reprints from

elsevier trends journals

Reprinted from

Trends in Cognitive Sciences

Elsevier Trends Journals
68 Hills Road
Cambridge
UK CB2 1LA
Tel +44 1223 315961
Fax +44 1223 321410

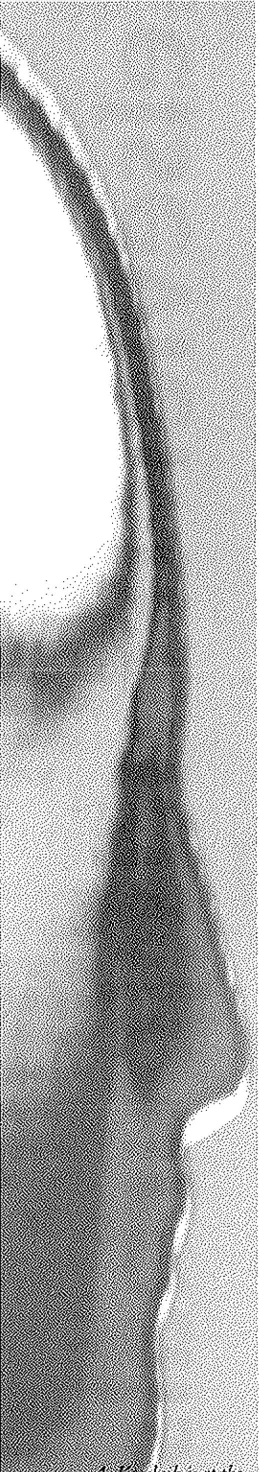

Risk-sensitivity: crossroads for theories of decision-making

Alex Kacelnik and Melissa Bateson

Most actions result in one of a set of possible outcomes. To understand how this uncertainty, or risk, affects animals' decision-making some researchers take a normative approach, asking how an animal should respond to risk if it is maximizing its fitness. Others focus on predicting responses to risk by generalizing from regularities in behavioural data, without reference to cognitive processes. Yet others infer cognitive processes from observed behaviour and ask what actions are predicted when these processes interact with risk. The normative approach (Risk-sensitivity Theory; RST) is unique in predicting a shift in a subject's response to risk as a function of its resource budget, but the predictions of this theory are not yet widely confirmed. In fact, evidence suggests a strong bias towards risk-proneness when delay to reward is risky and risk-aversion when amount of reward is risky, a pattern not readily explained by RST. Extensions of learning theory and of Scalar Expectancy Theory provide process-based explanations for these findings but do not handle preference shifts or provide evolutionary justification for the processes assumed. In this review we defend the view that risk-sensitivity must be studied with theoretical plurality.

A. Kacelnik is at the Department of Zoology, University of Oxford, South Parks Road, Oxford, UK OX1 3PS and M. Bateson is at the Department of Psychology: Experimental, Duke University, Durham, NC 27708, USA.

tel: +44 1865 271164
fax: +44 1865 271120
e-mail: alex.kacelnik@zoology.oxford.ac.uk

Under natural circumstances, most actions have a set of possible consequences, rather than a single well-defined outcome. A choice is described as risky if the probabilities of the different possible outcomes are known, but the precise outcome is not known. The best that an animal can do when faced with a risky choice is to use information about the probabilities of different outcomes to decide which action to perform. For evolutionary biologists committed to the study of decision-making in animals, this raises an intriguing and challenging question: has natural selection produced sensitivity to risk, and if such an adaptation exists, what form does it take? Posed in this form, the question of risk-sensitivity belongs in the normative field of behavioural ecology. However, recently it has become clear that a number of different fields of behavioural research have converged on the problem of how animals respond to risk, each bringing a different perspective on how the behavioural phenomena should be explained. Here, we review some of

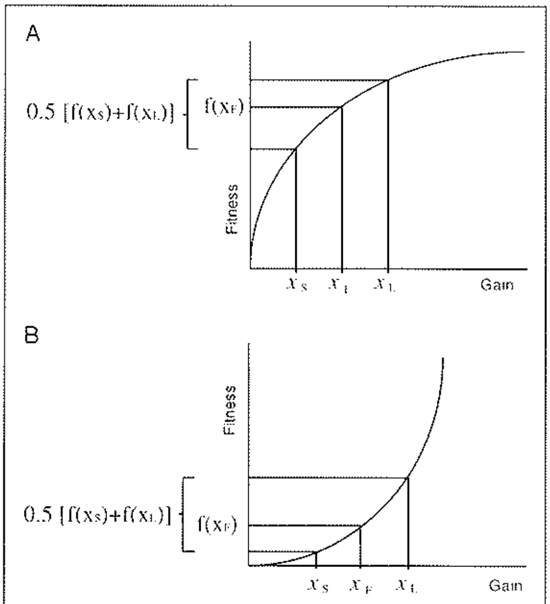

Fig. 1 The normative rationale of Risk-sensitivity Theory. Fitness is shown as either **(A)** a decelerating or **(B)** an accelerating function of food gains. The choice is between two food sources, one always giving x_F food gains and the other giving either a smaller amount, x_S, or a larger amount, x_L, with equal probability. The average food gains of both food sources are equal, but they differ in variance. In the upper panel the average fitness gains are higher for the fixed food source. In contrast, in the lower panel the opposite is true (modified from Ref. 25).

these contributions under the categories of normative, descriptive and process-based models.

Normative modelling: Risk-sensitivity Theory

Taking a normative approach implies establishing what an actor ought to do if it wishes to achieve a goal. Because natural selection must have played a significant role in the evolution of decision systems, and because risk is ubiquitous in nature, we might expect animals to be able to assess and include risk in their decisions, and to behave as if they had the goal of maximizing their darwinian fitness. Thus, there are strong biological justifications for taking a normative approach to risk-sensitivity.

This form of normative theorizing has important heuristic virtues[1–6]. The set of models that address how animals should respond to risk is known collectively as Risk-sensitivity Theory (RST; for reviews see Refs 3, 6 and the special issue of *American Zoologist* Vol. 36, September 1996). RST can be introduced by considering a subject choosing between two actions which yield gains that differ in both average magnitude and variance. The gains in whatever the actions produce (food, for instance) are not considered important in themselves, but because of their effects on darwinian fitness. For presentation purposes, we consider the case where the food gain from one option (fixed) has no variance, while the alternative option (called either risky or variable) yields either of two food gains with equal probability. If the fitness value of food gains of magnitude x follows a function f(x), then to maximize expected fitness the subject should prefer the fixed over the variable option when the fitness gained from the fixed option is greater than the average of the two possible fitness gains from the variable option. This can be expressed as follows:

$$f(x_F) > \frac{1}{2}[f(x_S) + f(x_L)] \qquad (1)$$

where the suffixes F, S and L indicate gains in the fixed option and the smaller and larger outcomes in the variable option, respectively. Notice that the equation is formulated for two equiprobable outcomes in the variable option, but the idea holds if the outcomes are weighted by different probabilities or there are more than two possible outcomes in the variable option.

Figure 1 shows that the inequality in Eqn 1 depends upon the shape of f(x) as well as on the magnitude of the outcomes and the probabilities of the outcomes. For the particular case when the gain from the fixed option is equal to the average gain from the variable option, that is,

$$x_F = \frac{1}{2}(x_S + x_l),$$

the inequality is true if f(x) is decelerated and is in the opposite direction when f(x) is accelerated (this property applies a result known as Jensen's inequality). Thus, the subject should prefer the fixed outcome (be 'risk-averse') if progressive increases in food gain give declining increases in value (Fig. 1A), prefer the risky option (be 'risk-prone') if progressive increases in food gain yield increasing increases in value (Fig. 1B) and be indifferent if f(x) is a linear function (not shown). Clearly, the shape of the function relating food gains to fitness, f(x), should have been crucial in determining the evolution of risk-sensitive decision-making. To assess the relevance of RST today we need to know the form taken by f(x) through evolutionary history. However, because the past relationship between food gains and fitness is inaccessible to research, evolutionary biologists are forced to use indirect approaches to discover the shape.

One possibility is to examine how food gains increase fitness under present circumstances and assume that these circumstances reflect those prevalent in evolutionary times. While being theoretically possible, this approach is not easy to implement because the time-scale of fitness measurements tends to be much greater than that of foraging decisions, making it very difficult to assign (with any accuracy) differences in fitness to differences in foraging strategy. Alternatively, the shape of the utility function by which the animal attributes subjective value to resource gains can be mapped experimentally. If it is assumed that the utility function was shaped by selection to reflect consequences in terms of darwinian fitness, then we can assume that the utility function reflects the shape of f(x) accurately[1,5]. Finally, the problem can be reasoned around by speculating on the most likely shape of the function. None of these three strategies is capable of producing results that refute or confirm RST, but if any of these strategies suggests that f(x) may have been non-linear then we gain confidence in the rationale underlying RST.

Usually, evolutionary biologists take the third strategy of predicting the shape of the function relating gains to fitness on the basis of reasoning. The shape of f(x) is likely to depend on the 'energy budget', or more generally, the resource budget of the subject, which describes the relationship between the subject's needs and the average gain it is currently receiving. If the subject is in a situation where the minimum gain required to survive lies between x_F and x_L, then the probability of survival when the food gains are x_S or x_F is zero, and the function is discontinuous but accelerated when evaluated at x_S, x_F and x_L. If, instead, a survival threshold lies between x_S and x_F and additional gains above this threshold make little difference to fitness, f(x) will be discontinuous but decelerated over the three possible values of x. The 'budget rule'[2–7] summarizes the optimum decision in each of these situations; it predicts risk-aversion when the fixed (or least variable option) provides enough gains to reach the critical threshold for survival and risk-proneness when it does not. There are exceptions to the budget rule, and fuller analyses can be found in more technical theoretical papers[3,8]. It should be noted that the budget rule may be of limited generality because it is possible that over evolutionary time organisms of a given species did not find themselves in circumstances where the ability to switch risk-preference gave a significant gain in fitness over persistently being risk-averse or risk-prone. McNamara[8] explored this issue and showed that while the ability to shift preference would always be best, under many scenarios, risk-aversion is almost as good, but inflexible risk-proneness can produce major fitness losses. The implication of this analysis is that if flexibility is not an option, then risk-aversion is best.

Thus, RST makes two main predictions. The first prediction is that subjects ought to be 'risk-sensitive' in a weak sense, namely, preferences should be affected by variance in gains and not just by average gain. The second, and stronger prediction is that preference should shift from risk-proneness to risk-aversion in response to experimental manipulations that change the shape of f(x) by altering the energy budget of the subject. We label this as a strong prediction because it is unique to RST. Evidence supporting this prediction constitutes strong evidence in favour of RST, however, failure to find support for this prediction should not be taken as evidence against RST because of the potential lack of generality of the assumptions leading to it.

We looked for evidence for and against these predictions in a review of 59 experimental studies of risk-sensitivity in non-humans. We divided reports according to whether or not they included manipulation of the subjects' energy budgets and also according to whether they used amount of gain or the delay to get it to generate risk[6]. The results are shown in Fig. 2. The main results are: (1) the majority of studies found risk-sensitivity in the weak sense; (2) when risk was effected through delay, animals were risk-prone and budget manipulations had no effect; and (3) when risk was effected through amount, there is a tendency towards risk-aversion, and of the studies that manipulated budget, about one-third found the switch in preference predicted by the budget rule. Note that, although RST predicts that risk in amount and delay should not be functionally equivalent[9], the theory does not predict the pattern of results that we found.

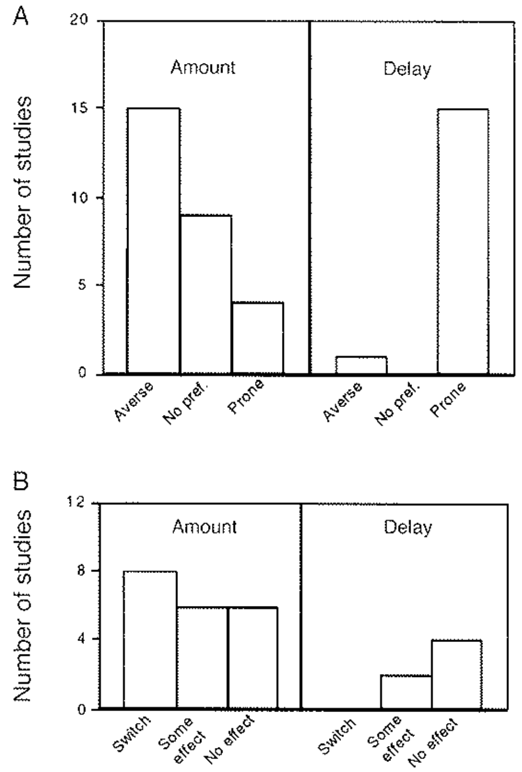

Fig. 2 The figures show the results of our survey of risk-sensitivity experiments. They are based on data from 28 animal species including insects, fish, birds and mammals. **(A)** Results from studies with no budget manipulations divided according to whether variance was in amount of gain or delay to gain. **(B)** Results from studies that manipulated budget, again divided according to whether variance was in amount or delay. 'Switch' indicates that the subjects switched from being risk-averse to risk-prone, according to the predictions of RST. 'Some effect' indicates that there was a shift in preference in the predicted direction but it was not statistically reliable. 'No effect' indicates that manipulations of budget did not change the direction or degree of preference (modified from Ref. 25).

Descriptive models

The behaviourist school of experimental analysis of behaviour prefers to make predictions by induction, generalizing from observed data. The extra-empirical theorizing of the kind described in the previous section goes against this spirit. Observed regularities in behavioural data are used to generate 'Laws' of behaviour (such as the influential Matching Law)[10,11] which in turn are used to generate predictions about how an animal will behave in so far unstudied circumstances. Such models are defended for their descriptive performance and not for their evolutionary or cognitive basis.

In the case of risk-sensitivity, behavioural analysts have studied the effects of variance in delay on choice with greater assiduity than the effects of variance in amount. They found that the following expression gives a good index of reinforcing value and, consequently, choice[12]:

$$V = \frac{A}{1 + kd} \qquad (2)$$

where A is proportional to reward amount, d is the delay to food and k is a parameter that is greater than zero. The essence of Eqn 2 is that value increases linearly with food amount and declines hyperbolically with delay. To apply Eqn 2 to risk-sensitivity, one needs to assume that an action yields a variable amount or delay and then postulate an averaging rule to compute the average value resulting from choosing the variable option. One possible averaging rule for an action with several outcomes is to calculate the mean across these outcomes and then apply Eqn 2 to the resulting mean amount or delay. An alternative is to apply Eqn 2 to each outcome and then calculate the mean of the resulting values. Because Eqn 2 is linear in amount and hyperbolic in delay, it does not make any difference which averaging rule is chosen when dealing with variable amounts, but it does when dealing with variable delays. If delays are averaged first and Eqn 2 is applied to the mean delay then reward sources are valued according to the mean of their outcomes, regardless of variance, but if the second rule is used, then a more variable food source has greater mean value than a fixed one with the same mean delay.

Thus, this approach predicts neutrality to variance in amount and either neutrality (first averaging rule) or risk-proneness (second averaging rule) when variability is in delay. There is no *a priori* reason to prefer one averaging rule over the other. If overall rate of reinforcement was paramount, then the first rule would be expected[13]. However, the second rule fits the data more closely[13–16], as can be seen from Fig. 2A.

Given that Eqn 2 is derived from data, it is hardly surprising that its predictive performance is good, but for the same reasons we feel that this research programme offers little explanatory power.

Process-based model 1: associative learning

Neither the normative nor the descriptive models described above make any effort to suggest how behaviour might be implemented by a plausible information-processing system. In contrast, this is precisely the goal of the models we have placed in this section and the next. These process-based models contain hypotheses about how individuals process information using 'cognitive' entities, such as associations between stimuli or representation of past experience. The model described in this section focuses on the processes by which the experimental subject acquires information about the options on offer.

Subjects in risk-sensitivity experiments are trained to assign significance to available alternatives in the same way as those in psychological studies of conditioning. They are exposed repeatedly to initially neutral stimuli (conditional stimuli; CS), followed by meaningful events (unconditional stimuli; US), such as food. In a typical risk experiment two CSs are used: one paired with a fixed outcome and the other paired with a probabilistic variable outcome. During testing, previously trained animals choose between the two CSs. Normative modellers ignore training because they make the assumption that natural selection has equipped animals to acquire the information they need in order to behave adaptively. However, learning cannot be instantaneous because, in the case of a variable option, a subject will have to experience it many times in order to extract the probabilities of the different outcomes. Thus, the rules by which the subjects learn what is associated with the different CSs may have some role in explaining the choices they make when presented with a pair of CSs. Specifically, we assume that the factors which promote learning about a CS will also promote preference for that CS in a choice test[6].

The shape of the functions relating amount and delay to learning is a lively topic[17] but, overall, learning performance is an increasing function of amount of reward and a decreasing function of delay between CS and US. Crucially, with respect to risk, neither of these two functions is linear. Increases in amount or delay of a fixed outcome provoke marginally decreasing changes in learning. To apply these ideas to training using stimuli with variable outcomes, again, we are faced with the problem of postulating an averaging rule, but now the learning process suggests a clear preference between averaging algorithms. The fundamental importance of the US in learning provides a justification for attributing value first and then computing averages. This makes sense because value should be gained by the CS each time a US is experienced (the control of learning by the US can be extended to extinction by postulating that the CS loses value when an expected US does not occur).

Given the non-linear effects of amount and delay, and the hypothetical averaging rule driven by the US, Jensen's inequality results in our predicting better learning performance for fixed than for variable amounts, but better learning performance for variable than for fixed delays. Figure 3 illustrates the latter case. These predictions are derived from learning functions with single outcomes, and the accuracy of the predictions for variable outcomes is not yet proven empirically, but as it stands, the model correctly predicts risk-aversion when amounts are variable and risk-proneness when delays are variable.

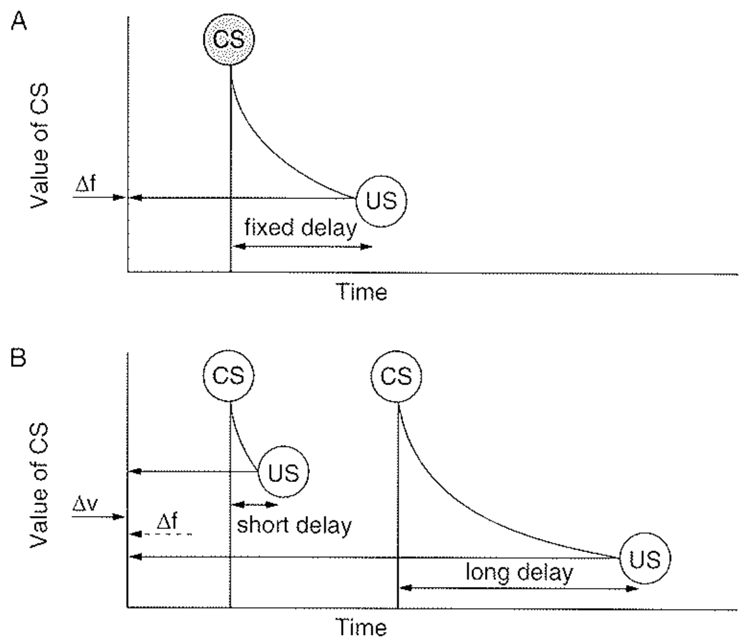

Fig. 3 A simple model showing how associative learning can produce risk-sensitivity when variance is in delay. (A) Shows a putative process for attributing value to a previously neutral stimulus (conditioned stimulus; CS) that is paired to a meaningful event (unconditioned stimulus; US) after a fixed delay f. The idea is that occurrence of the CS leaves a trace that decays non-linearly with time. When the US is experienced the CS gains a value Δf proportional to the state of its trace. **(B)** Uses the same model for a CS paired to a US that occurs half the times after a shorter delay and half the times after a longer delay, with mean f. In this case, the average change in the value of the CS is Δv. The CS associated with the variable delay should be preferred over the CS associated with the fixed delay because after similar amounts of exposure to the two options the CS associated with the variable delay will have a higher value. To account for the effects of variability in amount, this model needs to assume that the CS gains value as an increasing, decelerated function of magnitude of the US (modified from Ref. 25).

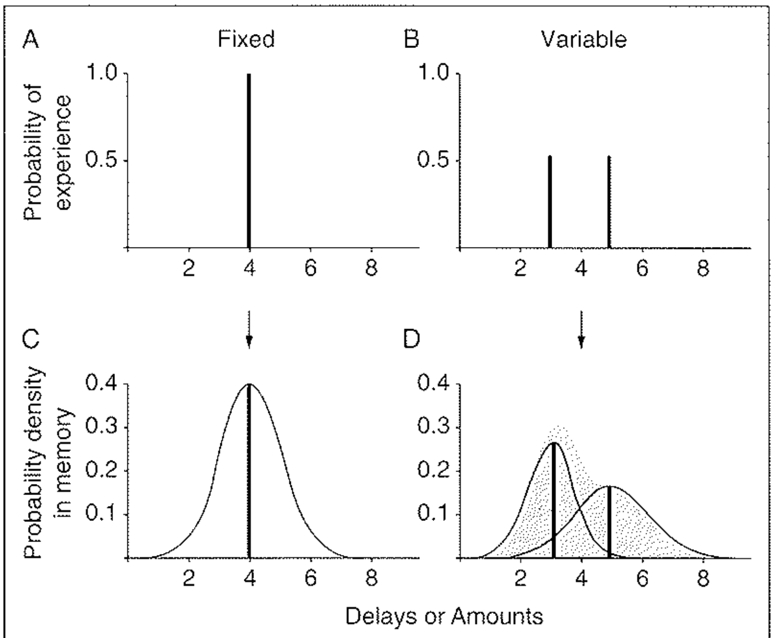

Fig. 4 Reboreda and Kacelnik's model for the effect of Weber's Law on risk sensitivity[22]. **(A,B)** Represent the experienced distribution of outcomes (these may be either delays or amounts) in (A) a fixed and (B) a variable option. **(C,D)** Represent the distributions that are assumed to be formed 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 differing accuracy with which the smaller and larger outcomes are represented. If choices result from comparing random samples from the two memory distributions, then in a majority of cases the sample from the symmetrical distribution will be larger. In the case of a choice between a fixed and variable food amount, this leads to risk-aversion, and in the case of a choice between a fixed and variable delay it leads to risk-proneness (reproduced from Ref. 14).

Process-based model 2: extending Scalar Expectancy Theory

So far, the models that we have described all assume that amounts and delays are perceived accurately and remembered without error. In this section we investigate the consequences of assuming a particular type of perceptual error. In its classical form, Weber's Law implies that bigger differences are required to discriminate stimuli of greater magnitude. Scalar Expectancy Theory (SET)[18,19] accounts for this phenomenon in the measurement of time intervals by assuming that time intervals are remembered with an accuracy proportional to their length. This assumption is supported by data from experiments in which subjects reproduce time intervals by showing an increase in food-related behaviour towards the typical time at which food is received. Timing shows a drop in absolute, but not relative, accuracy with the magnitude of the interval being timed. Specifically, in SET John Gibbon and collaborators[20,21] have postulated that when a subject stores the magnitude of a fixed interval in its reference memory it forms a bell-shaped probability density function with mean and standard deviation proportional to the magnitude of the interval. If a stimulus is associated with more than one interval, what is stored in reference memory is the sum of as many bell-shaped functions as there are intervals, each one with an area proportional to the relative frequency of that of the time interval being represented[21].

Questions for future research

- Strong risk-sensitivity is crucial for Risk-sensitivity Theory (RST) but is supported by weak evidence. Do the shifts in risk-preference for both amount and delay predicted by RST really occur? Perhaps smaller species, or species from harsher environments, are more likely to have evolved strong risk-sensitivity. Are there reliable differences in which species show the shifts?
- So far, we have not distinguished between different types of variability, however, this distinction is crucial: RST applies to variable outcomes that are also unpredictable (or risky), whereas the other models we discuss make similar predictions for both predictable (for example, alternation of good and bad outcomes) and unpredictable variance. Is risk-sensitivity controlled by the stochasticity of outcomes as assumed by RST or do animals treat predictable variability in the same way? Worryingly, our first stab at this issue[24] favours the latter.
- Is classical conditioning with food rewards of variable magnitude less effective than conditioning with fixed rewards of the same mean size? And is conditioning with a variable CS–US delay more effective than with a fixed delay of the same mean duration? Both of these predictions follow from the associative learning model of risk-sensitivity[6] but they have not been tested empirically.
- Can the extension of Scalar Expectancy Theory to food amounts[14,22] predict quantitative details of preferences?
- Finally, and more generally, can we devise a satisfactory epistemological framework to incorporate the contributions of cognitive science into normative models of behaviour?

Reboreda and Kacelnik extended the approach of Gibbon to the memory for food amounts[22] and claimed that SET might have a bearing on risk-sensitivity for both amounts and delays. Experimentally mapping the shape of the cognitive representation of food amounts is more difficult than for time intervals because there is no straightforward equivalent of a reproduction task. However, using discrimination tasks, it is possible to show that Weber's Law does apply to the perception of the size of fixed rewards[14,23] and the underlying representation of amounts could, therefore, be similar to the assumptions that SET makes for times. The description of the model that follows is assumed to apply to both amount and time.

According to the summation principle described above, the representation of an option with a variable outcome is skewed to the right because of the greater variance in the representation of the larger elements, whereas that of an option with a fixed outcome is symmetrical (Fig. 4). To apply these ideas to risk, we assume that the subject chooses between options by comparing random samples from the memory representations of the options. If two options have the same mean (amount or delay) but one has a fixed and the other a variable outcome, then the positive skew in the representation of the variable option implies that in more than half of the comparisons the variable option will yield a smaller sample than the fixed one. This process accounts for risk-proneness when delays are variable[21], and it extends naturally to predict risk-aversion when amounts are variable, because foraging animals prefer larger food amounts and shorter delays. The appeal of this model is derived from its economy: the same assumptions about the shape of cognitive representations account for the contrasting effects of delay and amount variability, a simplicity that none of the other models achieves.

Conclusions

We began this review by claiming that the problem of risk is likely to be ubiquitous in the evolution of animal decision systems, and went on to describe four areas of behavioural research, each of which addresses the problem of how risk affects decision making. The four theories we describe need not compete directly with each other because they address different questions: RST addresses the fitness consequences of choosing more or less risky options, the descriptive model seeks an algorithm that best describes the available data and the process-based theories address the cognitive mechanisms that underlie risk-sensitivity. We do not believe that any of these approaches substitutes for any other; all can contribute to a full understanding of risk-sensitivity. However, while theories offering various levels of explanation can co-exist, they should not exist in isolation, and should not ignore empirical findings from other fields of research. To date, the different approaches to risk-sensitivity have focused on explaining different aspects of the data and, consequently, they do not come together to form a single coherent explanation.

RST is unique in predicting shifts in preference with subjects' resource budgets (something we have called strong risk-sensitivity), but the evidence for these shifts is not yet conclusive, coming from a few studies and few species. If strong risk-sensitivity is confirmed, then the descriptive and process-based theories will have to be modified accordingly. By the same token, it seems necessary to adjust RST to cope with the dramatic difference in the effects of variability in amount and delay which the process-based theories explain easily. RST does show that variability in these two attributes has different fitness implications, but the behavioural predictions are so qualitative that experimental findings make little or no impact on the theory itself. As claimed by one of the best normative modellers working on risk: 'these models are better for thinking than for testing'. However, greater responsiveness to the results of empirical testing would do little harm.

Acknowledgements

Many of the ideas in this paper were developed in collaboration with John Gibbon, Dani Brunner and Juan Reboreda. Funding was provided by NATO grant 920308, Wellcome Trust grant 0461/Z/95 and a Wellcome Advanced Training Fellowship to MB.

References

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

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

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

4 Real, L.A., Ott, J. and Silverfine, E. (1982) On the tradeoff between mean and variance in foraging: an experimental analysis with bumblebees *Ecology* 63, 1617–1623

5 Kacelnik, A. (1997) Normative and descriptive models of decision making: time discounting and risk sensitivity, in *Characterizing Human Psychological Adaptations*, CIBA Symposium 208 (Bock, G. and Cardew, G., eds), pp. 51–70, Wiley

6 Kacelnik, A. and Bateson, M. (1996) Risky theories – the effects of variance on foraging decisions *Am. Zool.* 36, 402–434

7 Stephens, D.W. and Krebs, J.R. (1986) *Foraging Theory*, Princeton University Press

8 McNamara, J.M. (1996) Risk-prone behaviour under rules which have evolved in a changing environment *Am. Zool.* 36, 484–495

9 McNamara, J.M. and Houston, A.I. (1987) A general framework for understanding the effects of variability and interruptions on foraging behaviour *Acta Biotheoret.* 36, 3–22

10 Herrnstein, R.J. (1970) On the law of effect *J. Exp. Anal. Behav.* 13, 243–266

11 Davison, M. and McCarthy, D. (1988) *The Matching Law: a Research Review*, Erlbaum

12 Mazur, J.E. (1984) Tests of an equivalence rule for fixed and variable reinforced delays *J. Exp. Psychol. Anim. Behav. Proc.* 10, 426–436

13 Bateson, M. and Kacelnik, A. (1996) Rate currencies and the foraging starling: the fallacy of the averages revisited *Behav. Ecol.* 7, 341–352

14 Bateson, M. and Kacelnik, A. (1995) Preferences for fixed and variable food sources: variability in amount and delay *J. Exp. Anal. Behav.* 63, 313–329

15 Mazur, J.E. (1987) An adjusting procedure for studying delayed reinforcement, in *Quantitative Analyses of Behaviour: the Effect of Delay and of Intervening Events on Reinforcement Value* (Commons, M.L. *et al.*, eds), pp. 55–73, Erlbaum

16 Real, L., Ellner, S. and Harder, L.D. (1990) Short-term energy maximisation and risk-aversion in bumble bees: a reply to Possingham *et al. Ecology* 71, 1625–1628

17 Tarpy, R. (1997) *Contemporary Learning Theory and Research*, McGraw-Hill

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

19 Gibbon, J. (1991) Origins of scalar timing *Learn. Motiv.* 22, 3–38

20 Gibbon, J. and Church, R.M. (1984) Sources of variance in an information processing theory of timing, in *Animal Cognition* (Roitblat, H.L., Bever, T.G. and Terrace, H.S., eds), pp. 465–488, Erlbaum

21 Gibbon, J. *et al.* (1988) Scalar Expectancy Theory and choice between delayed rewards *Psychol. Rev.* 95, 102–114

22 Reboreda, J.C. and Kacelnik, A. (1991) Risk sensitivity in starlings: variability in food amount and food delay *Behav. Ecol.* 2, 301–308

23 Bateson, M. and Kacelnik, A. (1995) Accuracy of memory for amount in the foraging starling (*Sturnus vulgaris*) *Anim. Behav.* 50, 431–444

24 Bateson, M. and Kacelnik, A. (1997) Starlings preferences for predictable and unpredictable delays to food *Anim. Behav.* 53, 1129–1142

25 Bateson, M. and Kacelnik, A. Risk-sensitive foraging: decision making in variable environments, in *Cognitive Ecology* (Dukas, R., ed.), Chicago University Press (in press)