Risk sensitivity in starlings: variability in food amount and food delay

Starlings' preferences for constant versus variable food sources were studied in the laboratory. The constant alternative gave a fixed amount of food after a fixed delay. The variable alternative offered either a varying amount of food after a fixed delay (treatment A) or a fixed amount of food after a variable delay (treatment B). In both treatments the ratio of amount of food over trial length (the sum of intertrial interval plus delay and handling times) of the constant alternative equaled the average of the two ratios of the variable alternative. The variable ratios were 30% higher and 30% smaller than the fixed ratio. In free-choice trials (both options available in each trial), the subjects were risk-averse or indifferent in treatment A and indifferent or risk-prone in treatment B. In no-choice trials (only one source available per trial), the latency to respond was longer in the variable than in the constant source in treatment A and the opposite in treatment B. The greater preference for variability in time than for variability in reward amount is not consistent with either maximizing the ratio of expected energy over expected time or the expected ratio of energy over time for individual trials. There was a negative correlation between individual intake rate and degree of risk proneness for both kinds of variability. We present a model of choice based on an information-processing theory for temporal memory that accounts for the different effects of variability in delay and in amount but cannot explain the effects of intake rate. [Behav Ecol 1991;2:301–308]

Juan C. Reboreda
Sub-department of Animal Behaviour,
University of Cambridge,
Madingley, Cambridge, UK

Alejandro Kacelnik
King's College,
Cambridge, UK

The important effects of environmental variability on behavior have been recognized for some time both in foraging studies and in experimental psychology (Caraco et al., 1980; Gibbon et al., 1988; Mazur, 1984; McNamara and Houston, 1987; Pubols, 1962). Within foraging theory, it has been common to equate variability to uncertainty, and thus optimality models have been used to propose optimal foraging strategies when the fitness consequences of each action cannot be predicted deterministically. The most common approach has been to examine preference between habitats or patches offering various degrees of variability in prey size. Predictions are based on the function relating prey size to the resulting fitness gains. If this function is positively accelerated, namely, if larger food items give disproportionately high benefits, then a habitat with larger variance will give larger expected fitness gain than a less variable habitat with the same mean prey size (Jensen's inequality). The opposite is true if fitness gains are a negatively accelerated function of prey size. This consideration is encapsulated in the so-called "energy budget rule" (Stephens and Krebs, 1986): If a habitat provides energy at

an average rate that is lower than required for survival, then the function is assumed to be positively accelerated and animals are expected to prefer variability, while the opposite is true if average intake rate is above the minimum required. This switch in preference depending on the energy budget has been sought in several studies (see review in Real and Caraco, 1986; Stephens and Krebs, 1986) and has received support in recent studies (Caraco et al., 1990; Cartar and Dill, 1990; Young et al., 1990).

Variability in prey size is not the only possible form of variance experienced during foraging. Animals may, for example, encounter prey of equal size after variable searching times. McNamara and Houston (1987), using a dynamic model, have shown that this form of variability is not equivalent to variation in prey size. The main reason for this lack of equivalence is that animals normally forage under finite time horizons, and foraging cycles of searching followed by captures use up the available foraging time in addition to resulting in changes in energy state. However, time variability also leads to a predicted switch in preference dependent on energy budget.

What differs between forms of variability is the geometry of the boundary between risk-prone and risk-averse behavior in a space with energy state and remaining time as dimensions.

Few laboratories have attempted to find a switch in preference for variability in searching time. Ha et al. (1990) and Ha (1991) reported that grey jays (*Perisoreous canadensis*) showed persistent preference for a food source requiring a variable number of responses over an alternative offering the same amount of food after a constant number of responses, regardless of well-documented changes in energy budget. The number of responses required could be translated into variable searching times between prey. Their report is consistent with the outcome of many studies in psychological laboratories, where animals are often tested for preference between food sources offering delayed food reinforcements with or without variability. The overwhelming outcome in this kind of study is a strong preference for variable delays (Davison, 1969, 1972; Gibbon et al., 1988; Herrnstein, 1964; Killeen, 1968; Leventhal et al., 1959; Mazur, 1984; Pubols, 1962).

It might be argued that preference for delay variability could result from the schedules used in these studies. The standard technique in studies of variability in either prey size or delay to food is to have a constant and a variable alternative, where the latter has either a symmetric or positively skewed distribution of prey size or delay at each side of the constant alternative. These distributions could generate biases if animals categorize feeding episodes (say a chase and a capture) in terms of rate of gain (energy/time) per episode instead of computing independently energy and time. Such categorization would lead to preference for the alternative of variable delay because the expected ratio of energy over time would be higher in the variable side even though the ratio of expected energy over expected time is equated in the procedure. We illustrate this point with one example. Assume that a fixed food source delivers a reinforcement of size S following a delay τ. A subject using this option will experience a feeding rate of S/τ (assuming that no other times are involved). If a variable alternative offered reinforcements of the same size after delays of either $\tau - \delta$ or $\tau + \delta$, then a subject using it would experience a rate of

$$\frac{S}{[(\tau - \delta) + (\tau + \delta)]/2},$$

which is the same as in the fixed option. But if the subject remembers the experienced rate every time it collects a reward and averages these measurements, it will value the variable alternative according to

$$\frac{[S/(\tau - \delta)] + [S/(\tau + \delta)]}{2} = \frac{S\tau}{\tau^2 - \delta^2},$$

which is bigger than S/τ for all possible values of δ. This argument, which is based on direct perception of rates, is one possible (mechanistic) explanation for preference for delay variability in experiments such as those reported by Ha et al. (1990), Ha (1991), and Zabludoff et al. (1988). Notice that the argument may not apply if the variable source does not offer a symmetric distribution around the fixed source value. There is some evidence that variability per se cannot account for preference and that shape of the distribution in the variable option is important. In a study in which the variable alternative was programmed with a negatively skewed distribution ("backward exponential"), pigeons showed no bias (Gibbon et al., 1988). Variability in prey size would not produce this bias in subjects experiencing the same perceptual constraint because when numerators are concerned the ratio of the expectations is equivalent to the expectation of the ratio.

Subjects using single-cycle ratios as direct percepts of rate would be "per patch rate maximizers" instead of being long-term rate maximizers (Stephens and Krebs, 1986). It could be argued that these subjects, the same as some foraging theorists, commit the so-called "fallacy of the fallacy of the averages" (Gilliam et al., 1982; Templeton and Lawlor, 1981; Turelli et al., 1982). The possibility of trial-by-trial rate perception poses a problem for the interpretation of previous experiments. Variance is generally programmed between individual captures or individual searching times, and not by variance in the accumulated gain or time experienced in aggregates of many episodes. This has two consequences. First, variable alternatives offer little variance in the ratio of accumulated gain over accumulated time when considered over a large number of trials, and thus there are necessary additional assumptions to specify the animals' assessment of variability. Second, in experiments such as those by Ha et al. (1990) and Ha (1991), the expectation of the ratios was higher in the variable than in the fixed side, and this could have confounded the effect of risk sensitivity per se.

We present a study in which (1) preferences for variability in prey size and food delay are tested within subjects, with subjects in the same energy state and under the same time horizon and (2) variable alternatives in either size or delay are programmed to present equivalent deviations from the mean in rate of gain per cycle. The variable side offers ratios of gain/time per episode 30% above or below the ratio of the constant side.

METHODS

Subjects

The subjects were 12 starlings (*Sturnus vulgaris*) caught in the wild. After capture, we housed the birds in an outdoor aviary for 4–6 weeks. During this period we provided the birds with water and food (starter turkey crumbs) ad libitum and mealworms (*Tenebrio* sp.) twice a week. We then moved the birds indoors to individual cages measuring 0.9 m long × 0.45 m wide × 0.5 m high. The birds were divided into two groups of six, which were housed in separate rooms at 20 ± 2°C with a 12h:12h light–dark cycle. We deprived birds of food for 16 h before each experimental session and gave them a complement of food 30 min after each experiment, so that the total daily ration was 20 g.

Apparatus

The experiments were conducted in the house cages. Each cage had a removable panel with a centrally mounted food hopper and two response keys (Campden Instruments). The response keys (2.5 cm diameter) were 8 cm to the left and to the right of the food hopper. In front of the panel there was a perch 1 cm in diameter and 16 cm below the center of the pecking keys. A BBC Master microcomputer using SPIDER experimental control language (Paul Fray Ltd) controlled the stimulus events and response contingencies and recorded the data.

Training

The birds received 8 daily sessons (40 trials per session) of magazine training (4 s of access to food allowed every 30 s). After this period all the birds were retrieving food from the grain hopper. Then the delivery of food was paired with the key light. Each trial started with the transillumination of the pecking key with amber light and after 6 s there were 4 s of access to food. Then the magazine and key lights went dark and an intertrial interval (ITI) averaging 45 s (30 minimum, 60 maximum, uniform distribution) started. In this condition, despite the absence of any programmed consequence of pecking, the amber key light came to elicit key pecking responses in all birds. After 4 sessions (40 trials per session), the delivery of food was made conditional to at least one key peck (i.e., if the birds did not respond during the 6-s time window, the key light went dark and a new ITI started). The birds received another 17 sessions in this condition.
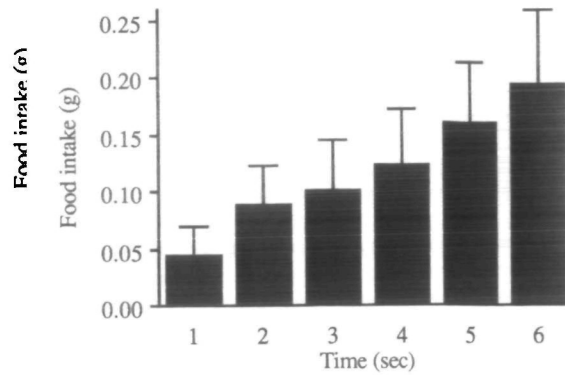
C̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶̶

During the last part of the training period, we determined the relationship between time of access to the feeder and food intake. In each of six conditions (1–6 s, 1-s steps) the birds received 20 trials. We determined the food intake during the last 10 trials. In half of the trials the bird had to respond in the left key and in the other half in the right key. The order of the treatments (2 per day) was 2, 1, 4, 3, 6, and 5 s. In two of the treatments (1 and 2 s) we compared the amount of food retrieved in rewards obtained by responses in the left versus the right key, and no differences were detected. In all the subjects the amount of food retrieved was a linear function of the time of access to the feeder (Figure 1 and Table 1). Hence, we assumed time of access to the feeder was equivalent to amount of food retrieved or reward magnitude.

Experiment

The constant and variable alternatives were indicated by the color of the keys. The birds had no previous experimental experience with the two colors used (red and green). Half of the subjects received red associated with the

Table 1

Linear regressions of food intake versus seconds of access to the feeder

Bird	Slope (mg/s)	R^2	p	Rate (g/h)
1	28.66	.65	.052	8.06
2	32.51	.71	.034	10.58
3	27.34	.91	.003	8.86
4	25.06	.93	.002	7.95
5	33.74	.81	.014	9.42
6	42.23	.95	.0009	12.63
7	33.26	.86	.0076	9.53
8	14.14	.94	.0013	3.88
9	14.14	.98	.0002	3.79
10	30.6	.95	.001	8.97
11	26.77	.86	.007	7.86
12	27.34	.97	.0002	8.78

Figure 2

Reward magnitude and total trial length used in the experiment. The symbol on the middle line shows the values in the constant alternative in both treatments. The other filled squares show the two possible outcomes in the treatment with variability in reward magnitude. Although delay in this treatment is fixed, trial length is slightly different between outcomes due to the different access time to the feeder. The open squares show the outcomes in the treatment with variability in delay to reward. The slopes of the lines shown are the per trial rates of gain. The values were chosen so that the per trial ratios of the variable outcomes deviate symmetrically from the constant option. In all cases the intertrial interval (τ) was 30 s. The ratios are calculated assuming a short latency to peck (R = 3 seconds).
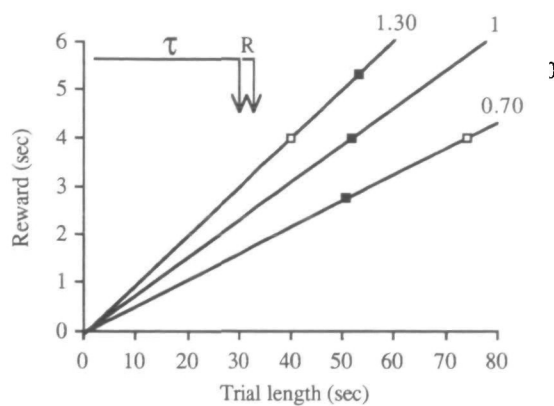
constant alternative and green with the variable one, whereas for the other half this was reversed. The side on which each alternative appeared was randomized trial by trial.

There were two treatments. In both of them the constant alternative gave a fixed amount of food after a fixed delay, whereas the variable alternative offered a varying reward magnitude after a fixed delay (treatment A) or a fixed reward after a variable delay (treatment B). In both treatments the programmed ratio of reward size over trial length of the constant alternative equaled the average of the two programmed ratios of the variable alternative. The programmed ratios of the variable alternative had a mixture of two equally frequent programmed ratios 30% higher and 30% smaller than the constant ratio, respectively (Figure 2).

A discrete-trials procedure, with a mean ITI of 30 s, was employed. Each session began with 12 no-choice trials during which only one of the keys was available for responding, followed by 20 free-choice trials during which subjects could choose to respond on either key. The no-choice trials were given in a semi-random order so that in each session the birds received six constant and six variable trials. The number of low- and high-ratio trials in the variable alternative was balanced.

Free-choice trials began with both keys blinking (0.7 s on, 0.3 s off). The first peck on either key switched the chosen key to a permanent light of the same color and turned the other key dark. In treatment A, after a fixed delay of 15 s the bird received (depending on its choice) either the constant reward (4 s of access to food) or the variable reward (2.7 or 5.3 s of access to food, with equal probability). In treatment B, after the choice peck, the bird experienced either a constant (15 s) or a variable (3 or 37.3 s) delay (depending on its choice) followed by 4 s of access to food. No-choice trials were similar to the free-choice trials, but they began with one key blinking and the other dark.

The 12 birds experienced both treatments (10 sessions each). Half of the subjects received first treatment A and then treatment B, whereas in the other half the order was reversed. Between treatments there were 5 sessions with only no-choice trials (32 per session) with the conditions of the next treatment.

All the subjects could obtain food at a rate which, if sustained for several hours, would have produced a daily energetic gain (see Table 1) in excess of daily energetic requirements (about 20 g per day).

Analysis

During the 12 no-choice trials at the beginning of each session, we recorded the latency to the first peck. Results from the last 3 sessions (8–10) were pooled and expressed as mean ± SEM for the constant and variable outcomes. Comparisons between treatments and alternatives were done with a two-way ANOVA for repeated measures. For the free-choice trials we pooled for each individual the 60 trials from the last 3 sessions and calculated the proportional choice (p) of the constant alternative. We used the normal approximation to the binomial distribution to categorize each p value as variability averse, indifferent, or variability prone toward the alternative. In a similar way we tested for side preferences.

RESULTS

Free-choice trials

Table 2 presents the proportion of free-choice trials in which the subjects chose the constant alternative under both experimental treatments. In treatment A (variability in reward magnitude), five subjects were significantly risk averse, four were indifferent, and three were significantly risk-prone, whereas in treatment B (variability in delay), seven subjects were significantly risk-prone and five were indifferent.

Six subjects showed differences between treatments in statistically significant preference for variability. Three of the five birds that preferred constant over variable reward amount preferred variable over constant delays, and the other two were indifferent to variability in delay. One bird that was indifferent to variability in reward preferred variability over constancy in delay. Three of the six subjects that did not show changes between treatments (birds 3, 7, and 12) showed significant side preferences. The remaining three subjects showed preference for variability in both treatments, but because the variable alternative was associated with the

same color for rewards and delays, it is not possible to exclude color preferences (the subjects followed red in both treatments).

No-choice trials

We conducted a two-way repeated measures ANOVA (treatment and alternative) for the latencies to the first peck during the no-choice trials. There was a significant effect of the interaction treatment–alternative ($F_{1,11} = 13.71$, $p < .03$). Within-treatment comparisons between means showed that latency to peck was higher in the constant than in the variable alternative when variability was in delay ($p < .01$, two-tailed) and lower in the constant than in the variable alternative when variability was in reward, although this comparison had lower significance ($p < .08$, two-tailed). Mean latencies are shown in Figure 3.

In the treatment with variable reward magnitude, the subjects always had 15 s to wait after the first peck. The number of pecks during this delay can provide an additional index of preference. Mean number of pecks showed the same tendency as the latency to the first peck: it was higher in trials with fixed reward (66.47 ± 11.95) than in those with variable reward (48.85 ± 10.27) ($p < .07$, two tailed).

Effect of energy budget

We did not manipulate the energy budget to study the possibility of a preference switch. Nevertheless, as Table 1 shows, there was considerable natural variation between birds in the efficiency with which they collected food from the feeders, and this resulted in some birds eating more than three times as fast as others. Variation in intake rate did not reflect ad libitum food intake and can probably be interpreted as varying efficiency in scooping food from the feeder and not from difference in requirements. We used this variation to test whether risk aversion increased across birds as a function of intake rate. Correlations were calculated using arcsine square root of proportions, and birds 3, 7, and 12 were excluded due to their side bias. There was a positive correlation between intake rate and proportion of choices of the constant alternative in free-choice trials in both the variable reward ($r = .64$, $p < .05$) and the variable delay ($r = .61$, $p < .05$) treatments.

DISCUSSION

We observed a greater preference for delay over reward variability in both free-choice and no-choice trials. In free-choice trials, the subjects showed a weak preference for the constant reward and stronger preference for the

Table 2

Proportion and p value (normal approximation to the binomial distribution) of free-choice trials with response in the fixed source
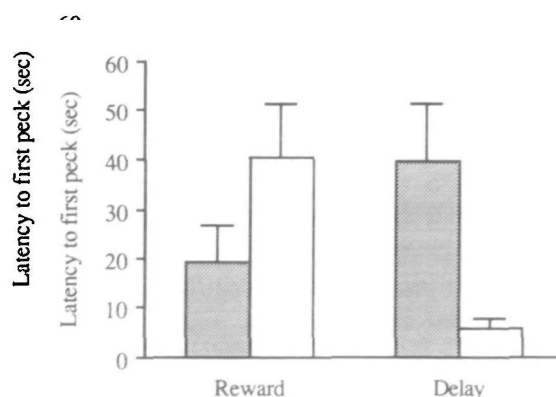
Bird	Variability in reward		Variability in delay		Δ proportion
	Proportion	p	Proportion	p	
1	0.567	.18	0.283[b]	.001	0.284
2	0.633[a]	.03	0.067[b]	.001	0.566
3	0.6	.08	0.5	.45	0.117
4	0.983[a]	.001	0.467	.35	0.516
5	0.983[a]	.001	0.45	.26	0.533
6	0.817[a]	.001	0.383[b]	.05	0.434
7	0.433	.18	0.517	.45	−0.084
8	0.017[b]	.001	0[b]	.001	0.017
9	0.217[b]	.001	0[b]	.001	0.217
10	0.317[b]	.003	0.283[b]	.001	0.034
11	0.767[a]	.001	0[b]	.001	0.767
12	0.517	.45	0.517	.45	0

[a] Risk-averse.
[b] Risk-prone.

variable delay. In no-choice trials, the latency to the first peck was higher in the variable than in the constant alternative when using reward variability and the opposite when the variability was in delay. Similarly, in the treatment with variability in reward magnitude, the number of pecks during the 15-s delay was higher in fixed-reward than in variable-reward trials.

The difference cannot be explained by rate maximization. In the treatment with variability in reward size, both options offered the same programmed rate of gain. In the treatment with variable delays, we programmed the variable delays asymmetrically with respect to the constant delay so that short and long delays resulted in equal deviations in terms of the energy/time ratio, but in a lower long-term rate in the variable than in the constant option. Thus, animals maximizing rate subject to the constraint of trial-by-trial perception of rate would have been neutral in both treatments. Animals maximizing rate but categorizing alternatives by their long-term rate

Figure 3
Latencies to peck in no-choice trials (last three sessions pooled) in treatment A (variability in reward magnitude) and treatment B (variability in length of delay). The bars correspond to the mean ± SEM across the 12 birds. Shaded bars: constant alternative; open bars: variable alternative.

would have avoided the variable delay option while being indifferent to reward variability.

Two other features of the data are problematic for rate maximization. First, it is hard to envisage an optimality reason for showing any latency to start a trial: this only causes a loss of intake without any change in variability of outcomes. Our birds showed considerable pauses before pecking, and in no-choice trials they showed a difference in latency that paralleled the other measures of preference. These pauses create interpretation difficulties: by varying the pauses the birds altered the experienced rates; less preferred options in no-choice trials led to longer pauses and thus lower experienced rates. It is possible that preferences in free-choice trials took into account these resulting rates. One expression of preference (choice in free-choice trials) can thus be seen as rational subject to an (irrational) expression of preference in another measure (delay to start no-choice trials). A similar problem was faced by Killeen et al. (1981) in a test of central place foraging. They found that by varying working rate, the rats created a concave gain function and that patch exploitation seemed to be adaptive with respect to this experienced gain function. Similar to previous authors working on risk sensitivity, we chose to leave aside this difficulty by examining behavior with respect to programmed features of the schedule. Second, although only one peck was required to obtain food, the birds pecked throughout the waiting period and the rate of pecking again paralleled other measures of preferences. This superfluous pecking is probably caused by autoshaping (Reboreda and Kacelnik, in press).

Standard risk sensitivity theory also fails to accommodate our results. According to a simple application of the energy budget rule for risk-sensitive behavior (Stephens and Krebs, 1986), our animals ought to have been risk-averse in both treatments. This is because during experimental sessions all birds could obtain food from the constant alternative at a rate that, if projected toward the end of the day, would have produced greater accumulated intake than required. A realistic implementation of this rule would require an argument for the appropriateness of this projection, but this level of speculation is probably not justified by our results. Nevertheless, the between-subject comparison showed that subjects that were collecting smaller rewards during the experiment and thus were experiencing a lower intake rate were significantly more risk-prone than more efficient animals. This correlation provides indirect support for the predicted changes in risk proneness according to risk-sensitivity theory, although literal application of this theory would require a modified version of the energy budget rule.

The contrasting effects of variability in reward amount and in delay to food could be related to discounting, namely, the convex devaluation of reward value with expected delay (Kagel et al., 1986a). This approach says that delayed reinforcements may be "devalued" by the estimated probability that unexpected intervening events may occur before the delay ends. If discounting is decelerated, then greater certainty of capture would lead to overvaluation of short delay food items and thus generate strong preference for food sources associated to variable delays to food. These ideas are closely related to descriptive empirical models proposed by operant psychologists, in particular by a series of papers by Mazur (1984, 1986, 1987, 1989), which we discuss below as one possible mechanistic interpretation.

Subjective value of delayed reinforcement

Mazur (1984) suggested that the subjective value of an expected reward can be described by the following equation:
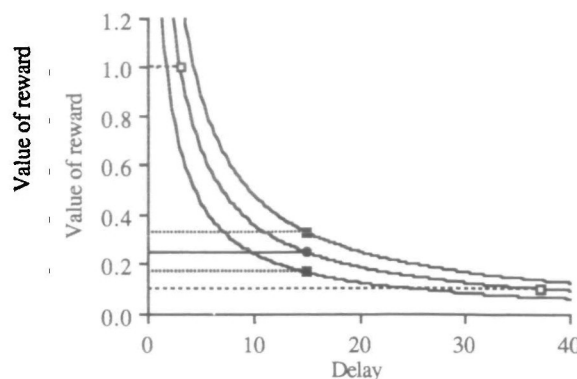
$$V = \frac{A}{1 + (K \cdot D)}$$

where V is the value of a reward delivered after a delay of D seconds, A is a measure of the amount of reinforcement, and K is a free parameter with dimension of time^{-1} that depends on the animal under study. If K is positive, then the value of the reward is a decreasing, negatively accelerated function of the delay between the choice and the reward. Mazur (1984, 1987) found that a pigeon's preference could be accommodated by values of K around 1 s^{-1}. The consequences for our situation are depicted in Figure 4, using this value of K. The figure uses the delays and rewards programmed in our experiment. The ordinate of Figure 4 shows the subjective value of each reward at the time of the choice. The value of the constant alternative is its projection at this axis, and the value of the variable alternative is the mean of the projections of the two possible outcomes. Because the equation is linear in A but convex in D, the mean value at the time of choice is higher for the variable delay option than for the constant option, and equal between the variable reward option and the constant option. The model is consistent with preference for variability in delays but does not predict any preference for variability in reward. In addition, the hypothesis is not related to any suggested underlying cognitive process. The value of the alternatives is presumably compared at the time of
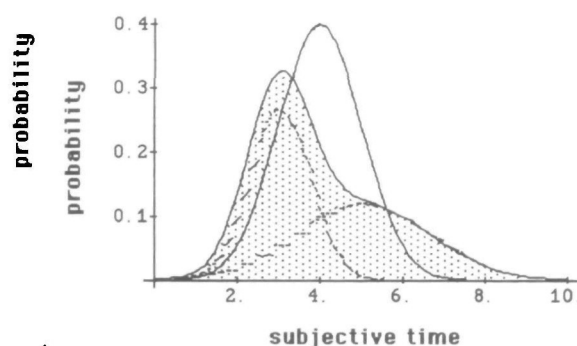
choice, whereas the delay is experienced at the time of the food reward, thus requiring some memory mediation, which is not fully specified.

Scalar timing

Our results can also be analyzed in terms of a theory of information processing for time. Scalar expectancy theory (Gibbon, 1977) is based on the principle that when an animal stores in its memory a time interval, it does so with a certain weight, or probability density, around the true experienced value. This probability distribution is assumed to be normally distributed around the real-time value of the interval, with a coefficient of variation characteristic of each subject but constant with changes in the length of the interval. According to this hypothesis, mixtures of intervals are stored in memory as aggregates of distributions around the values actually experienced. Figure 5 depicts this idea for the memory of two alternatives, one representing a constant interval and the other a mixture of equally frequent intervals to the left and right of the former. The memorized distribution of the constant option is bell shaped, but the memorized distribution of the mixture is the aggregate of two distributions, each bell shaped and with half the area of the constant option. Because the coefficient of variation is preserved, the distributions get flatter as their central values increase, and the aggregate of the two distributions for the variable side shows considerable skew. In a foraging context, scalar expectancy has been applied to sampling (Shettleworth et al., 1988), choice for delayed reinforcement in pigeons (Gibbon et al., 1988), and patch residence time (Brunner, 1990; Kacelnik et al., 1990). We assume that for each choice the subjects compare two values, one sample from the memory distribution for the fixed option and another sample from the aggregate memory for the variable distribution, and that they prefer the shorter delay to food or the longer food reward. Because of the skew in the memory for the variable mixture, the comparison of samples from memory will result in more frequent cases in which the sample from the mixed distribution is smaller than the sample from the constant alternative. In consequence, the subject ought to prefer more often the mixed distribution when it is choosing among delays to food and the constant distribution when it is choosing among food rewards of different durations. This is what we found, even though in our experiments the mean of the real-time mixture of delays was longer than the constant interval. Scalar timing seems to provide an elegant, process-based explanation for preference for

amount when the latter is generated by different times of exposures to food.

The best account of our results is thus provided by a model of preference for shorter delays or longer rewards under an information-processing constraint, consistent with a memory system with a constant coefficient of variation and a decision process based on single samples from the memory for each alternative. This model is for a form of constrained rate maximization, rather than a model of preference or avoidance of variability. Our model in its present form, however, does not predict shifts in risk sensitivity according to energy budget. These shifts have been reported in several species (viz. Barnard and Brown, 1985; Caraco, 1981, 1983; Caraco et al., 1980, 1990; Young et al., 1990 but see also Ha et al., 1990; Kagel et al., 1986b) and have been supported indirectly by our correlational evidence between subjects, implying that the scalar memory account cannot yet provide a comprehensive explanation.

and an anonymous referee for helpful discussion and comments on a previous version.

REFERENCES

Barnard CJ, Brown CAJ, 1985. Risk sensitivity foraging in common shrews (*Sorex araenus* L.). Behav Ecol Sociobiol 16:161–164.

Brunner D, 1990. Psychological mechanisms of optimal foraging (Doctoral dissertation). Cambridge: Cambridge University.

Caraco T, 1981. Energy budgets, risk and foraging preferences in dark-eyed juncos (*Junco hyemalis*). Behav Ecol Sociobiol 8:213–217.

Caraco T, 1983. White-crowned sparrows (*Zonotrichia leucophrys*) foraging preferences in a risky environment. Behav Ecol Sociobiol 12:63–69.

Caraco T, Blankenhorn WU, Gregory GM, Newman JA, Reger GM, Zwicker SM, 1990. Risk-sensitivity: ambient temperature affects foraging choice. Anim Behav 39:338–345.

Caraco T, Martindale S, Whittam TS, 1980. An empirical demonstration of risk-sensitivity foraging preferences. Anim Behav 28:820–830.

Cartar RV, Dill LM, 1990. Why are bumble bees risk sensitive foragers? Behav Ecol Sociobiol 26:121–127.

Davison M, 1969. Preference for mixed-interval vs. fixed interval schedules. J Exp Anal Behav 12:247–252.

Davison M, 1972. Preference for mixed-interval vs. fixed interval schedules: number of component intervals. J Exp Anal Behav 17:169–176.

Gibbon J, 1977. Scalar expectancy theory and Weber's law in animal timing. Psychol Rev 84:279–325.

Gibbon J, Church RM, Fairhurst S, Kacelnik A, 1988. Scalar expectancy theory and choice between delayed rewards. Psychol Rev 95:102–114.

Gilliam JF, Greene RF, Pearson NE, 1982. The fallacy of the traffic policeman: a response to Templeton and Lawlor. Am Nat 119:875–878.

Ha JC, 1991. Risk sensitive foraging: the role of ambient temperature and foraging time. Anim Behav 41:528–529.

Ha JC, Lehner PN, Faley SD, 1990. Risk-prone foraging behaviour in captive grey jays *Perisoreus canadensis*. Anim Behav 39:91–96.

Herrnstein RJ, 1964. Aperiodicity as a factor in choice. J Exp Anal Behav 7:179–182.

Kacelnik A, Brunner D, Gibbon J, 1990. Timing mechanisms in optimal foraging: some applications of scalar expectancy theory. In: Behavioural mechanisms of food selection (Hughes RN, ed). NATO ASI Series, Series G: Ecological Sciences, vol. G 20. Berlin: Springer-Verlag; 63–81.

Kagel JH, Green L, Caraco T, 1986a. When foragers discount the future: constraint or adaptation? Anim Behav 34:271–283.

Kagel JH, MacDonald D, Battalio R, White S, Green L, 1986b. Risk aversion in rats (*Rattus norvegicus*) under varying levels of resource availability. J Comp Psychol 100:95–100.

Killeen PR, 1968. On the measurement of reinforcement frequency in the study of preference. J Exp Anal Behav 11:263–269.

Killeen PR, Smith JP, Hanson SJ, 1981. Central place foraging in *Rattus norvegicus*. Anim Behav 29:64–70.

Leventhal AM, Morrell RF, Morgan EJ Jr, Perkins CC Jr, 1959. The relation between mean reward and mean reinforcement. J Exp Psychol 59:284–287.

Mazur JE, 1984. Test of an equivalence rule for fixed and variable reinforcer delays. J Exp Psychol Anim Behav Proc 10:426–436.

Mazur JE, 1986. Fixed and variable ratios and delays: further tests of an equivalence rule. J Exp Psychol Anim Behav Proc 12:116–124.

Mazur JE, 1987. Estimation of indifference points with an adjusting-delay procedure. J Exp Anal Behav 49:37–47.

Mazur JE, 1989. Theories of probabilistic reinforcement. J Exp Anal Behav 51:87–99:

McNamara JM, Houston AI, 1987. A general framework for understanding the effects of variability and interruptions on foraging behaviour. Acta Biotheor 36:3–22.

Pubols BH Jr, 1962. Constant versus variable delay of reinforcement. J Comp Physiol Psychol 55:52–56.

Real L, Caraco T, 1986. Risk and foraging in stochastic environments: theory and evidence. Annu Rev Ecol Syst 17:371–390.

Reboreda JC, Kacelnik A, in press. Reinforcement, autoshaping and cooperation in a two-player game between starlings. In: The nature of reinforcement. Quantitative analyses of behavior, vol. 13 (Commons M, ed). Hillsdale, New Jersey: Lawrence Erlbaum and Associates.

Shettleworth SJ, Krebs JR, Stephens DW, Gibbon J, 1988. Tracking a fluctuating environment: a study of sampling. Anim Behav 36:87–105.

Stephens DW, Krebs JR, 1986. Foraging theory. Princeton: Princeton University Press.

Templeton AR, Lawlor LR, 1981. The fallacy of averages in ecological optimization theory. Am Nat 117:390–393.

Turelli M, Gillespie JH, Schoener TW, 1982. The fallacy of the fallacy of the averages in ecological optimization theory. Am Nat 119:879–884.

Young RJ, Clayton H, Barnard CJ, 1990. Risk-sensitive foraging in bitterlings, *Rhodeus sericus*: effects of food requirement and breeding site quality. Anim Behav 40:288–297.

Zabludoff SD, Wecker J, Caraco T, 1988. Foraging choice in laboratory rats: constant vs. variable delay. Behav Proc 16:95–110.