Risky Theories—The Effects of Variance on Foraging Decisions[1]

ALEX KACELNIK AND MELISSA BATESON

*Department of Zoology, University of Oxford, South Parks Road,
Oxford OX1 3PS, United Kingdom*

SYNOPSIS   This paper concerns the response of foraging animals to variability in rate of gain, or risk. Both the empirical and theoretical literatures relevant to this issue are reviewed. The methodology and results from fifty-nine studies in which animals are required to choose between foraging options differing in the variances in the rate of gain available are tabulated. We found that when risk is generated by variability in the amount of reward, animals are most frequently risk-averse and sometimes indifferent to risk, although in some studies preference depends on energy budget. In contrast, when variability is in delay to reward, animals are universally risk-prone. A range of functional, descriptive and mechanistic accounts for these findings is described, none of which alone is capable of accommodating all aspects of the data. Risk-sensitive foraging theory provides the only currently available explanation for why energy budget should affect preference. An information-processing model that incorporates Weber's law provides the only general explanation for why animals should be risk-averse with variability in amount and risk-prone with delay. A theory based on the mechanisms of associative learning explains quantitative aspects of risk-proneness for delay; specifically why the delay between choice and reward should have a stronger impact on preference than delays between the reward and subsequent choice. It also explains why animals should appear to commit the "fallacy of the average," maximising the expected ratio of amount of reward over delay to reward when computing rates rather than the ratio of expected amount over expected delay. We conclude that only a fusion of functional and mechanistic thinking will lead to progress in the understanding of animal decision making.

INTRODUCTION

We begin with three examples to introduce the behavioral phenomenon we are seeking to understand. First, if bumblebees (*Bombus edwardsi*) are given a mixed array of two types of artificial flowers one of which provides a constant volume of 0.1 μl of nectar, and the other a variable volume with 1 μl on 10% of encounters and nothing on the other 90%, the bees show a strong preference for the constant type (Waddington *et al.*, 1981). In contrast, if pigeons (*Columba livia*) are given a choice of two keys, arranged such that pecks on one lead to a 15-s fixed delay to water and pecks on the other lead to a variable delay to water averaging 15 s, the birds prefer the variable delay key (Case *et al.*, 1995). And thirdly, when yellow-eyed juncos (*Junco phaeonotus*) are given a choice of two feeding stations, one of which offers four millet seeds on every visit and the other a variable number that is either one or seven seeds with 50% probability, the birds' preference depends on the ambient temperature: at 1°C they prefer the variable option, but at 19°C they prefer the constant option (Caraco *et al.*, 1990).

In all three cases these animals are faced with two options with apparently equal rates of gain but different variances. The variation is experimentally programmed to occur at the level of individual choices such

that it is impossible for a forager to know exactly what it will obtain if it chooses the variable option; all it can theoretically learn is the probability distribution of the possible outcomes of making such a choice. In the foraging literature this type of environmental stochasticity is usually referred to as risk. The word risk is also used sometimes in connection with danger of predation, but in the present context we use it exclusively to indicate variance. The above examples are not isolated cases, but are representative of a whole area of research by both psychologists and behavioural ecologists devoted to understanding the effects of risk on the foraging decisions made by animals. In general, if an animal is allowed to choose between two foraging options that differ in risk, the chances are that it will show a preference for one option over the other even if this means choosing the option offering a somewhat lower average rate of gain. In other words animals tend to be risk-sensitive.

From an evolutionary point of view these results are important because they challenge the basic assumptions of classical optimal foraging theory. In Charnov's (1976a, b) two seminal models it is assumed that the forager is adapted to a homogeneous habitat with stochastic but stable properties in which fitness is an increasing function of the rate of energy gain per unit of time spent foraging. Both of these models provide deterministic rules that cope with environmental variability by operating with only the averages of the relevant variables such as the average energy content of prey items, the average time taken to handle prey items or the average travel time between patches. However, the results described above make it clear that considering only average values of the variables involved in foraging decisions is insufficient, since if the foragers only had access to estimates of the average nectar volume, the average time delay to water or the average number of seeds they should have treated the fixed and variable options as equivalent.

Within behavioral ecology, the dominant evolutionary framework for explaining animals' responses to variance is risk-sensitive foraging theory. This body of theory depends on assuming that there is a non-linear relationship between the rate of gain associated with each prey item and fitness (for a review see McNamara and Houston, 1992; Smallwood, 1996). Given a non-linear function, $F(x)$, Jensen's inequality states that for a range of values of x the average of $F(x)$ is not equal to $F(\bar{x})$ where $\bar{x}$ is the average of the values of x. This produces risk-sensitivity as follows. If a forager is faced with two patches offering equal average rates of gain and if the function relating rate of gain to fitness is positively accelerated, then, because deviations above the average give greater fitness gains than deviations below the average, it pays to be risk-prone. Conversely if the function is negatively accelerated then by the same reasoning symmetrical deviations from the average cause fitness losses, and it pays to be risk-averse.

A paradigmatic risk-sensitive foraging model (Stephens, 1981) considers a small bird that has to reach a minimum threshold of reserves by nightfall if it is to survive until morning. This threshold provides the necessary non-linear relationship between rate of gain and fitness. As in the junco example, it is assumed that the bird is faced with two foraging options offering the same expected number of seeds with one fixed and the other variable. If the fixed option offers a rate of gain sufficiently high to take the bird above the threshold (i.e., it is on a positive energy budget) then it should be risk-averse, whereas if the rate is not sufficiently high (i.e., the bird is on a negative energy budget) then the bird's only chance of survival is to be risk-prone and gamble on the variable option giving an above average rate of gain. This has been summarised in the daily energy budget rule (Stephens, 1981) which states that a forager on a positive budget should be risk-averse but a forager on a negative budget risk-prone.

Our aim in this paper is to review some of the theoretical frameworks proposed to account for the phenomenon of risk sensitivity. While we show that risk-sensitive foraging theory has had some success, our primary aim is to demonstrate that there are a number of alternative frameworks that deserve equal attention. We bring together

several lines of thought based on functional, mechanistic and descriptive considerations that, while being clearly connected, have remained isolated from each other. In some cases we show that these theories make different predictions about how foragers should respond to risk. For example, they differ in their predictions relating to the effects of the energy budget of the subjects, which dimension is risky (since this can be either the amount of a commodity or the foraging time associated with obtaining it), and even details of the procedure such as how stimuli are used to signal the different options. Before launching into the theory, we therefore start by reviewing what is currently known about the responses of real animals to risky choices.

RISK-SENSITIVE FORAGING: A CRITICAL REVIEW OF THE EVIDENCE

Our aim in this section is to identify whether there are any consistent patterns in the responses of animals to risk. The review encompasses both the psychological and ecological literatures, however we have excluded the vast literature on human gambling behavior since we are doubtful whether some of the factors known to affect human decision making, such as for example the exact words used to frame a question, or the dress code of a casino (for examples see Kahneman and Tversky, 1984; Wagenaar, 1988), have any direct equivalents in animals. Specifically, we have chosen to review the results of experiments in which animals are given a choice between two foraging options that offer different levels of variance in the rate of gain available. In most cases the options are arranged such that the average rate is the same in both options (by rate we mean the currency of classical optimal foraging theory which is long-term rate, defined as the ratio of expected gain divided by expected foraging time), however in some experiments the averages may also differ in order that information can be obtained on how animals trade-off average rate of gain against variance in rate of gain. In these latter cases it is often possible to extrapolate whether the subjects would be risk-averse or risk-prone

were the average rate in the two options in fact the same.

As well as recording the preferences of the animals, studies were scored according to a number of variables that may influence the extent or direction of risk sensitivity. Table 1 includes the following information from each study: the species; the dimension that was variable or risky; the choice procedure employed; and finally, whether the study involved manipulation of the energetic (or water) budget of the subjects. We are particularly interested in whether the predictions of the energy budget rule described above are born out by experimental evidence since, as we will argue later, this prediction is currently unique to risk-sensitive foraging theory. There are many problems inherent in trying to extract general patterns from comparative data, and in this case it is particularly likely that the varying methodologies employed in the different studies affected how the animals behaved. We therefore begin by discussing two aspect of methodology that are likely to affect animals' responses to risk.

The dimension that is risky

Risk can be generated experimentally in a number of ways that can be broadly divided into varying either the amount of food available in a foraging option or the time delay associated with obtaining food. While both types of manipulation are capable of generating risk in the rate of gain associated with an option, it is important for both theoretical and practical reasons to differentiate between the two types of risk. From a theoretical point of view the distinction is important because risk-sensitive foraging theory makes different predictions for variable amounts and delays in some circumstances (see later), and also some of the alternative theories we consider differ in their treatment of the two sorts of variance.

From a practical point of view the means used to generate risk can affect the rate of gain experienced by the subjects. It is important to establish empirically that there is a linear relationship between the amount of food programmed by the experimenter and that taken by the animal. If we start by con-

sidering variability in amount of reward, Table 1 includes studies in which variability is in the number of similarly sized food items, the duration of access to a food hopper, and the volume or concentration of nectar delivered. Although many of the experiments claim that the mean amounts obtained from constant and variable options were equal, in the cases where amount is controlled by the length of access given to a food hopper this is rarely verified (e.g., Essock and Rees, 1974; Menlove et al., 1979). In fact, it has been demonstrated that the amount of grain consumed by pigeons feeding from a hopper is not always a linear function of the length of hopper access (Epstein, 1981). An exception is the study on starlings by Reboreda and Kacelnik (1991) in which the necessary linear relationship between duration of access and grams of food consumed is found over the range of durations they used in an experiment. As with amount, delay to obtaining food can also be controlled in a number of ways. In the language of operant psychology the studies in Table 1 use either interval, time or ratio schedules (operant terminology is explained in the Appendix). However, the programmed delay need not necessarily have a linear relationship to the delay the animal experiences. This is especially likely to be a problem with interval and ratio schedules in which the length of the delay experienced will depend on the rate at which the animal responds.

A second practical problem is that experiments that explicitly claim to vary amount almost invariably also vary the time associated with acquiring the food, because large rewards will often take longer to deliver, to handle and to consume than small ones. This source of variability is generally ignored, and can lead to apparent risk sensitivity in a long-term rate maximizer if there is a non-linear relationship between the size of a food item and the time it takes to handle. For example, Possingham et al. (1990) have shown that a bee maximising its long-term rate of intake should be risk-prone to variability in volume, given the observed increasing decelerating relationship between the volume of nectar taken from a flower and the time spent on the

flower (Hodges and Wolf, 1981). This effect of handling time will be most important in animals such as bees where it comprises a substantial portion of each foraging cycle. Effects of handling time are probably minimised in the experiments in which amount is varied by manipulating the concentration of a fixed volume of nectar (e.g., Wunderle and O'Brien, 1985; Banschbach and Waddington, 1994; Waddington, 1995), although even here there may be differences in the time taken to acquire or process nectar of different concentrations. It is notable that recent experiments on bees that use this technique to equalise the mean energy received from fixed and variable flowers have not found any risk sensitivity (Banschbach and Waddington, 1994; Waddington, 1995; Perez and Waddington, 1996).

A final remark relates to the pre-1990 studies of Caraco and his colleagues (Caraco et al., 1980; Caraco, 1981; Caraco, 1983) and others who have followed their methods (e.g., Moore and Simm, 1986). In an attempt to maintain experimental animals on the stated energy budget, the following inter-trial interval was made proportional to the size of the food reward just obtained, thus maintaining the intake rate of the animal at a set level regardless of the variance in the size of rewards. However, a secondary effect of this was that both delay to obtaining reward and amount of reward were variable in the variable option. Thus these experiments confound the effects of variability in amount and delay. Caraco et al. (1990) subsequently addressed this criticism by using ambient temperature rather than rate of intake to control the energy budgets of the subjects, thereby placing the animals on a positive or negative energy budget without changing the amount or timing of feedings.

In summary, it is important to notice how risk is generated in an experiment, since the means used can influence both the theoretical model that is appropriate to the problem and the long-term rate of intake experienced by the subjects. Some cases of apparent risk sensitivity may be explained simply as the result of long-term rate maximising.

The choice procedure and measure of preference

A variety of methods have been used to present animals with choices of two foraging options and measure their resulting preferences. The study by Cartar (1991) is unique in utilising naturally occurring differences in the variance in profitability of two wild flower species; all others rely on generating the reward distributions artificially. In the simplest technique, the subject is allowed to forage freely on an array of dishes or artificial flowers of two different types, generally identified by coloured markers. Preference is measured as the proportion of visits made to the flowers or dishes of each type. This procedure is often used with bees (Real, 1981; Waddington *et al.*, 1981), and seed-eating and nectarivorous birds (Caraco, 1982; Wunderle and O'Brien, 1985; Barkan, 1990; Tuttle *et al.*, 1990). It has the advantage of mimicking natural foraging well, but it allows limited control of the rate at which the subjects forage and the reward distributions they experience if they revisit depleted sites. These problems can be solved by using just two feeding stations, one for each option, that can be replenished after each visit (*e.g.*, Croy and Hughes, 1991; Banschbach and Waddington, 1994). A further extension involves scheduling the availability of food at the two patches according to a discrete trials procedure, in which there is an inter-trial interval (*e.g.*, Caraco *et al.*, 1980). Procedures of this type allow precise control of the animals' experience at the cost of reduced resemblance to natural foraging. A major advantage is that an animal can be forced to experience both options in trials in which no choice is given before those trials in which choice is measured. The development of side preferences, in which an animal repeatedly uses one feeding station independently of the reward schedule it offers, is however often a problem (Ha *et al.*, 1990; Ha, 1991; Reboreda and Kacelnik, 1991).

Many choice procedures are implemented using operant techniques; birds can be trained to peck keys or hop on perches for food rewards, and mammals such as rats

can be trained to press levers. Choice experiments usually involve two operanda, and the animals are taught that either the position of an operandum or a particular cue, such as the colour of light projected onto a pecking key, is associated with a particular distribution of rewards. Different schedules of reinforcement can be programmed on the operanda. In schedules where the animal makes a choice with a single response, preference is measured as the proportion of times each of the options is chosen. Sometimes additional measures that may correlate with preference are also taken, such as the latency to accept each option in forced trials or pecking rate during a delay (Reboreda and Kacelnik, 1991; Bateson and Kacelnik, 1995a). A number of studies have used titration procedures in which either the amount of food or the delay in the fixed option is adjusted until a point is found where the subjects are indifferent between the fixed and variable option (*e.g.*, Mazur, 1984; Mazur, 1986; Bateson and Kacelnik, 1996). This technique has the advantage of providing a quantitative measure of the value subjects attribute to the variable option. It may also provide a more sensitive measure of preference.

Apart from the discrete trials procedure mentioned above in which the subjects may have to peck a key or press a lever a single time to register their choice, there are two other types of operant choice procedure that are frequently used: concurrent schedules and concurrent-chains schedules. In concurrent schedules, a separate schedule of reinforcement is independently programmed on each operandum, and the subject is allowed to switch between the two, being rewarded whenever the requirements of the schedule on either of the operanda are met. Preference is measured as the proportion of the total time or total responses allocated to each option. Concurrent schedules can be used to test preference for fixed versus variable amounts by making the two schedules identical apart from the amount of reinforcement given (*e.g.*, Hamm and Shettleworth, 1987). Alternatively, if the delay to reward in one of the schedules is fixed and in the other it is mixed or variable, the technique can be used to measure preference for

variability in delay or number of responses required to obtain reinforcement. For example, Ha *et al.,* used concurrent variable ratio schedules, one of which had a higher variance than the other, to measure preference for variability in the number of responses required per reward (Ha *et al.,* 1990; Ha, 1991).

The concurrent-chains procedure involves an initial link in which two identical ratio or interval schedules are concurrently available on two operanda. When the requirement on one of these schedules is met, the other operandum becomes unavailable and the subject has to complete the terminal link schedule on the current operandum. For risk-sensitivity experiments the terminal links are generally a fixed-interval versus a mixed-interval or variable-interval (Herrnstein, 1964; Davison, 1969; Davison, 1972), a fixed-ratio versus a mixed- or variable-ratio (Fantino, 1967; Morris, 1986; Ahearn and Hineline, 1992), or a fixed-time versus a mixed- or variable-time (Rider, 1983*a*; Zabludoff *et al.,* 1988). Preference is measured as the proportion of initial link responses made on each operandum, thus giving a continuous measure of the strength of preference for the fixed and variable terminal links in each trial.

There is some evidence to suggest that the preference obtained from an experiment is likely to be affected by the type of choice procedures and index of preference used. For example, Caraco (1982) compared the results obtained from white crowned sparrows choosing between fixed and variable sources of seeds in an experiment with discrete choice trials programmed at two feeding stations with the results from another experiment in which the birds foraged freely from an array of dishes covered with differently coloured paper discs. He found that in both experiments the birds were risk-averse, but for similar variance to mean ratios this preference was more extreme in the discrete trials paradigm. Hamm and Shettleworth (1987) argue on the basis of their results that preferences are also likely to be more extreme when measured with a discrete trials procedure than with a concurrent schedule. However this claim is difficult to assess from their data because factors other

than just the choice procedure, such as the variance of the variable option, were changed between their two experiments. Rider (1983*b*) tested preference for the same fixed-ratio (FR50) versus mixed-ratio (MR1/99) schedule using both a concurrent procedure and a concurrent-chains procedure. He found strong risk aversion in the former experiment and risk proneness in the latter. However, the risk aversion in the concurrent schedules experiment is easily explained because the rat only has to make a single response to determine whether the mixed-ratio is currently 1 or 99, and if it is 99 it can switch to responding solely on the fixed-ratio lever. This experiment demonstrates that the results obtained from concurrent schedules need to be interpreted with caution.

General patterns in risk-sensitive foraging

Having highlighted some of the possible caveats in the interpretation of the available evidence we proceed with the results of the studies. Fifty-nine papers are summarised in Table 1, and while this is not an exhaustive list we believe it is a fair representation of the literature. The table includes studies on a total of 28 species consisting of 10 insects (all bees and wasps), two fish, 12 birds and four mammals. Risk-sensitive preferences are observed in all of these taxonomic groups. However, in the light of our earlier comments regarding the studies on bees, it is possible that some of the risk sensitivity reported in this group could in fact be explained as an outcome of unequal average rates of gain in the two options. All studies involve artificial manipulation of either the subjects or their food sources; only one study uses wild unmanipulated animals living in their natural environment (black-capped chickadees, Barkan, 1990), and only one study uses a natural unmanipulated food source (two species of wild flowers, Cartar, 1991). Thirty-eight studies investigate the response of animals to variability in amount of food, 18 to variability in delay to food, and three studies compare the two types of variability. Studies differ in whether or not they investigate the effects of manipulating energy budgets; we shall start by

TABLE 1. *Summary of experiments investigating animals' responses to variability in their food sources. Studies are divided into those on insects, fish, birds and mammals.*

Species	Quantity variable or risky, and nature of reward	Choice procedure
Insects		
Paper wasp (*Vespula vulgaris*).	Volume of 40% sucrose	Free foraging on array of artificial flowers
Bumblebee (*Bombus sandersoni*)	Volume of 40% sucrose	Free foraging on array of artificial flowers
Bumblebee (*Bombus edwardsi*)	Volume of 33% sucrose	Free foraging on array of artificial flowers
Bumblebee (*Bombus pennsylvanicus*)	Volume of sucrose soln	Free foraging on array of artificial flowers
Bumblebee (*Bombus occidentalis*)	Volume of sucrose soln NB Unequal means	Artificial flowers
Bumblebee (*Bombus melanopygus, B. mixtus and B. sitkensis*)	Net profitability, however, most variance originates from nectar volume	Free foraging on unmanipulated naturally occurring flowers (seablush and dwarf huckleberry)
Honey-bee (*Apis mellifera ligustica*)	Sucrose concentration	Discrete choices between two pairs of artificial flowers
Bumblebee (*Bombus fervidus*)	Sucrose concentration	Discrete choices between two pairs of artificial flowers
Fish		
Bitterling (*Rhoderus sericus*)	Number of trout pellets	Discrete choices of two feeding patches
Fifteen-spined stickleback (*Spinachia spinachia*)	Number of prey items (*Artemia*)	Free foraging at two feeding patches
Birds		
Blue jay (*Cyanocitta cristata*)	Number of mealworm halves	Discrete choice trials
Grey jay (*Persoreus canadensis*)	Ratio to obtain 97-mg cat pellets	Concurrent (VR, VR)
Grey jay (*Persoreus canadensis*)	Ratio to obtain 97-mg cat pellets	Concurrent (VR, VR)
White-crowned sparrow (*Zonotrichia leucophrys*)	Number of millet seeds	Discrete choices of two feeding stations
White-crowned sparrow (*Zonotrichia leucophrys*)	Number of millet seeds	Free foraging on array of food dishes
White-crowned sparrow (*Zonotrichia leucophrys*)	Number of millet seeds	Discrete choices of two feeding stations
White-throated sparrow (*Zonotrichia albicollis*)	Number of thistle seeds	Free foraging on array of food dishes
Dark-eyed junco (*Junco hyemalis*)	Number of millet seeds	Discrete choices of two feeding stations
Dark-eyed junco (*Junco hyemalis*)	Number of millet seeds	Discrete choices of two feeding stations
Yellow eyed junco (*Junco phaeonotus*)	Number of millet seeds	Discrete choices of two feeding stations
Yellow eyed junco (*Junco phaeonotus*)	Number of millet seeds	Discrete choices of two feeding stations
Yellow-rumped warbler (*Dendroica coronata*)	Number of *Tenebrio* larvae	Discrete choices of two feeding stations

TABLE 1. *Extended.*

Manipulation of energetic status of subjects	Results: response to risk	Reference (numbers are used for identification in Figures 1 and 2)
No manipulation	Risk-averse	Real (1981)—1
No manipulation	Risk-averse	Real (1981)—1
No manipulation	Risk-averse	Waddington et al. (1981)—2
No manipulation	Risk-averse	Real et al. (1982)—3
Sucrose added to or removed from colony reserves	Risk-averse when reserves enhanced, but risk-prone/no preference when depleted	Cartar and Dill (1990)—4
Sucrose added to or removed from colony reserves	No preference when reserves enhanced, but risk-prone when depleted	Cartar (1991)—5
Honey combs removed from or added to colony	No preference in either treatment	Banschbach and Waddington (1994)—6
No manipulation	No preference	Waddington (1995)—7
Prior deprivation and feeding rate manipulated	Risk-averse with short deprivation but risk-prone with long deprivation	Young et al. (1990)—8
Prior deprivation or satiation manipulated	Risk-averse when sated but risk-prone when hungry	Croy and Hughes (1991)—9
No manipulation Maintained at 80% free feeding weights	Risk-averse	Clements (1990)—10
Mean ratio requirement manipulated under closed economy	Risk-prone independent of treatment	Ha et al. (1990)—11
As above plus manipulation of day length and ambient temperature	Risk-prone independent of treatment	Ha (1991)—12
Maintained on positive budgets	Risk-averse	Caraco (1982)—13
Maintained on positive budgets	Risk-averse	Caraco (1982)—13
Prior deprivation and inter-trial interval manipulated	Risk-averse on positive budgets and risk-prone on negative budgets	Caraco (1983)—14
Maintained on positive budgets	Risk-averse	Tuttle et al. (1990)—15
Prior deprivation and inter-trial interval manipulated	Risk-averse on positive budgets, no preference on balanced budgets and risk-prone on negative budgets	Caraco (1981)—16
Maintained on positive budgets	Risk-averse	Caraco and Lima (1985)—17
Ambient temperature manipulated	Risk-averse on positive budgets and risk-prone/no preference on negative budgets	Caraco et al. (1990)—18
Prior deprivation and inter-trial interval manipulated	Risk-averse on positive budgets and risk-prone on negative budgets	Caraco et al. (1980)—19
Day length used to manipulate premigratory state	Pre-migratory birds risk-averse and postmigratory risk-prone	Moore and Simm (1986)—20

TABLE 1. *Continued.*

Species	Quantity variable or risky, and nature of reward	Choice procedure
European starling (*Sturnus vulgari.*)	Length of access to hopper with turkey crumbs (amount)	Discrete choice trials scheduled on two keys (FI, FI)
European starling (*Sturnus vulgaris*)	Interval to turkey crumbs	Discrete choice trials (FI, VI)
European starling (*Sturnus vulgaris*)	Number of units of turkey crumbs	Discrete choice trials scheduled on two keys (FI, FI)
European starling (*Sturnus vulgaris*)	Interval to turkey crumbs	Discrete choice trials scheduled on two keys (FI, VI)
European starling (*Sturnus vulgaris*)	Interval to turkey crumbs NB unequal means	Discrete choice trials scheduled on two keys (FI, MI) with adjusting FI
European starling (*Sturnus vulgaris*)	Interval to turkey crumbs NB unequal means	Discrete choice trials scheduled on two keys (FI, VI) with adjusting FI
Bananaquit (*Coereba flaveola*	Volume of 30% sucrose soln	Free foraging on array of artificial flowers
Bananaquit (*Coereba flaveola*)	Concentration of 10 µl sucrose soln	Free foraging on array of artificial flowers
Bananaquit (*Coereba flaveola*)	Volume of 0.47 M sucrose soln	Free foraging on array of artificial flowers
Bananaquit (*Coereba flaveola*)	Volume of sucrose soln	Free foraging on array of artificial flowers
Great tit (*Parus major*)	Number of fly pupae	Discrete choices of two feeding stations
Black-capped chickadee (*Parus atricapillus*)	Number of sunflower seeds	Wild birds freely foraging on array of holes in logs
Pigeon (*Columba livia*)	Interval to hopper access	Concurrent chains scheduled on two keys (FI, VI)
Pigeon (*Columba livia*)	Amount of hopper access	Concurrent (VI, VI) scheduled on two keys
Pigeon (*Columba livia*)	Amount of hopper access	Concurrent (FR, FR) scheduled on two keys
Pigeon (*Columba livia*)	Ratio to obtain mixed grain	Discrete choice trials (FR, MR) scheduled on two keys
Pigeon (*Columba livia*)	Number of pellets	Discrete choice trials scheduled on two keys
Pigeon (*Columba livia*)	Ratio to obtain reward	Concurrent chains (FR, VR) scheduled on two keys
Pigeon (*Columba livia*)	Interval to hopper access	Concurrent (FI, MI)
Pigeon (*Columba livia*)	Interval to reward	Concurrent chains (FI, MI)
Pigeon (*Columba livia*)	Interval to hopper access	Concurrent chains (FI, MI)
Pigeon (*Columba livia*)	Ratio to obtain hopper access	Concurrent chains (FR, MR)
Pigeon (*Columba livia*)	Number of 75-mg pellets	Concurrent (VI, VI)
Pigeon (*Columba livia*)	Number of 20-mg pellets	Discrete choice trials
Pigeon (*Columba livia*)	Amount of hopper access	Concurrent chains (FI, FI)
Pigeon (*Columba livia*)	Interval to scoop of water (0.3 or 1.00 ml depending on condition)	Concurrent chains (FI, VI)
Pigeon (*Columba livia*)	Time to 3-s hopper access NB unequal means	Discrete choice trials (FT, MT), (FT, VT) or (FT, RT) Adjusting procedure

TABLE 1. *Extended. Continued.*

Manipulation of energetic status of subjects	Results: response to risk	Reference (numbers are used for identification in Figures 1 and 2)
No intended manipulation, however subjects differed in their ability to exploit hopper access time	Risk-averse/no. preference Correlation betwen aversion and amount of food obtained across subjects	Reboreda and Kacelnik (1991)—21
No intended manipulation, however subjects differed in their ability to exploit hopper access time	Risk-prone. Negative correlation between proneness and amount of food across subjects	Reboreda and Kacelnik (1991)—21
No manipulation	Risk-averse/no preference	Bateson and Kacelnik (1995a)—22
No manipulation	Risk-prone	Bateson and Kacelnik (1995a)—22
No manipulation	Risk-prone	Bateson and Kacelnik (1996)—23
Number of units of turkey crumbs manipulated	Risk-prone on both positive and negative budgets	Bateson and Kacelnik (in press)—24
No manipulation	Risk-averse/no preference	Wunderle and O'Brien (1985)—25
No manipulation	No preference	Wunderle and O'Brien (1985)—25
No manipulation Budgets probably negative	Risk-averse	Wunderle et al. (1987)—26
No manipulation Budgets neutral	Risk-averse	Wunderle and Cotto-Navarro (1988)—27
No manipulation Budgets probably positive	Risk-averse	Stephens and Ydenberg (1982)—28
Rate of intake manipulated, but budgets always positive	Risk-averse	Barkan (1990)—29
Maintained at 80% free feeding weights	Risk-prone	Herrnstein (1964)—30
Maintained at 80% free feeding weights	No preference	Staddon and Innis (1966)—31
Maintained at 80% free feeding weights	Risk-prone/no preference	Essock and Rees (1974)—32
Maintained at 75–80% free feeding weights	Risk-prone	Morris (1986)—33
Maintained at 80% free feeding weights	Risk-prone	Young (1981)—34
Maintained at 80% free feeding weights	Risk-prone	Ahearn et al. (1992)—35
Maintained at 80% free feeding weights	Risk-prone	Cicerone (1976)—36
Maintained at 80% free feeding weights	Risk-prone	Davison (1969)—37
Maintained at 80% free feeding weights	Risk-prone	Davison (1972)—38
Maintained at 80% free feeding weights	Risk-prone	Fantino (1967)—39
Mean length of VI manipulated	No preference in all treatments	Hamm and Shettleworth (1987)—40
Maintained at 80% free feeding weights	Risk-prone	Mazur (1984, 1986)—43, 44
Maintained at 80% free feeding weights	Risk-averse	Menlove et al. (1979)—41
Prior state, volume of rewards and length of sessions all manipulated	Risk-prone in all treatments	Case et al. (1995)—42
Maintained at 80% free feeding weights	Risk-prone	Mazur (1984, 1986)—43, 44

TABLE 1. *Continued.*

Species	Quantity variable or risky, and nature of reward	Choice procedure
Pigeon (*Columba livia*)	Amount of hopper access NB unequal means	Discrete choice trials Adjusting procedure
Pigeon (*Columba livia*)	Amount of hopper access NB unequal means	Discrete choice trials Adjusting procedure
Mammals		
Rat (*Rattus norvegicus*)	Number of units (0.24 g) of food or number of drops of water	Discrete choices in E-maze
Rat (*Rattus norvegicus*)	Time to 45-mg pellets	Discrete choices in Y-maze (FT, MT)
Rat (*Rattus norvegicus*)	Time to 45-mg pellets	Discrete choices in maze (FT, MT)
Rat (*Rattus norvegicus*)	Number of 45-mg pellets	Discrete choices in maze
Rat (*Rattus norvegicus*)	Time to 45-mg pellets	Concurrent chains (FT, MT)
Rat (*Rattus norvegicus*)	Ratio to obtain 45-mg pellets	Concurrent (FR, MR)
Rat (*Rattus norvegicus*)	Ratio to obtain 45-mg pellets	Concurrent chains (FR, MR)
Rat (*Rattus norvegicus*)	Number of food pellets	Discrete choice trials (FR, FR)
Rat (*Rattus norvegicus*)	Number of scoops of water	Discrete choice trials
Rat (*Rattus norvegicus*)	Time to 45-mg pellets	Concurrent chains (FT, MT)
Rat (*Rattus norvegicus*)	Number of 45-mg NB unequal means	Discrete choice trials on two levers
Rat (*Rattus norvegicus*)	Number 45-mg pellets NB unequal means	Discrete choice trials on two levers
Rhesus monkey (*Macaca mulatta*)	Number of popcorn kernels	Discrete choice trials
Common shrew (*Sorex araneus*)	Number of mealworm segments	Discrete choice trials
Common shrew (*Sorex araneus*)	Number of mealworm segments	Discrete choice trials
Round-eared elephant shrew (*Macroscelides proboscideus*)	Number of mealworms	Discrete choice trials from two trays

summarising the results of the studies that did not.

Studies that did not manipulate energy budgets

We list 22 studies in which the response of animals to variability in amount is investigated with no manipulation of energy budgets. A range of different responses to variability is observed: the majority of these studies report risk aversion, while some report results close to indifference, and a few report risk proneness (Fig. 1). It is possible that some of this variation could be attributed to the energetic status of the subjects in line with the predictions of the energy budget rule. Of the risk-averse animals some were certainly on positive budgets

TABLE 1. *Extended. Continued.*

Manipulation of energetic status of subjects	Results: response to risk	Reference (numbers are used for identification in Figures 1 and 2)
Maintained at 80% free feeding weights	Risk-prone	Mazur (1985)—45
Maintained at 80% free feeding weights	No preference	Mazur (1989) Experiment 1.—46
Rate of intake manipulated	No preference with low rate of intake, but risk-prone with high rate	Leventhal *et al.* (1959)—47
No manipulation, but deprived before sessions	Risk-prone	Pubols (1962)—48
No manipulation	Risk-prone	Logan (1965)—49
No manipulation	Risk-averse/no preference	Logan (1965)—49
No manipulation	Risk-prone	Rider (1983a)—50
Maintained at 80% free feeding weights	Risk-averse, but see text for why this is misleading	Rider (1983b)—51
Maintained at 80% free feeding weights	Risk-prone	Rider (1983b)—51
Number of forced trials before choices manipulated	Risk-averse in all treatments	Battalio *et al.* (1985)—52
Number of trials in a session manipulated	Risk-averse on positive and balanced budgets, no preference on negative budgets	Kagel *et al.* (1986b)—53
Body weight and mean time delay manipulated	Initially risk-averse, but subsequently no preference/risk-prone when 85% free-feeding weight, but risk-prone/no preference when below 85% NB changes in body weight were confounded with changes in the variance of the variable option	Zabludoff *et al.* (1988)—54
Number of trials per session manipulated	Increasing risk-aversion with fewer trials	Hastjarjo *et al.* (1990)—55
Number of pellets per reward manipulated	Increasing risk-aversion with larger rewards	Hastjarjo *et al.* (1990)—55
No manipulation	Risk-averse/no preference	Behar (1961)—56
Length of ITI manipulated	Risk-averse on positive budgets and risk-prone on negative budgets	Barnard and Brown (1985)—57
No manipulation Subjects probably on negative budgets	Risk-prone/no preference	Barnard *et al.* (1985)—58
Prior deprivation, ITI and ambient temperature all manipulated	Risk-averse/no preference on positive budgets, and no preference on negative budgets	Lawes and Perrin (1995)—59

(*e.g.,* Caraco, 1982; Tuttle *et al.*, 1990; Caraco and Lima, 1985; Stephens and Ydenberg, 1982), however in one study the subjects were probably on negative energy budgets (Wunderle *et al.*, 1987). Of the risk-prone animals, Barnard *et al.*'s (1985) shrews were probably on negative energy budgets, and the other three studies (Essock

and Rees, 1974; Young, 1981; Mazur, 1985) were on pigeons maintained at 80% of their free-feeding weights that could also have been on negative budgets. However, the difficulty of trying to explain the variation seen is exemplified by the studies on pigeons: despite the fact that similar procedures were employed and all the birds
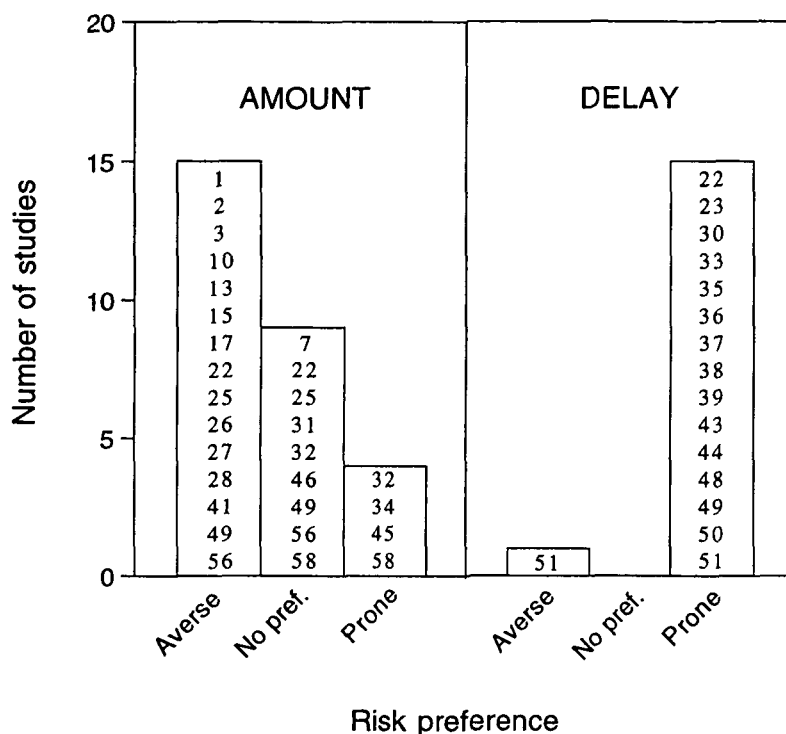
FIG. 1. The effect of the variable dimension (amount or delay) on the direction of risk-sensitive preferences. The data are from those studies in Table 1 that did not manipulate the energy budgets of the subjects. In cases where a study is ambiguous in its findings (*e.g.*, "risk-averse/no preference" in Table 1) both outcomes are included. The numbers within the bars indicate which studies in Table 1 contribute to the bar.

were maintained at 80 or 85% of their free-feeding weights, two experiments found risk aversion (Menlove *et al.*, 1979; Hamm and Shettleworth, 1987), three indifference (Staddon and Innis, 1966; Essock and Rees, 1974; Hamm and Shettleworth, 1987) and three risk proneness (Essock and Rees, 1974; Young, 1981; Mazur, 1985). This variation can not be explained by any obvious differences in the schedules used.

We list 15 studies in which the effects of variability in delay to obtain reward were investigated with no manipulation of energy budgets (Fig. 1). The results show that with interval, time and ratio schedules, animals are almost universally risk-prone (the only exception being the concurrent schedule study of Rider (1983*b*) discussed earlier, in which a misleading measure of preference was used). Given that the majority of these studies were carried out by psychologists using pigeons maintained at as low as 75%

of their free-feeding body weights, it is possible that the subjects may have been on negative energy budgets. However, this seems unlikely because pigeons are relatively large birds that can be maintained on 75–80% of their free-feeding weights for long periods of time, and it is not clear how these levels of deprivation relate to natural body weights and requirements in the wild. Also since the majority of the daily ration is generally received in the experiment, the rate of intake experienced must be sufficient to result in a positive energy budget.

Three studies (Logan, 1965; Reboreda and Kacelnik, 1991; Bateson and Kacelnik, 1995*a*) have compared responses to variability in amount and delay. All have found that animals (rats and starlings) are risk-averse when variability is in amount but risk-prone when variability is in delay, supporting the general trend observed between studies (Fig. 1).

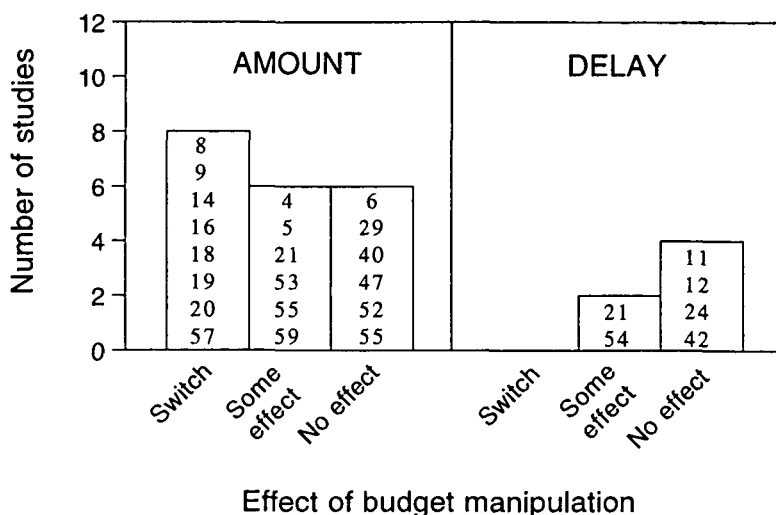|	AMOUNT			DELAY		
8						
9						
14	4	6				
16	5	29				
18	21	40			11	
19	53	47			12	
20	55	52		21	24	
57	59	55		54	42	

Effect of budget manipulation

FIG. 2.   The size of the effect of budget manipulations. The data are from those studies in Table 1 that manipulated the energy budgets of the subjects. In cases where a study is ambiguous in its findings both outcomes are included. The numbers within the bars indicate which studies in Table 1 contribute to the bar.

Effects of energy budget manipulations

We list 24 studies that manipulated the energy budgets of the subjects: 18 of these investigated risk-sensitive preferences when variability was in amount, five when variability was in delay and one study when variability was in both dimensions (Fig. 2). Of the experiments with variability in amount, 14 found some evidence for a shift toward risk-proneness when energy budgets were reduced and towards risk aversion when they were increased, although only eight of these studies showed a complete switch in preference between significant risk proneness and significant risk aversion. Despite the apparent level of support for the energy budget rule shown by the literature, we suspect that the real number of failures to obtain the predicted switch in preference with budget is actually greater, since studies that fail to reject the null hypothesis of lack of effect of budget are bound to be harder to publish and are probably often not submitted at all. Some failures may also be salvaged for publication if a convincing post hoc argument can be constructed for why a switch in preference would not have been predicted. A possible example is Barkan's (1990) study involving treatments with two very different rates of intake, in which he justifies the continuing risk aversion of the

birds on the basis of careful calculations showing that they were always on positive budgets.

Of the few experiments with variability in delay, all have met with failure to demonstrate convincing shifts in preference. Ha and colleagues (Ha *et al.,* 1990; Ha, 1991) tried without success to induce a switch in preference in gray jays as did Bateson and Kacelnik (in press) in starlings. Similar failure was met by Case *et al.* (1995) using water as a reinforcer for pigeons and manipulating water budgets. A study on rats by Zabludoff *et al.* (1988) found some evidence for a switch in preference. The rats became risk-prone as their body weights were reduced from 85% to 75% of their free-feeding weights. However, this experiment is difficult to interpret because the decrease in body weight is confounded with an increase in the variance of the more variable option, and other studies have shown increasing risk proneness with increasing variance in delay to reward. More convincingly, Reboreda and Kacelnik (1991) report a correlation between decreasing risk proneness and increasing efficiency at extracting food during a period of hopper access in different individual starlings; however this was not manipulated experimentally. Thus, there is little direct evidence

that energy budget affects risk sensitivity when variability is in delay to receiving food. Instead, risk proneness for delay seems universal.

There is some suggestion in the birds and mammals that body weight might explain which species respond to budget manipulations. Of the studies that did not find evidence for the predicted shift most were on larger bird species such as pigeons (Hamm and Shettleworth, 1987; Case *et al.*, 1995), jays (Ha *et al.*, 1990; Ha 1991) and starlings (Bateson and Kacelnik, in press) or rats (Leventhal *et al.*, 1959; Battalio *et al.*, 1985; Kagel *et al.*, 1986*b*; Hastjarjo *et al.*, 1990). It is possible that light species with few reserves might be more likely to have been subject to selection for short-fall minimisation and thus energy budget-associated switches in risk-sensitivity. In an attempt to control for phylogenetic confounds such as differences in metabolic rate between birds and mammals, we analysed the relationship between body weight and the effects of budget on risk sensitivity across the bird species only, and found that changes in budget were more likely to result in appropriate switches in foraging preferences in small birds (Fig. 3). However, a closer inspection of the data reveals that all of the studies with variability in delay have been conducted on larger species such as pigeons, jays and starlings. Thus it is not clear at present whether the lack of an effect of budget in these studies is due to the dimension that was varied or the body weights of the subjects since these two variables are confounded.

To summarise our conclusions:

● There are enough well controlled studies for us to conclude that risk sensitivity is a real phenomenon. In a few cases (the bees in particular) the animals may have experienced unequal average rates. However this does not apply in all studies.

● There is a general difference in animals' response to variability in amount and in the delay to reward: animals are more often risk-averse when variability is in amount but risk-prone when variability is in delay.

● There is growing support for energy bud-
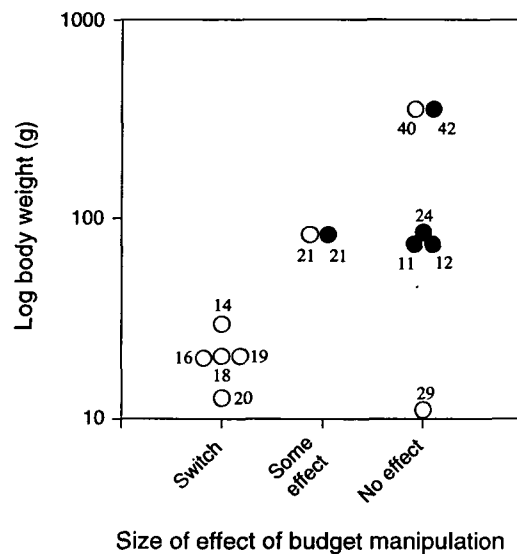

Size of effect of budget manipulation

FIG. 3. The size of the response of bird species to a manipulation of budget plotted against body weight. 'Switch' indicates a switch in preference as predicted by the energy budget rule, and 'Some effect' indicates a change in the right direction but not a complete switch in preference, 'No effect' indicates either no effect of budget or an effect in the opposite direction to that predicted. Body weights were taken from Dunning (1993). The numbers refer to the studies in Table 1. The empty circles indicate studies that investigated variability in amount the filled circles variability in delay.

get having some role in determining the direction of risk-sensitive preferences when variability is in amount: animals on positive energy budgets are often more risk-averse and animals on negative budgets more risk-prone.

● There is little evidence for a comparable effect when variability is in delay. However, experiments with variability in delay have not been performed in the same (low body weight) species that show a switch in preference with variability in amount.

PROBLEMS WITH TESTING RISK-SENSITIVE FORAGING THEORY

It is clear that the energy budget rule is insufficient to explain the patterns of risk-sensitivity in the literature. However, despite the fact that the majority of experimental tests have focused on the energy budget rule, this rule is not a universal pre-

diction of all risk-sensitive models. Stephens' original risk-sensitive foraging model has spawned a number of variants that explore the effects of modifying various of his original assumptions. It is not our intention to give an exhaustive review of the theory here, since this has been done elsewhere (*e.g.*, see McNamara and Houston, 1992), but rather to give a taste of the current level of sophistication of the theory and complexity of its predictions.

One of the most important constraints in Stephens' model lies in the fact that the forager is only allowed to make a single foraging decision, and is then required to stick with this for the remainder of the day. This led to the criticism that risk proneness would be very rare because it would only occur when the forager had a probability of dying of over 50% that day (Krebs *et al.*, 1983). However, if the forager is allowed to make sequential decisions that can vary according its current state, which may be affected by the outcome of previous decisions, then risk proneness becomes far less dangerous and therefore more likely. This is because a risk-prone forager that has a run of good luck that takes it back onto a positive trajectory can switch to risk aversion rather than chancing the possibility that it will drift back below the budget line (Houston and McNamara, 1982; McNamara, 1983; McNamara, 1984). Whether a single or sequential choice model is more realistic for a given foraging situation will depend on the degree to which an animal commits itself when it makes a choice. Single choice models may be more appropriate to large-scale patch choice decisions, whereas sequential choice models will be appropriate to modelling prey choice within a patch.

A second assumption of Stephens' model is that the only way to die is by failing to meet the critical level of reserves by nightfall. There is no possibility of starving while foraging, which for small mammals such as the shrew is a very real danger. If the model is modified so that a forager is assumed to forage continuously to stay above the lethal limit of reserves, and if the mean net gain is positive, then the optimal policy is always to be risk averse (McNa-

mara and Houston, 1992). The introduction of unpredictable interruptions to foraging such as bad weather or presence of a predator can further change this prediction. For example, Barnard *et al.* (1985) have shown that the need to insure against the possibility of interruptions makes risk proneness optimal at intermediate levels of reserves. Continuous foraging models are further modified if a constraint is placed on the maximum level of reserves a forager can store (McNamara and Houston, 1990). Houston and McNamara (1985) have combined the possibility of death by falling below a lethal boundary while foraging and the need to build up reserves to survive the night in a single model. In this case the optimal policy is a compromise been those in the separate models: to be risk-averse at all level of reserves except for a wedge shaped region in the reserves-time space near to dusk when it is optimal for an animal below the budget line to be risk-prone in order to get above the critical level of reserves.

A third limitation of Stephens' original model, and in fact all of the others mentioned so far, is that the optimality criterion has been restricted to maximization of probability of survival. There may be animals, particularly insects, that have the option of using energy acquired for immediate reproduction. McNamara *et al.* (1991) have used a model in which reproduction occurs above a certain level of reserves, and results in a reduction in reserves, to show that the policy that maximises lifetime reproductive success is different from one that minimises mortality. In the reproduction models risk-proneness can occur at high levels of reserves since there are conditions under which a risk-prone decision could take a forager over a threshold above which it can reproduce (see also Bednekoff, 1996). Given certain parameters, it can be shown that as reserves increase, the optimal policy changes successively from risk aversion (to escape an immediate lethal boundary) to risk proneness (to have a chance of meeting a daily requirement) back to risk aversion (when the requirement can be reached but there is no chance of reproduction) and finally back to risk proneness again (when this offers the possibility of reproduction).

A final deficiency with the models discussed so far is that they only consider variance in amount of food. Zabludoff *et al.* (1988) have modelled choice between a foraging option in which there is a fixed delay to food and one in which the delay has the same mean but is variable. They show that if the forager must reach a critical level of reserves by nightfall, then the energy budget rule describes the optimal policy in this situation. However, McNamara and Houston (1987) have modelled the effect of both variability in amount and delay on foraging preferences. They considered an animal foraging within a finite time horizon (a day for example), by the end of which it must have exceeded a certain critical level of reserves to survive. They concluded that when variability is in amount the animal should be risk-averse if its current level of reserves is sufficient for it to be above the critical threshold at dusk, and risk-prone if they are below. However, for delay the picture is more complex, and there are typically four regions in the reserves-time space, two in which it is optimal to be risk-averse for delay, and two in which it is optimal to be risk-prone. Thus, there are regions in which it is optimal to be risk-prone in time but risk-averse in reserves and vice versa. Variability in amount and time have different effects because although both affect variance in rate of intake, for an animal with a short time horizon variable delays eat unpredictably into the foraging time left in the day, whereas variable amounts do not.

The message from this brief overview of risk-sensitive foraging theory is that, as Houston and McNamara have stressed (Houston, 1991; McNamara *et al.*, 1991; McNamara and Houston, 1992), there is no single model of risk-sensitive foraging and no single prediction. Notably, the energy budget rule is not a universal prediction of risk-sensitive foraging theory. This complexity has a number of ramifications for how risk-sensitive foraging theory can be tested. As pointed out by Houston and McNamara (Houston, 1991; McNamara and Houston, 1992) failures to support the energy budget rule are inconclusive, since it is always possible to claim that the wrong model has been tested. In our view this

damages rather than supports the value of risk-sensitive foraging theory as an explanatory framework since it is often difficult to ensure that the correct model is being tested *a priori*. A related problem is that it is practically impossible to make quantitative predictions from risk-sensitive foraging models. Even if the correct model can be chosen for a given situation, there will still be too many unknowns to predict quantitatively how an animal should respond to a given manipulation. A related problem is that it is difficult to assess the fitness benefits bought by a shift in preference of a given magnitude. For instance, Caraco *et al.*'s (1990) juncos switched from an average of 60% preference for variability under a cold regime to an average of 37% preference for it under thermoneutral conditions, but we have no idea of the magnitude of the benefit accrued by this switch.

A further problem with testing risk-sensitive foraging theory, and in fact most normative models, is that it is necessary to make assumptions about the mechanisms an animal uses to assess important variables such as energy budget or rate of intake. While the many different manipulations that have been tried are all theoretically capable of modifying energy budgets, this does not mean that a shift in budget will necessarily be registered by the animal, since this depends on the proximate mechanism it uses to assess its budget. For example, in the natural habitat ambient temperature may be easy to measure and provide a reliable correlate of energy requirements. However, an animal that has evolved a rule-of-thumb to change its risk preference using ambient temperature may fail to register a change in energy requirements induced by another means such as a period of restricted access to food. At present we do not know how animals assess their energetic status, meaning that a failure to demonstrate an effect of an energy budget shift on preference cannot be interpreted as implying that the energy budget rule would not predict behavior in the wild. A similar argument could be applied to the assessment of variance in rate of gain. This problem is only partially solved by doing experiments in the field (*e.g.*, Cartar, 1991),

since under more natural conditions it is very difficult to control and measure the experience of individual foragers, thus introducing a different set of problems.

From this discussion it becomes clear that it will be exceedingly difficult to test the ecological validity of risk-sensitive foraging theory by rejecting predictions of models, since the rejection of a model can always be attributed to some cause other than the general validity of the theory. Two different tacks can be taken: either to test predictions of risk-sensitive foraging theory that are common to all of the models, or to test single predictions that while not common to all risk-sensitive foraging models are unique to risk-sensitive foraging theory, and therefore if confirmed will suggest that the theory is applicable to behavioral decisions in at least some circumstances. This latter approach suffers from the problem that other theoretical frameworks may subsequently be found that will predict the same pattern of behaviour.

A candidate for a universal prediction of risk-sensitive foraging theory is that all risk-sensitive foraging models depend on the environmental variance being risky or unpredictable rather than just variable. Thus we can predict that if the observed preferences have evolved for the reasons proposed by risk-sensitive foraging theory, an optimal animal should be risk-sensitive only in the face of unpredictable variance and not predictable variance. We have recently tested this idea in starlings with the result that the birds appear not to treat predictable and unpredictable variable delays differently in the manner predicted by risk-sensitive foraging theory (Bateson and Kacelnik, in press). However, our prediction may in practice be invalid, because predictable variance is probably uncommon in the natural world making it possible that animals have evolved a rule-of-thumb of treating all variability as risk whether it is predictable or unpredictable.

The energy budget rule emerges as the best candidate for a prediction that is currently unique to risk-sensitive foraging theory, and we suggest that a sound empirical demonstration of this would be strong evidence for relevance of the theory. Caraco

et al.'s (1990) demonstration of the predicted switch in juncos has been heralded as just such a test (*e.g.*, Houston and McNamara, 1990), however we claim below that the test has some weaknesses and deserves replication and extension.

How much did the juncos know and when did they know it?

Risk sensitivity predictions assume full knowledge of the stable environmental statistics (for a new model that does not see McNamara, 1996). It is our view that attempts ought to be made to demonstrate knowledge independently from preferences, lest the observed choices are due to uncertainty about the parameters of the problem and not to risk sensitivity. Here we ask how many trials a forager has to experience before it can estimate the probabilities of the alternative outcomes of a variable option with a given level of certainty. Consider a foraging option with two alternative outcomes, a small prey item programmed with probability p and a large prey item programmed with probability $q = (1 - p)$. Experimental animals have to estimate these probabilities from experience. Let the animal experience n independent trials of this option, on a proportion p_s of which it gets a small item and q_s a large item. The confidence interval for the proportion, p_s is given by

$$p_s \pm z \sqrt{\frac{p_s q_s}{n}}$$

where z is the z-score corresponding to the desired level of confidence *e.g.* for 95% confidence, $z = 1.96$. If the level of accuracy is defined as $\pm d$, where

$$d = z \sqrt{\frac{p_s q_s}{n}}$$

then by rearranging the above expression it is possible to show that the number of trials, n, necessary to say that $P(p_s - d < p < p_s + d) \geq 0.95$ is

$$n = p_s q_s \left(\frac{z}{d}\right)^2 .$$

For given values of z and d, the value of

TABLE 2. *Percentage confidence in an estimate of p from an observed proportion based on different numbers of trials for three levels of accuracy.*

Number of trials (n)	Accuracy of estimate (d)		
	0.05	0.1	0.2
10	24.8	47.3	79.4
20	34.5	63.0	92.6
40	47.3	79.4	98.9
80	62.9	92.6	100.0
160	79.4	98.9	100.0
320	92.6	100.0	100.0
640	98.9	100.0	100.0

n will be at a maximum when the product $p_s q_s$ is at its maximum, i.e., $0.5 \cdot 0.5 = 0.25$. Assuming this scenario, to achieve 95% confidence that the true probability, p, lies within ± 0.1 of the proportion experienced, p_s, the animal would need to experience, and remember, the outcomes of 96 trials. This applies to most risk-sensitivity experiments because the two outcomes of the variable option are often programmed to occur an equal number of times in a session to ensure the animal gets equal experience of both. Table 2 shows the percentage confidence in estimates of p based on different numbers of trials and at different levels of accuracy.

Caraco *et al.* (1990) tested juncos' preferences by giving each bird 40 trials a day for three days. The first 16 trials each day were forced trials of which half were of the variable option, therefore in the extreme case where an animal chooses the fixed option exclusively in the choice trials it will only have experienced the variable option 24 times in the entire experiment. In fact Caraco *et al.* report that many of the birds exhibited similar behavior on all three days of the test, suggesting that they had reached asymptotic performance after only eight experiences of the variable option. It can be seen from Table 2 that the animals cannot have had a confident estimate of the parameters of the variable option given this amount of experience. McNamara (1996) has shown that uncertainty of the parameter can change the predictions of a risk-sensitive foraging model with the general effect of making risk proneness less likely. Two measures that can be taken to combat this difficulty are first, to provide evidence that

the subjects' preferences are stable, and second, to demonstrate that they know the possible outcomes in the two options on offer. In the case of reward size this could be done by examining whether when given extra seeds the subject pauses after having collected the programmed number, and in the case of delay to reward pecking rates can be used to demonstrate whether the subject knows when food is due (e.g., Bateson and Kacelnik, 1995a).

To summarise our arguments in this section:

● Risk-sensitive foraging theory is a suite of models that make different predictions about how animals should respond to variability in both amount and delay depending on the precise biological scenario assumed. The energy budget rule is not a universal prediction of all risk-sensitive foraging models.

● It will therefore be difficult to test the general ecological validity of risk-sensitive foraging theory by rejecting specific models.

● The evidence for switches in preference with energy budget is the strongest supporting evidence for risk-sensitive foraging theory, however studies need to be replicated since even the most cited have weaknesses on close inspection.

ALTERNATIVE APPROACHES

Time discounting based on the probability of being interrupted

Our purpose in this section is to discuss one effect of variability in delay to food. However, to introduce the argument we start by considering a human subject who is given a one-off choice between gaining a reward immediately and another, of the same magnitude, to be delivered some time later. In this context a one-off choice is one that occurs once only in the life of the subject. Note that at this stage we are not introducing any programmed variance in the delays to reward.

One expects an immediate reward to be preferred over a delayed reward for a variety of reasons. First, for suitable time scales and rewards such as money or offspring in an expanding population, there is

an added value of early gains because the money could be made to work or the offspring could reproduce, so that at the end of the longer delay, a reward obtained earlier has increased by gaining compound interest. Second, there is the probability of loss by interruption. If a promised reward has a chance of being lost during the waiting time, the longer the delay the greater will be the cumulative probability that the reward does not arrive. A delayed reward will have a lower expected value than an immediate one because expected value is the result of multiplying its worth by the probability of its being delivered. This argument implies loss of value of the delayed reward, while the first one (interest) involves gain in value of the early reward. Other logically distinct reasons can be suggested for effects of delay on the value of single rewards. For instance, if reward value depends on the subject's state, and state may change during the waiting time, the value associated with a delayed reward may have greater uncertainty because of unpredictability in the state of the subject itself in the future. Finally, the delayed reward may be less desirable if the waiting time cannot be freely employed for other purposes. This would apply if the choice was one of a series rather than a one-off event. By choosing an immediate reward, time is saved that can be used for making more choices in the future. Similarly, when a reward is lost because of an interruption, this releases time for pursuing further rewards. For present purposes we do not distinguish between these various effects of delay, because although they differ in nature, the effects on the relative value of immediate and delayed rewards may be similar.

The value of a delayed reward V, can be written as $V = F(A, d)$ where A is the worth of the reward if obtained immediately, d is the delay from the point of judgement and F is the discounting function ($F(A, 0) = A$). There is a sizeable theoretical and empirical literature on the shape of the discounting function, F (for a recent discussion see Myerson and Green, 1995). A normative approach (Samuelson, 1937), predicts the shape of F given certain additional assumptions. If the reason underlying discounting

is the probability of loss through interruption, the discounting functions for rewards of different magnitude are equal, the chances of the reward being lost are constant per unit of waiting time, and the subject is fully informed of all the parameters of the problem, then F ought to be a negative exponential function, since V depends on the cumulative probability that a loss might occur between the time of the choice and the time of reward delivery. This approach has dominated foraging theory (*e.g.*, Kagel *et al.*, 1986a; Stephens and Krebs, 1986; McNamara and Houston, 1987).

There have been some attempts to explore how humans actually discount single delayed events. In a recent example, Myerson and Green (1995; see also Green and Myerson, 1996) gave subjects a choice between two notional sums of money, s1 to be delivered immediately and s2 to be delivered after a delay, d. By systematically varying s1 for each pair of values of s2 and d they found the value of s1 at which each subject switched preference from one option to the other. This gives an estimate of the relative value of the immediate and the delayed rewards from which the shape of the time discounting function, F, can be obtained. In common with other similar studies this experiment shows that the shape of the discounting function is hyperbolic rather than exponential, meaning that for a given absolute increase in d the fraction of value lost decreases as d increases (*cf.* an exponential function where the proportional reduction in value with an increase in d is independent of d). Therefore, the empirically determined discounting function is inconsistent with the function predicted under the assumptions outlined above. Specifically, the function that most closely fits available data is of the form

$$V = A/(a + kd), \qquad (1)$$

where a and k are constants with a small and k close to unity.

We now extend the argument to include a one-off choice between a fixed and a variable delay option. In the fixed option a reward of fixed magnitude A is promised for delivery after a fixed delay d_f, while in the variable option a reward of fixed magnitude

A is promised after a delay of either $d_f - \delta$ or $d_f + \delta$, with equal probability. The expected delay in the variable option is $d_v = 0.5[(d_f - \delta) + (d_f + \delta)] = d_f$. Note that because we are dealing with a one-off choice, the subject is told the parameters of the variable option rather than experiencing the variance personally. The discounted value of the fixed option is $V_f = F(A, d_f)$, but the subject may employ different algorithms to assign value to the variable option. For instance, it may use

$$V_{v1} = F(A, d_v) \qquad (2a)$$

or

$$V_{v2} = 0.5[F(A, d_f - \delta) + F(A, d_f + \delta)] \qquad (2b)$$

In Equation 2a the subject discounts reward value using the expectation of the delay, while in Equation 2b it averages the discounted value of each possible outcome. Thus $V_{v1} = V_f$, but for concave-up discounting functions, such as the empirically derived hyperbola or the predicted negative exponential, by Jensen's inequality, $V_{v2} > V_{v1} = V_f$. Thus, to understand how variability in delay affects preference, both the shape of the discounting function and the algorithm used for computing expected value (whether subjects use Equation 2a or 2b) come into play.

In the case of a one-off choice between a fixed and a variable amount option the situation is different. Consider a fixed option with a reward of fixed magnitude A_f promised for delivery after a fixed delay d_f, and a variable option with a reward of either $A_f + \delta$ or $A_f - \delta$, with equal probability, promised after a fixed delay of d. As above the expected magnitude of the variable reward is $A_v = 0.5[(A_f - \delta,d) + (A_f + \delta,d)] = A_f$, and the discounted value of the fixed reward is $V_f = F(A_f, d)$. However, as long as F is a linear function of A, as is the case for most of the putative equations for value, then $F(A_v, d) = 0.5[F(A_f + \delta) + F(A_f - \delta)]$. Therefore variability in amount *per se* should have no effect on preference. If, however, A and d are correlated, as will be the case if large food items take longer to handle than small ones, then variability in

amount can affect preference through its effects on delay. We return to this issue later.

We now turn to whether subjects should use Equation 2a or 2b, regardless of the shape of the discounting function. A subject facing the one-off choice used in this example should use Equation 2b, because it is the expected values of the two possible outcomes that count. However, when a series of repeated choices is involved the problem differs because the time lost during the delay to reward in one choice affects the time at which the next choice can be made. Thus, when rate of reward per unit time is to be maximized, as in most of the problems addressed by classical optimal foraging theory, the use of Equation 2a is to be expected.

In summary, by combining the empirical evidence for a concave-up time discounting function and the normative argument for using the average of discounted payoffs (Equation 2b) we predict that in a one-off choice a reward promised after a variable delay ought to be preferred to a reward promised after a fixed delay equal to the arithmetic mean of the possible delays in the variable option. In spite of the variety of closely related tasks in the human literature on self-control, we have not yet found empirical tests of this prediction.

REPEATED CHOICES AND THE RATE ALGORITHM

Hyperbolic discounting without variability

Many experiments on non-human animals have examined the effect of length of delay and variability in delay on reward value. These experiments have resulted in hyperbolic time discounting functions (Mazur, 1987; Rodriguez and Logue, 1988) and in preference for variable delays (Herrnstein, 1964; Davison, 1972; Gibbon et al., 1988; Reboreda and Kacelnik, 1991; Bateson and Kacelnik, 1995a). Consequently, it is tempting to generalise the ideas from the previous section and treat the animal results as equivalent to the human work on time discounting using one-off choices. There are however problems in translating between human and animal work because one-off choices with variance, as described

above, cannot be offered to animals. Humans can be told the parameters for a one-off choice, but animal subjects must be trained by repeated exposure to the two alternatives. These are generally identified by different stimuli followed by programmed delays and reward amounts, thus choices are among arbitrary conditioned stimuli rather than the rewards themselves. Training gives the subject information about the parameters of the options on offer and the opportunity to learn whether or not trials are ever interrupted during the delay. Given that animals do not treat probabilistic rewards (partial reinforcement) as they treat certain rewards (Mazur, 1985), it is clear that they can adjust to various probabilities of interruptions to foraging. It therefore seems unparsimonious to assume that hyperbolic discounting in animals that have been maintained on deterministic schedules for long periods is a response to the probability of interruption. While it could be argued that animals are hardwired to expect some baseline level of interruptions, we investigate whether there are alternative explanations for hyperbolic discounting in animals faced with a sequence of choices.

In a sequence of choices, immediate rewards are more valuable than delayed ones because the time saved can be used to pursue further rewards. Longer delays imply a loss of foraging opportunity. An appropriate measure of value in such situations is rate of gain, the ratio of expected gain over expected time which is the currency of classical optimal foraging theory. Similarly to Equation 1, Value = Rate = $F(A, d) = A/(t + d)$, where t includes time intervals other than d, such as handling time, inter-trial interval or travel time. Therefore rate drops hyperbolically with d. There is no need to invoke any variability to predict hyperbolic discounting.

Repeated choices and variability in delay

Once we have a sequence of choices, we can consider the effects of introducing variability in d. For a stimulus that is paired with a variable delay, even in the absence of uncertainty, the distinction between Equations 2a and 2b applies: value can be calculated either by discounting the average

reward size by the average of the delays, or by calculating the average of the discounted outcomes. In the context of rate maximization, the algorithm resulting from applying Equation 2a is long-term rate, also named the Ratio of the Expectations of amount and time (RoE), and from Equation 2b, short-term rate or the Expectation of the Ratios of amount and time (EoR, Bateson and Kacelnik, 1995a; Bateson and Kacelnik, 1996). These currencies were initially discussed in the foraging literature in the context of the so-called "fallacy of the averages" (Templeton and Lawlor, 1981; Gilliam et al., 1982; Turelli et al., 1982).

For one-off choices we claim that Equation 2b is the appropriate currency, but with a sequence of choices the problem is different. Here RoE (*i.e.*, Equation 2a) characterises the benefit obtained over a foraging period and therefore it is this measure of rate that should be maximized. According to RoE, two food sources have the same value if one gives a reward after a fixed delay and another gives the same reward after a variable delay of the same average duration. However, experimental results show that animals have a strong preference for food sources with variable delays, and the exact value of this preference is compatible with assuming that variable food sources are characterised by something close to Equation 2b, or EoR (Mazur, 1984; Mazur, 1986; Bateson and Kacelnik, 1996). Also relevant here, the literature shows that bees are risk-averse when there is variability in the volume of nectar present in flowers (Real, 1981; Waddington, 1981), and Harder and Real (1987) have claimed that this finding is compatible with bees using EoR. The correlation between nectar volume, A, and handling time, d, in bees results in A appearing in both the numerator and denominator of F resulting in a concave down function relating the volume of nectar taken to intake rate. This relationship does not lead to risk sensitivity if bees are maximising RoE, but if they are maximising EoR then risk aversion is predicted by Jensen's inequality (see Caraco et al., (1992) for a full analysis of the effects of covariance between A and d on risk sensitivity under maximization of EoR). The magni-

tude of this effect of variability in amount will depend on the proportion of total foraging time that is spent handling the prey item. In bees this is large, however in many of the bird studies the effects of handling time are likely to be swamped by longer inter-trial intervals, and variability in amount is unlikely to lead to detectable risk-aversion even if there is a correlation between amount and handling time.

We do not know precisely the cost of maximising EoR rather than RoE, and it could be argued that perhaps in natural environments it may make little difference which algorithm an animal uses. If there is no variation in the denominator, $t + d$, Equations 2a and 2b become identical. However, measurements of the distribution of inter-prey intervals experienced by starlings foraging on natural pastures suggest that this argument cannot help us to understand why animals appear to maximise EoR, since the birds would have experienced significantly reduced rates of energetic intake if foraging decisions were controlled by this algorithm (Bateson and Whitehead, 1996).

Differential weight of various time components

From the perspective of rate maximization, all the time allocated to a feeding event is unavailable for performing other actions. Time lost pursuing a prey item and time spent handling it cause the same loss of foraging opportunity, and this is why the denominator in the equation for rate is the sum of all the expected times in a foraging cycle, travel, searching, handling etc. This applies whether the subjects use RoE (as expected for rate maximising) or EoR (as required to fit the evidence). Nevertheless, it has been shown (Snyderman, 1987) that times intervening between the arbitrary conditioned stimulus and the reward have much greater impact on the value of an option than those following the reward. A food item associated with a short choice-reward delay and a long post-reward delay is preferred over one in which the lengths of the two delays are reversed such that the choice-reward delay is long and the post-reward delay short. This strong impact of

choice-reward delays is also found in studies using two sources, one of which has variable delays. For example, in an experiment with starlings one stimulus was followed by a fixed delay before food and another by either of two delays (Bateson and Kacelnik, 1996). There were two additional times in the foraging cycle, the reward time (or handling time) and the inter-trial interval. The fixed delay was adjusted by titration until the subjects were indifferent between the two food sources. We calculated the predicted fixed delay at which the two stimuli ought to be equally valuable according to six algorithms resulting from applying either RoE or EoR, each with three combinations of time components (namely including only choice-reward delay, adding the handling time and further adding the inter-trial interval). The point of indifference was very close to that predicted by EoR maximization computed with only the choice-reward delay.

Interruptions could be brought into the picture by claiming that this predominance of the period preceding the reward is due to the possibility of natural interruptions, since in nature interruptions may be more likely during this period than at other times. We do not favour this approach for two reasons. First, there is no evidence for this nor is it likely that evidence will be forthcoming. Relative susceptibility to interruptions during various foraging components is bound to be highly situation-specific and inaccessible to research. Second, the results are obtained after long-term training without interruptions. Since laboratory animals can respond appropriately when interruptions do occur, it is parsimonious to avoid the involvement of this *ad hoc* postulate.

Summarising our arguments about time discounting and rate algorithms:

- Time discounting based on interruptions is only relevant to one-off choices, and it is not supported by the finding that the discounting function appears to be hyperbolic rather than exponential and appears to be insensitive to the lack of interruptions.
- For sequences of choices without variability, rate maximization predicts a hy-

perbolic relationship between value and delay.

- When variability is present, rate maximization favours Equation 2a (RoE); but the empirical results favour Equation 2b (EoR).
- Maximization of EoR explains the hyperbolic time discounting function and risk-proneness when variability is in delay. It can also explain risk aversion when variability is in amount if there is a correlation between amount and handling time and if handling time makes up a significant portion of the foraging cycle.
- All of the rate maximising arguments fail to account for the predominance of the delay between choice and the delivery of reward over other times in the foraging cycle.

Associative learning and delay variability

In this section we offer a mechanistic explanation for risk proneness for delay, for the use of EoR and for the special relevance of the interval between stimulus and reward. We suggest that all of these findings can be explained by considering the processes by which animals learn about causal relationships in the environment (see Montague *et al.* (1995) for a related argument). Our idea is that the general principles of associative learning have evolved under broader selective pressures than those acting on foraging decisions, and that they lead to deviations from optimality in some foraging tasks.

The link with associative learning comes from the fact that the training of subjects in foraging experiments follows a protocol compatible with that used in standard studies of Pavlovian conditioning. In a conditioning experiment the subject is exposed to an originally neutral, conditioned stimulus (CS) and after some time delay it receives a meaningful, unconditioned, stimulus (US) such as a food reward. Rewards that follow shortly after the stimulus onset strengthen the association with the CS more efficiently than those that occur after a longer delay. This makes functional sense since such a mechanism is more likely to allow the animal to respond to causal relationships in

its environment (Dickinson, 1980). The function relating both the rate of learning and its asymptotic strength to the length of the CS-US delay is also non-linear, and appears to be hyperbolic. The hyperbolic relationship can be explained if, for example as has been shown by Gibbon *et al.* (1977), learning is proportional to the ratio of the inter-trial interval to the CS-US delay. This is adaptive, since this ratio indicates how much extra information the CS provides about the timing of an impending reward. Alternatively, a hyperbolic relationship could be explained if the strength of the CS-US association depends on the ratio of the size of a prey item to the CS-US delay. If learning occurs each time a reward is received, the value of the association with the CS will approximate Equation 2b, because to learn according to Equation 2a, the subject would have to remember the set of CS-US delays, compute the average on these intervals and then attribute value to the CS. It is unlikely that these intervals should be stored before the animal knows of their significance. Given a hyperbolic relationship, Jensen's inequality suggests that an animal should learn more efficiently about a CS that is separated from its US by a variable delay than it will about a CS separated from the US by a fixed delay of the same mean duration. Working on the assumption that differences in the efficiency of learning about different CSs may translate into preferences in a choice test between the CSs, this idea could explain animals' preferences for options with variable delays to reward. If the subject is trained with two stimuli, one followed by a fixed CS-US delay and another followed by either of two delays with the same mean as the fixed one, as in our starling experiment mentioned above, the attribution of value by the reward that follows a variable delay will be greater, and the subject will show preference for that CS.

Following this interpretation, the reason why the interval between stimulus and reward is more influential than other periods in the cycle is that learning occurs at the time when food is encountered by retrospective assignment of value to the CS. The stimulus does not gain any associative val-

ue at the end of the handling time or the inter-trial interval. Therefore, this account copes with Snyderman's (1987) finding that the CS-US interval is of special importance. Thus, in summary, the learning model explains the features of animals' responses to variable delays, but it cannot currently explain responses to variability in amount, although some modifications could allow this (see Montague *et al.*, 1995). Also, it cannot be used to explain the results from one-off choice experiments in humans where no learning is involved.

Psychophysics and memory

Our final perspective takes variability into the head of the forager. So far we have ignored the problems of measurement in that none of the ideas discussed above make any reference to the accuracy with which information is processed. Here we consider how errors in perception or memory will affect animals' responses to fixed and stochastic environments and how they could generate sensitivity to variance in either amount or delay.

An attempt to analyse the effects of measurement error on optimal prey choice is provided by Yoccoz *et al.* (1993). They model the effects of perceptual error by assuming explicit relationships between the real energy content (G) and handling time (T) of each prey item and what the forager estimates these to be (X and Y respectively). The random variables X and Y are constructed from the random variables G and T by adding normally distributed errors with a mean of zero such that $X = G + E_G$ and $Y = T + E_T$. Introducing this type of perceptual error into a sequential encounter diet choice model generates several predictions concerning the behavioural tactics foragers should employ. Of interest in the present context, foragers are predicted to be sensitive to the variance in the expected energy content of food items, preferring either environments with highly predictable energy content of food items or fairly unpredictable energy content of food items even when the resources are characterised by the same mean values. We shall not expand on these predictions because we question the validity of the assumptions used to derive

them. The model assumes that perceptual error is independent of the magnitude of the quantity being perceived, and as we shall show below this is contrary to what is known to be the case.

Weber's law

Put simply, Weber's law says that the difference in stimulus magnitude required to see two stimuli as different is proportional to the mean value of the stimuli. This minimum difference is called the Just Noticeable Difference or JND. The law is compatible with modelling subjective estimates as increasing in proportion to the logarithm of the stimulus magnitude. We shall not discuss details of deviations of this law and the problems of any attempt to model sensation. It will suffice here to accept that the law is successful in a variety of domains and to go on to examine some of its consequences.

Let us start by assuming that an animal forms a representation of the magnitude of some stimulus, such as the size or the inter-capture interval for a prey type. Let us further assume for the moment that there is no variability in the objective experience, namely the prey has a fixed size and it takes always the same time to find. On each encounter with that prey type, in accordance to Weber's law the subject will register a value within the range of its true value plus or minus one JND. The distribution of values experienced will be close to a normal distribution centred on the true value of the stimulus, and with width proportional to it (because of the proportionality of the JND). This notion introduces a form of variability that is generated by the individual itself, but that leads to interesting interactions with environmental variance, as we shall see below.

Empirical evidence consistent with this account comes from tasks in which the subject reproduces the value of a foraging parameter from its memory, or is required to make a decision based on its memory for a foraging parameter. Because of their relevance for risk-sensitivity we describe two examples, one related to delays to food and the other to amounts of food. In the first (Kacelnik *et al.*, 1990; Brunner *et al.*,

1992), starlings were given a task simulating foraging in patches with fixed inter-prey interval and sudden depletion. Within each patch, an unpredictable number of prey could be obtained until the patch was exhausted, and the interval between successive prey was constant. The birds could only detect patch exhaustion by the interval since the last food delivery: if the typical inter-prey interval was exceeded, then the patch would yield no more prey and the only option was to give up and travel to a new patch. The birds learned this task and showed giving up times that were proportional to the inter-prey intervals, and the standard deviations of their giving up times were proportional to the mean intervals, as required if Weber's law applies to time delays.

In a second experiment (Bateson and Kacelnik, 1995b), starlings faced a choice between two sources of reward. One was fixed (the standard) and the other gave an amount of food that increased when the standard was preferred and diminished when the opposite happened. The oscillation of the amount in the adjustable alternative gives a measure of the accuracy of discrimination. Further, because the discrimination is not between two amounts of food placed in front of the animal, but between two stimuli (coloured pecking keys) associated with these amounts, the task depended on memory. By changing the value of the standard, it was possible to examine the accuracy of this discrimination. As expected, the spread of the oscillations of the adjustable alternative increased when the value of the standard was increased as required if Weber's law also applies to amount of food.

Scalar expectancy theory

The most immediate implication of Weber's law to the problem of variability in foraging is to give a new explanation for why food sources with variability in delay are preferred to alternatives with equal mean delay and no variance, and conversely why food sources offering fixed amount of food are generally preferred to alternatives offering variable amounts with the same mean (Reboreda and Kacelnik, 1991; Bate-

son and Kacelnik, 1995a; see also Perez and Waddington, 1996). Reboreda and Kacelnik have developed a model based on the observation that Weber's law affects memory for both amount and delay. In their model, which is based on Scalar Expectancy Theory (Gibbon, 1977; Gibbon et al., 1984; Gibbon et al., 1988), the memory formed for time intervals or amounts of food looks like a normal distribution with mean equal to the real value of the stimulus and standard deviation proportional to the mean. The constant of proportionality relating the mean to the standard deviation (i.e., the coefficient of variation) is assumed to be fixed for a given subject in a given experiment.

In the analyses of variability discussed above, variable food sources offer equal frequency of two outcomes. According to Reboreda and Kacelnik's version of SET, a coloured cup signalling a food source offering 1 or 7 seeds with equal probability would be remembered as the joint memories of 1 and 7 seeds. Similarly, a stimulus followed by food after either 5 or 15 s will be remembered as the joint memories of these two intervals. However, because of Weber's Law, a uniformly distributed mixture of intervals or amounts (as in these experiments) is represented in memory as a positively skewed distribution. This skewed distribution is generated by combining the normal distributions that represent each of the constituent elements of the mixture (Fig. 4).

Subjects are assumed to choose between two options by retrieving a sample from memory for each option, comparing these samples, and preferring the option offering the better sample (bigger reward size or shorter delay). Because the memory representations of variable options are skewed to the right, whereas memory representations of fixed options with the same mean are symmetric around the true mean, in more than half of comparisons the sample from the representation of the variable option will be smaller than the sample from the representation of the fixed option. For options with equal long-term rates of food intake the model predicts that an option offering variability in delay to food should be
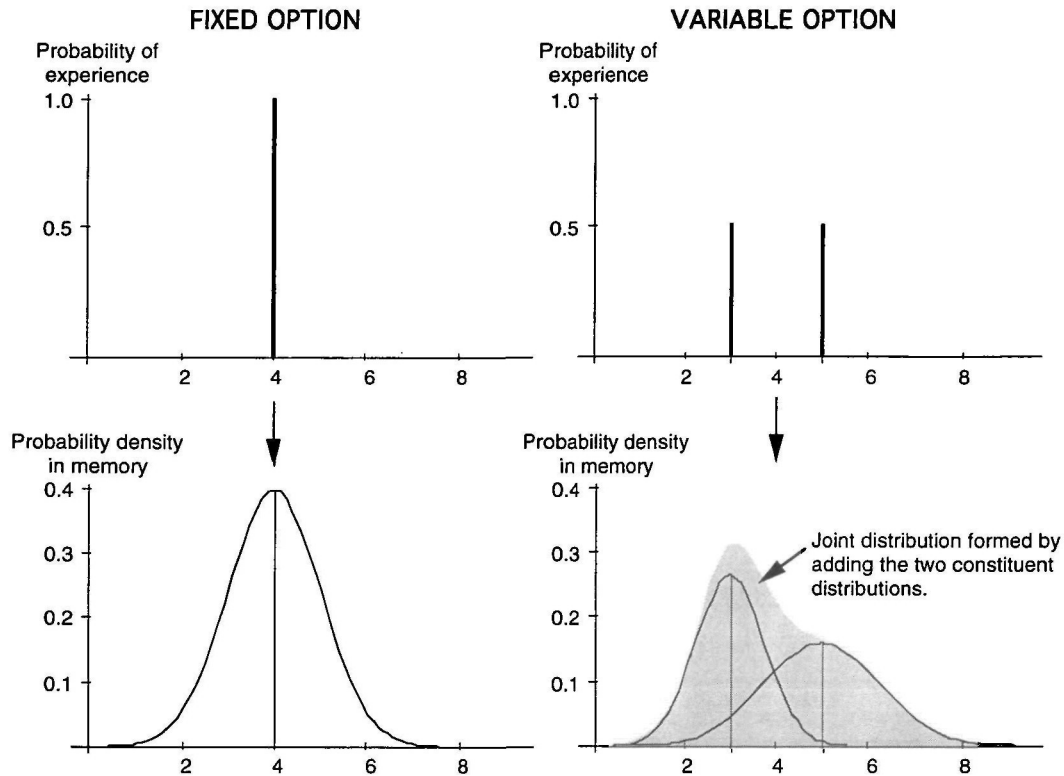
FIG. 4. The upper two panels represent the experienced distribution of outcomes (amounts or delays) in a fixed (left) and variable (right) option. The lower tow panels represent the distributions that are assumed to be formed in memory as a result of the above experiences. Note the skew in the distribution of the memory for the variable option that results from the constant accuracy with which the constituent stimuli are represented (this Fig. is reproduced from Bateson and Kacelnik, 1995a).

chosen more often than a fixed alternative and that a source offering variability in amount of food should be chosen less often than a fixed alternative.

Although the assumptions of this account are supported by the data showing that Weber's law applies to memory for both amount and delays and its predictions for preference for variability in delay (see also Gibbon et al., 1988) and aversion to variability in amount have been confirmed, the model fails in some respects, which we list below. (i) Although the model predicts that animals should choose a fixed amount of food more often than a variable amount in our studies the risk-aversion for amount was not found in the frequency of choices but on other measures of preference such as the latency to choose and the number of pecks made. It is likely that these results

depend on the particular value of amount variance used in the tests and thus further experiments with larger variance are required (Bateson and Kacelnik, 1995a). (ii) Regarding delays, the qualitative agreement with the model was much stronger, but there is a quantitative discrepancy. The model predicts that animals should be indifferent between a fixed and variable option when the fixed option is equal to the geometric mean of the two outcomes in the variable option (Bateson and Kacelnik, 1995a), whereas subsequent research has demonstrated that indifference occurs at approximately the harmonic mean (equivalent to Equation 2b or EoR) (Mazur, 1984; Mazur, 1986; Bateson and Kacelnik, 1996). (iii) The model cannot accommodate the observed effects of energy budget on preference. As discussed above, these results

are not as well documented as static preference for delay variability, but they do exist. While the model is unable to produce a switch in preference from risk aversion to risk proneness or vice versa, it could accommodate shifts in the extent of preference (such as those reported by Reboreda and Kacelnik, 1991) if it is assumed that energy budget could affect the degree of perceptual error. (iv) In its present form the model does not predict what animals should choose if the alternatives differ in the variance of both amount and delay simultaneously, since it provides no way of combining predictions relating amount and delay. A more recent version of SET (Brunner *et al.*, 1994) addresses some of the above problems, including that of the indifference point and the integration of amount and time, however it also brings with it new limitations that we do not have space to discuss here.

DISCUSSION AND CONCLUSIONS

Our aims in this paper have been twofold: first to review the current state of our knowledge of how animals respond to environmental variability and second to describe and compare a number of competing explanations for this behavior, that although clearly related have seldom been brought together. One outcome of the literature review is the identification of methodological problems that need to be considered when interpreting or planning experiments on risk-sensitivity. We found that both the presence and the direction of risk-sensitive preferences can be influenced by what could appear to be trivial differences in the choice procedure, the schedule of reinforcement, the means used to deliver rewards or the index used to measure the subjects' preferences. However, in spite of these difficulties we did find general patterns in risk-sensitive preferences that candidate theoretical explanations need to address. In brief, when variability is in the amount of reward:

- Animals are usually risk-averse, sometimes indifferent to risk and rarely risk-prone.
- There is some evidence that energy budget affects the direction of preference.

And, when variability is in the delay to reward:

- Animals are always risk-prone.
- There is currently little evidence to suggest that energy budget affects the degree of risk proneness.
- The time between a subject focusing on a stimulus or making a choice and receiving an associated reward has a larger impact on the value of a foraging option than other times in the foraging cycle.
- Animals are indifferent between a fixed and a variable delay when the expected ratio of amount over the stimulus-reward interval (EoR) is equal in the two options.

Our review highlights substantial gaps in the existing literature. There are few studies that compare variability of both amount and delay in the same species and under the same budget conditions. Of particular interest, studies investigating the effects of variability in amount have used smaller species than those investigating variability in delay. Therefore there has not yet been an adequate test of whether energy budget can influence the direction of risk-sensitive preferences when variability is in delay to reward. To date no studies have investigated the dynamics of risk-sensitive preferences. This could be interesting since risk-sensitive foraging theory and our mechanistic model based on associative learning make opposite predictions here. Risk-sensitive foraging theory predicts that after a run of good luck on a variable option animals will have improved their state more than after a run of bad luck, will be more likely to be on a positive budget, and thus more likely to be risk-averse. In contrast, the learning model predicts that animals should have a stronger preference for options that have recently rewarded them well.

The theories that we have considered to explain the above phenomena rely on two main types of explanation: functional (or normative) arguments that consider the circumstances under which risk-sensitive behavior is adaptive, and causal arguments that consider what is known of the general mechanisms used by animals to perceive, learn and remember information about the

environment. Risk-sensitive foraging theory and the interruption-based rationale for time discounting are strictly functional, whereas the other models combine causal and functional reasoning. In this latter category, empirically observed risk sensitivity is treated as a side effect of behavioural mechanisms that are not direct or exclusive adaptations to cope with risk. Favouring one of these models does not imply taking an anti-optimality stance, because the mechanisms responsible for the observed sensitivity to variance are interpreted as products of evolutionary trade-offs due to selection at some broader level. For instance, our associative learning account produces risk proneness for variability in delay by assuming that the non-linear drop in the efficiency with which the association between an arbitrary stimulus and a meaningful event is learnt, has evolved because of its suitability for detecting causal relations in the environment. This non-linearity is present whether or not there is variability, but in the presence of variability stronger associations are formed as a by-product. In a different example, our discussion of rate algorithms assumes that the selective force has been for rate maximization but that because animals use a particular algorithm (EoR) to compute their average rate of intake they show sensitivity to variability in delay. Similarly, the perceptual errors postulated in the models of Yoccoz et al. (1993) and of Reboreda and Kacelnik (1991) can also be seen as the result of an evolutionary trade-off. It is likely that the errors could be reduced if more time or neural processing capacity were devoted to the representation of the relevant parameters, but that it is not adaptive to allocate resources in such a way. These mechanistic hypotheses are not *ad hoc* additions solely for the purposes of explaining observed results. The hypotheses are often based on known features of animal cognition, and will often lead to the generation of new predictions that can be tested empirically justifying this approach to optimality as a progressive program of research (Kacelnik and Cuthill, 1987; Mitchell and Valone, 1990).

Our stress on causal accounts of risk sensitivity over the purely functional models may be seen as neither mainstream behavioural ecology nor particularly desirable, especially if it led to relinquishing the search for sophisticated and detailed behavioural adaptations. In this case we advocate the approach because when judged overall, a consideration of the role of mechanism has contributed more to behavioural research on risk than unfettered adaptationist thinking. Part of the reason for this is that risk-sensitive foraging theory is not a single model but a variegated set of models pertinent to different biological scenarios. These models make different predictions about how an animal should respond to risk depending on a number of variables (McNamara and Houston, 1992). The existence of these alternative models makes risk-sensitive foraging theory both sophisticated and powerful, but also very hard to test since failures can often be attributed to a mismatch between the experimental situation and the scenario in which the animal's behavior has evolved. It is virtually impossible to have accurate knowledge of the ecological parameters required to select the appropriate model to test, let alone make quantitative or even qualitative predictions for a given species. This has the consequence that negative results cannot be used to reject risk-sensitive foraging theory. A potential response to this problem is to find and test predictions that are general to all risk-sensitive foraging models. One such prediction is that risk-sensitive preferences are due to uncertainty (true risk) and not just to variance which can in some circumstances be predictable.

A second reason why risk-sensitive models may have been hard to test lies in the fact that adaptive behavior must be implemented by behavioural mechanisms that will themselves have evolved in specific environments. Unless the conditions of a risk-sensitive foraging experiment are close to those under which an adaptive response to risk evolved the mechanisms may fail to produce an adaptive outcome. This problem could be relevant for the assessment of variables such as energy budget, variance in rate of gain or probability of interruption. These considerations are bound to be significant given that the vast majority of risk-

sensitive foraging experiments are performed in the artificial conditions of laboratories.

Mechanistic models have the obvious advantage that they can be falsified under laboratory conditions. Our claims about the rate-maximising algorithm under constraint (the use of EoR) would be abandoned if our birds maximised the arithmetic rather than the harmonic mean in a variable delay experiment. No speculation about the irrelevance of the lab can save a failed mechanistic model. However, causal explanations can be seen as unsatisfactory on their own since they shift the explanatory burden elsewhere: showing that risk proneness for delay may be a by-product of the way animals learn begs the question of why animals learn this way.

To date none of the available theories adequately accounts for all the crucial behavioural observations in studies of the effects of variability. Risk-sensitive foraging theory does not provide any insight into why risk-aversion should be more common when variability is in amount and risk proneness when it is in delay. The fact that shifts in preference have sometimes been observed when budgets are manipulated remains the single empirical observation that prevents us from abandoning the theory entirely. It is therefore crucial that these key experiments are replicated and extended. The interruption model, EoR and the associative learning model can all account for risk proneness when variability is in delay, but they are unable to account for the general observation of risk aversion when variability is in amount. The associative learning model explains why the time delay between focusing on a stimulus and receiving reward should be more important than other times in the foraging cycle. Although EoR maximising can explain risk aversion to variability in amount if there is a positive correlation between amount and handling time and if handling time is a large proportion of the foraging cycle. The SET model predicts risk aversion when variability is in amount and risk proneness when variability is in delay, but fails to predict indifference at the harmonic mean, and the shift in preference with energy budget. Of course it is

possible that the phenomena described in this review, although superficially related, may be the products of very different underlying mechanisms, or have evolved under different selective pressures, in which case they may never yield to one unifying explanation.

We close by asking the question of whether risk-sensitive foraging theory is still a useful theoretical tool to guide behavioural research. Our tentative reply is that the jury is still out. We can conclude that in this area of research, in common with most others, the achievements of purely evolutionary or purely mechanistic theorising are severely limited. A persistence with Tinbergen's (1963) call for simultaneous attention to both the functions and mechanisms of behavior is not optional but imperative; we can only make progress by following his lead.

REFERENCES

Ahearn, W. and P. N. Hineline. 1992. Relative preferences for various bivalued ratio schedules. Anim. Learn. Behav. 20:407–415.

Banschbach, V. S. and K. D. Waddington. 1994. Risk-sensitivity in honey-bees—no consensus among individuals and no effect of colony honey stores. Anim. Behav. 47:933–941.

Barkan, C. P. L. 1990. A field test of risk-sensitive foraging in black-capped chickadees (Parus atricapillus). Ecology 71:391–400.

Barnard, C. J. and C. A. J. Brown. 1985. Risk sensitive foraging in common shrews (Sorex araneus L.). Behav. Ecol. Sociobiol. 16:161–164.

Barnard, C. J., C. A. J. Brown, A. Houston, and J. M. McNamara. 1985. Risk-sensitive foraging in common shrews: An interruption model and the effects of mean and variance in reward rate. Behav. Ecol. Sociobiol. 18:139–146.

Bateson, M. and A. Kacelnik. 1995a. Preferences for fixed and variable food sources: Variability in amount and delay. J. Exp. Anal. Behav. 63:313–329.

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

Bateson, M. and A. Kacelnik. Starlings' preferences for predictable and unpredictable delays to food: Risk-sensitivity or time discounting? Anim. Behav. (In press).

Bateson, M. and A. Kacelnik. 1996. Rate currencies and the foraging starling: The fallacy of the averages revisited. Behav. Ecol. (In press).

Bateson, M. and S. C. Whitehead. 1996. The energetic costs of alternative rate currencies in the foraging starling. Ecology 77:1303–1307.

Battalio, R. C., J. H. Kagel, and D. N. McDonald. 1985. Animals choices over uncertain outcomes: Some initial experimental results. Amer. Econ. Rev. 75:596–613.

Bednekoff, P. A. 1996. Risk-sensitive foraging, fitness, and life histories: Where does reproduction fit into the big picture? Amer. Zool. 36:471–483.

Behar, I. 1961. Learned avoidance of non reward. Psych. Rep. 9:43–52.

Brunner, D., J. Gibbon, and S. Fairhurst. 1994. Choice between fixed and variable delays with different reward amounts. J. Exp. Psych.: Anim. Behav. Processes 20:331–346.

Brunner, D., A. Kacelnik, and J. Gibbon. 1992. Optimal foraging and timing processes in the starling, Sturnus vulgaris: Effect of inter-capture interval. Anim. Behav. 44:597–613.

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

Caraco, T. 1982. Aspects of risk-aversion in foraging white crowned sparrows. Anim. Behav. 30:719–727.

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

Caraco, T. and S. L. Lima. 1985. Foraging juncos: Interaction of reward mean and variability. Anim. Behav. 33:216–224.

Caraco, T., W. U. Blanckenhorn, G. M. Gregory, J. A. Newman, G. M. Recer, and S. M. Zwicker. 1990. Risk-sensitivity: Ambient temperature affects foraging choice. Anim. Behav. 39:338–345.

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

Caraco, T., A. Kacelnik, N. Mesnick, and M. Smulewitz. 1992. Short-term rate maximization when rewards and delays covary. Anim. Behav. 44:441–447.

Cartar, R. V. 1991. A test of risk sensitive foraging in bumble bees. Ecology 72:888–895.

Cartar, R. V. and L. M. Dill, 1990. Why are bumble bees risk sensitive foragers? Behav. Ecol. Sociobiol. 26:121–127.

Case, D. A., P. Nichols, and E. Fantino. 1995. Pigeon's preference for variable-interval water reinforcement under widely varied water budgets. J. Exp. Anal. Behav. 64:299–311.

Charnov, E. L. 1976a. Optimal foraging: Attack strategy of a mantid. Am. Nat. 110:141–151.

Charnov, E. L. 1976b. Optimal foraging: The marginal value theorem. Theor. Popul. Biol. 9:129–136.

Cicerone, R. A. 1976. Preference for mixed versus constant delay of reinforcement. J. Exp. Anal. Behav. 25:257–261.

Clements, K. C. 1990. Risk-aversion in the foraging blue jay, Cyanocitta cristata. Anim. Behav. 40:182–195.

Croy, M. I. and R. N. Hughes. 1991. Effects of food supply, hunger, danger and competition on choice of foraging location by the fifteen-spined stickleback, Spinachia spinachia L. Anim. Behav. 42:131–139.

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

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

Dickinson, A. 1980. Contemporary Animal Learning Theory. Cambridge University Press, Cambridge.

Dunning, J. B. 1993. CRC handbook of avian body masses. Boca Raton, Florida.

Epstein, R. 1981. Amount consumed as a function of magazine-cycle duration. Behav. Anal. Lett. 1:63–66.

Essock, S. M. and E. P. Rees. 1974. Preference for and effects of variable as opposed to fixed reinforcer duration. J. Exp. Anal. Behav. 21:89–97.

Fantino, E. 1967. Preference for mixed versus fixed-ratio schedules. J. Exp. Anal. Behav. 10:35–43.

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

Gibbon, J., M. D. Baldock, C. Locurto, L. Gold, and H. S. Terrace. 1977. Trial and intertrial durations and autoshaping. J. Exp. Psych.: Anim. Behav. Processes 3:264–284.

Gibbon, J., R. M. Church, and W. H. Meck. 1984. Scalar timing in memory. In J. Gibbon and L. Allan (eds.), Timing and Time Perception, pp. 52–77. New York Academy of Sciences, New York.

Gibbon, J., R. M. Church, S. Fairhurst, and A. Kacelnik. 1988. Scalar Expectancy Theory and choice between delayed rewards. Psychol. Rev. 95:102–114.

Gilliam, J. F., R. F. Green, and N. E. Pearson. 1982. The fallacy of the traffic policeman: A response to Templeton and Lawlor. Amer. Nat. 119:875–878.

Green, L. and J. Myerson. 1996. Exponential versus hyperbolic discounting of delayed outcomes: Risk and waiting time. Amer. Zool. 36:496–505.

Ha, J. C. 1991. Risk-sensitive foraging: The role of ambient temperature and foraging time. Anim. Behav. 41:528–529.

Ha, J. C., P. N. Lehner, and S. D. Farley. 1990. Risk-prone foraging behaviour in captive grey jays Perisoreus canadensis. Anim. Behav. 39:91–96.

Hamm, S. L. and S. J. Shettleworth. 1987. Risk aversion in pigeons. J. Exp. Psych.: Anim. Behav. Processes 13:376–383.

Harder, L. and L. A. Real. 1987. Why are bumble bees risk-averse? Ecology 68:1104–1108.

Hastjarjo, T., A. Silberg, and S. R. Hursh. 1990.

Risky choice as a function of amount and variance in food supply. J. Exp. Anal. Behav. 53:155–161.

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

Hodges, C. M. and Wolf, L. L. 1981. Optimal foraging in bumblebees: Why is nectar left behind in flowers? Behav. Ecol. Sociobiol. 9:41–44.

Houston, A. and J. McNamara. 1985. The choice of two prey types that minimises the probability of starvation. Behav. Ecol. Sociobiol. 17:135–141.

Houston, A. I. 1991. Risk-sensitive foraging theory and operant psychology. J. Exp. Anal. of Behav. 56:585–589.

Houston, A. I. and J. M. McNamara. 1982. A sequential approach to risk-taking. Anim. Behav. 30:1260–1261.

Houston, A. I. and J. M. McNamara. 1990. Risk-sensitive foraging and temperature. Trends Ecol. Evol. 5:131–132.

Kacelnik, A. and I. C. Cuthill. 1987. Starlings and optimal foraging theory: Modelling in a fractal world. In A. C. Kamil, J. R. Krebs, and H. R. Pulliam (eds.), Foraging theory. Plenum Press, New York.

Kacelnik, A., D. Brunner, and J. Gibbon. 1990. Timing mechanisms in optimal foraging: Some applications of scalar expectancy theory. In R. N. Hughes (ed.), NATO ASI Series, Series G: Ecological Sciences, Vol. G 20, pp. 63–81. Springer-Verlag, Berlin.

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

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

Kahneman, D. and A. Tversky. 1984. Choices, values and frames. Am. Psych. 39:341–350.

Krebs, J. R., D. W. Stephens, and W. J. Sutherland. 1983. Perspectives in optimal foraging. In A. H. Brush and G. A. Clark (eds.), Perspectives in ornithology, pp. 165–216. Cambridge University Press, New York.

Lattal, K. A. 1991. Scheduling positive reinforcers. In I. H. Iversen and K. A. Lattal (eds.) Experimental analysis of behavior, part 1, pp. 87–134. Elsevier, Amsterdam.

Lawes, M. J. and M. R. Perrin. 1995. Risk-sensitive foraging behaviour of the round-eared elephant shrew (Macroscelides proboscideus). Behav. Ecol. Sociobiol. 37:31–37.

Leventhal, A. M., R. F. Morell, E. J. Morgan, and C. C. Perkins. 1959. The relation between mean reward and mean reinforcement. J. Exp. Psychol. 59:284–287.

Logan, F. A. 1965. Decision making by rats: Uncertain outcome choices. J. Comp. Physiol. Psychol. 59:246–251.

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

Mazur, J. E. 1985. Probability and delay of reinforce-

ment as factors in discrete-trial choice. J. Exp. Anal. Behav. 43:341–351.

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

Mazur, J. E. 1987. An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin and H. Rachlin (eds.), Quantitative Analyses of Behaviour: The effect of delay and of intervening events on reinforcement value, pp. 55–73. Erlbaum, Hilsdale, New Jersey.

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

McNamara, J. 1983. Optimal control of the diffusion coefficient of a simple diffusion process. Math. Oper. Res. 8:373–380.

McNamara, J. M. 1984. Control of a diffusion by switching between two drift-diffusion coefficient pairs. SIAM J. Cont. 22:87–94.

McNamara, J. M. 1996. Risk-prone behaviour under rules which have evolved in a changing environment. Amer. Zool. 36:484–495.

McNamara, J. M. and A. I. Houston. 1987. A general framework for understanding the effects of variability and interruptions on foraging behaviour. Acta Biotheoretica 36:3–22.

McNamara, J. M. and A. I. Houston. 1990. Starvation and predation in a patchy environment. In B. Shorrocks and I. R. Swingland (eds.), Living in a patchy environment, pp. 23–43. Oxford University Press, Oxford.

McNamara, J. M. and A. I. Houston. 1992. Risk-sensitive foraging: A review of the theory. Bull. Math. Biol. 54:355–378.

McNamara, J. M., S. Merad, and A. I. Houston, 1991. A model of risk-sensitive foraging for a reproducing animal. Anim. Behav. 41:787–792.

Menlove, R. L., H. M. Inden, and E. G. Madden. 1979. Preference for fixed over variable access to food. Anim. Learn. Behav. 7:499–503.

Mitchell, W. A. and T. J. Valone. 1990. The optimization research program: Studying adaptations and their function. Quart. Rev. Biol. 65:43–52.

Montague, P. R., P. Dayan, C. Person, and T. J. Sejnowski. 1995. Bee foraging in uncertain environments using predictive hebbian learning. Nature 377:725–728.

Moore, F. R. and P. A. Simm. 1986. Risk-sensitive foraging by a migratory bird (Dendroica coronata). Experientia 42:1054–1056.

Morris, C. J. 1986. The effects of occasional short (FR1) reinforcement ratios on choice behaviour. Psych. Rec. 36:63–68.

Myerson, J. and L. Green. 1995. Discounting of delayed rewards: Models of individual choice. J. Exp. Anal. Behav. 64:263–276.

Perez, S. M. and Waddington, K. D. 1996. Carpenter bee (Xylocopa micans) risk-indifference and a review of nectarivore risk-sensitivity studies. Amer. Zool. 36:435–446.

Possingham, H. P., A. I. Houston, and J. M. McNamara. 1990. Risk-averse foraging in bees: A comment on the model of Harder and Real. Ecology 71:1622–1624.

Pubols, B. H. 1962. Constant versus variable delay to reinforcement. J. Comp. Physiol. Psychol. 55:52–56.

Real, L. A. 1981. Uncertainty and pollinator-plant interactions: The foraging behavior of bees and wasps on artificial flowers. Ecology 62:20–26.

Real, L. A., J. Ott, and E. Silverfine. 1982. On the trade-off between mean and variance in foraging: An experimental analysis with bumblebees. Ecology 63:1617–1623.

Reboreda, J. C. and A. Kacelnik. 1991. Risk sensitivity in starlings: Variability in food amount and food delay. Behav. Ecol. 2:301–308.

Rider, D. P. 1983a. Preference for mixed versus constant delays of reinforcement: Effect of probability of the short, mixed delay. J. Exp. Anal. Behav. 39:257–266.

Rider, D. P. 1983b. Choice for aperiodic versus periodic ratio schedules: A comparison of concurrent and concurrent-chains procedures. J. Exp. Anal. Behav. 40:225–237.

Rodriguez, M. L. and A. W. Logue. 1988. Adjusting delay to reinforcement: Comparing choice in pigeons and humans. J. Exp. Psych.: Anim. Behav. Processes 14:105–117.

Samuelson, P. A. 1937. A note on the measurement of utility. Rev. Econ. Stud. 4:155–161.

Smallwood, P. D. 1996. An introduction to risk-sensitivity: Using Jensen's inequality to clarify evolutionary arguments of adaptation and constraint. Amer. Zool. 36:392–401.

Snyderman, M. 1987. Prey selection and self-control. *In* M. L. Commons, J. E. Mazur, J. A. Nevin and H. Rachlin (eds.), *Quantitative Analyses of Behaviour*, pp. 283–308. Lawrence Erlbaum Associates, Hillsdale, New Jersey.

Staddon, J. E. R. and N. K. Innis. 1966. Preference for fixed *vs.* variable amounts of reward. Psychon. Sci. 4:193–194.

Stephens, D. W. 1981. The logic of risk-sensitive foraging preferences. Anim. Behav. 29:628–629.

Stephens, D. W. and J. R. Krebs. 1986. *Foraging theory.* Princeton University Press, Princeton.

Stephens, D. W. and R. C. Ydenberg. 1982. Risk aversion in the great tit. *In* D. W. Stephens *Stochasticity in foraging theory: Risk and information,* Unpublished DPhil Thesis, Oxford University.

Templeton, A. R. and L. R. Lawlor. 1981. The fallacy of the averages in ecological optimisation theory. Am. Nat. 117:390–393.

Tinbergen, N. 1963. On aims and methods of ethology. Z. Tierpsychol. 20:410–433.

Turelli, M., J. H. Gillespie, and T. W. Shoener. 1982. The fallacy of the fallacy of the averages in ecological optimization theory. Amer. Nat. 119:879–884.

Tuttle, E. M., L. Wulfson, and T. Caraco. 1990. Risk-aversion, relative abundance of resources and foraging preferences. Behav. Ecol. Sociobiol. 26:165–171.

Waddington, K. D. 1995. Bumblebees do not respond to variance in nectar concentration. Ethology 101:33–38.

Waddington, K. D., T. Allen, and B. Heinrich. 1981. Floral preferences of bumblebees (*Bombus edwardsii*) in relation to intermittent versus continuous rewards. Anim. Behav. 29:779–784.

Wagenaar, W. A. 1988. *Paradoxes of gambling behaviour.* Erlbaum, Hove.

Wunderle, J. M. and Z. Cotto-Navarro. 1988. Constant *vs.* variable risk-aversion in foraging bananaquits. Ecology 69:1434–1438.

Wunderle, J. M. and T. G. O'Brien. 1985. Risk-aversion in hand reared bananaquits. Behav. Ecol. Sociobiol. 17:371–380.

Wunderle, J. M., M. Santa-Castro, and N. Fletcher. 1987. Risk-averse foraging by bananaquits on negative energy budgets. Behav. Ecol. Sociobiol. 21:249–255.

Yoccoz, N. G., S. Engen, and N. C. Stenseth. 1993. Optimal foraging: The importance of environmental stochasticity and accuracy in parameter estimation. Am. Nat. 141:139–157.

Young, J. S. 1981. Discrete-trial choice in pigeons: Effects of reinforcer magnitude. J. Exp. Anal. Behav. 35:23–29.

Young, R. J., H. Clayton, and C. J. Barnard. 1990. Risk sensitive foraging in bitterlings, *Rhodeus sericus:* Effects of food requirement and breeding site quality. Anim. Behav. 40:288–297.

Zabludoff, S. D., J. Wecker, and T. Caraco. 1988. Foraging choice in laboratory rats: Constant *vs.* variable delay. Behav. Processes 16:95–110.

APPENDIX

Definitions of different types of schedules

Interval schedule (I): As soon as the programmed time interval has elapsed the first response made is reinforced.

Time schedule (T): reinforcement occurs after a programmed time interval, independent of the behavior of the subject.

Ratio schedule (R): reinforcement occurs as soon as the subject has completed the programmed number of responses.

All of these schedule types can be either fixed (F), variable (V) or mixed (M), where a fixed schedule is always the same programmed value, a variable schedule is any of a range of values usually described by their mean, and a mixed schedule is one of two values given probabilistically. For example, FI30 specifies that the first response after 30 s has elapsed is reinforced, VR20 that reinforcement is possible after an unknown number of responses with a mean of 20 have been completed, and MT5/10 that reinforcement is equally likely to occur after either 5 s or 10 s. For more details, a reference work such as Lattal (1991) should be consulted.