A2. Therefore, in the experiments described in this article, the predictions based on rate are identical to those based on efficiency.

In the delay treatment of Experiment 2, the only value in Equations A1 and A2 that is not fixed is that of D_{adj} , since this is the value obtained from the titrations. Thus, by solving Equations A1 and A2 for D_{adj} we can calculate the value of D_{adj} predicted if the birds are using each of the rate currencies. Similarly, in the amount treatment of Experiment 2 and in Experiment 3, the only value not fixed is that of A_{adj} , since this is the value obtained from the titrations. Thus, by solving Equations A1 and A2 for A_{adj} we can calculate the value of A_{adj} predicted if the birds are using each of the rate currencies.

Since the latency to peck, L , is controlled by the birds, we had to make some assumption about its length in order to make the predictions that include L in the calculations. For the predictions in Tables 2 and 3 we used an arbitrary value of 1 s for L . To compute the predictions without the intertrial interval, latency, or feeding time, we solved the same equations omitting the terms I , L , and F as appropriate.

For example, in the delay treatment of Experiment 2 the adjusting delay, D_{adj} , predicted by maximization of EoR computed with delay and feeding time only is calculated as follows. The values programmed in the experiment are substituted into Equation A2, which is then solved to give D_{adj} :

$$\frac{5}{D_{adj} + 5} = \frac{1}{2} \left(\frac{5}{2.5 + 5} + \frac{5}{60.5 + 5} \right)$$

$$D_{adj} = 8.46.$$

We thank the Wellcome Trust for financial support (research grant no. 036990 to A.K.) and Miguel Rodriguez-Girones for helpful discussion.

REFERENCES

Bateson M, Kacelnik A, 1995a. Accuracy of memory for amount in the foraging starling (*Sturnus vulgaris*). *Anim Behav*. 50:431–443.
Bateson M, Kacelnik A, 1995b. Preferences for fixed and variable food

- sources: variability in amount and delay. *J Exp Anal Behav* 63:313–329.
- Bateson M, Whitehead SC, in press. The energetic costs of alternative rate currencies in the foraging starling. *Ecology*.
- Brunner D, Gibbon J, Fairhurst S, 1994. Choice between fixed and variable delays with different reward amounts. *J Exp Psychol Anim Behav Processes* 20:331–346.
- Brunner D, Kacelnik A, Gibbon J, 1992. Optimal foraging and timing processes in the starling, *Sturnus vulgaris*: effect of inter-capture interval. *Anim Behav* 44:597–613.
- Caraco T, Kacelnik A, Mesnik N, Smulewitz M, 1992. Short-term rate maximization when rewards and delays covary. *Anim Behav* 44:441–447.
- Catania CA, Reynolds GS, 1968. A quantitative analysis of the responding maintained by interval schedules of reinforcement. *J Exp Anal Behav* 11:327–385.
- Cuthill IC, Haccou P, Kacelnik A, 1994. Starlings (*Sturnus vulgaris*) exploiting patches: response to long-term changes in travel time. *Behav Ecol* 5:81–90.
- Cuthill IC, Kacelnik A, Krebs JR, Haccou P, Iwasa Y, 1990. Starlings exploiting patches: the effect of recent experience on foraging decisions. *Anim Behav* 40:625–640.
- Gibbon J, Church RM, Fairhurst S, Kacelnik A, 1988. Scalar expectancy theory and choice between delayed rewards. *Psychol Rev* 95:102–114.
- Gilliam JF, Green RF, Pearson NE, 1982. The fallacy of the traffic policeman: a response to Templeton and Lawlor. *Am Nat* 119:875–878.
- Green L, Fisher EB, Perlow S, Sherman L, 1981. Preference reversal and self-control: choice as a function of reward amount and delay. *Behav Anal Lett* 1:244–256.
- Ha JC, 1991. Risk-sensitive foraging: the role of ambient temperature and foraging time. *Anim Behav* 41:528–529.
- Ha JC, Lehner PN, Farley SD, 1990. Risk-prone foraging behaviour in captive grey jays *Perisoreus canadensis*. *Anim Behav* 39:91–96.
- Harder L, Real LA, 1987. Why are bumble bees risk-averse? *Ecology* 68:1104–1108.
- Kacelnik A, 1984. Central place foraging in starlings (*Sturnus vulgaris*). I. Patch residence time. *J Anim Ecol* 53:283–299.
- Kacelnik A, Houston AI, 1984. Some effects of energy costs on foraging strategies. *Anim Behav* 32:609–614.
- Kacelnik A, Todd IA, 1992. Psychological mechanisms and the marginal value theorem: effect of variability in travel time on patch exploitation. *Anim Behav* 43:313–322.
- Kagel JH, Green L, Caraco T, 1986. When foragers discount the future: constraint or adaptation? *Anim Behav* 34:271–283.
- Mazur JE, 1984. Tests of an equivalence rule for fixed and variable reinforcer delays. *J Exp Psychol Anim Behav Processes* 10:426–436.
- Mazur JE, 1986. Fixed and variable ratios and delays: further tests of an equivalence rule. *J Exp Psychol Anim Behav Processes* 12:116–124.
- McDiarmid GC, Rilling ME, 1965. Reinforcement delay and reinforcement rate as determinants of schedule preference. *Psychonom Sci* 2:195–196.
- McNamara JM, Houston AI, 1987. A general framework for understanding the effects of variability and interruptions on foraging behaviour. *Acta Biotheor* 36:3–22.
- McNamara JM, Houston AI, 1992. Risk-sensitive foraging: a review of the theory. *Bull Math Biol* 54:355–378.
- Possingham HP, Houston AI, McNamara JM, 1990. Risk-averse foraging in bees: a comment on the model of Harder and Real. *Ecology* 71:1622–1624.
- Real LA, 1991. Animal choice behavior and the evolution of cognitive architecture. *Science* 253:980–986.
- Real L, Ellner S, Harder LD, 1990. Short-term energy maximisation and risk-aversion in bumble bees: a reply to Possingham et al. *Ecology* 71:1625–1628.
- Reboreda JC, Kacelnik A, 1991. Risk sensitivity in starlings: variability in food amount and food delay. *Behav Ecol* 2:301–308.
- Schmid-Hempel P, Kacelnik A, Houston AI, 1985. Honeybees maximise efficiency by not filling their crop. *Behav Ecol Sociobiol* 57:509–521.
- Snyderman M, 1983. Optimal prey selection: partial selection, delay of reinforcement and self control. *Behav Anal Lett* 3:131–147.
- Stephens DW, Krebs JR, 1986. Foraging theory. Princeton, New Jersey: Princeton University Press.
- Templeton AR, Lawlor LR, 1981. The fallacy of the averages in ecological optimization theory. *Am Nat* 117:390–393.
- Todd IA, Kacelnik A, 1993. Psychological mechanisms and the marginal value theorem: dynamics of scalar memory for travel time. *Anim Behav* 46:765–775.
- Turelli M, Gillespie JH, Shoener TW, 1982. The fallacy of the fallacy of the averages in ecological optimization theory. *Am Nat* 119:879–884.
- Ydenberg RC, Welham CVJ, Schmid-Hempel R, Schmid-Hempel P, Beauchamp G, 1994. Time and energy constraints and the relationship between currencies in foraging theory. *Behav Ecol* 5:28–34.
- Zabludoff SD, Wecker J, Caraco T, 1988. Foraging choice in laboratory rats: constant vs. variable delay. *Behav Process* 16:95–110.