

Choice between small certain and large uncertain reinforcers

JAMES E. MAZUR

Harvard University, Cambridge, Massachusetts

When choosing between two alternatives that deliver the same amount of food per trial in the long run, organisms are called *risk-averse* if they choose a small certain reinforcer over a larger probabilistic reinforcer. They are called *risk-prone* if they choose the larger probabilistic reinforcer. This experiment attempted to predict whether rats would be risk-prone or risk-averse on the basis of their separate choices between reinforcers differing in probability and reinforcers differing in amount. Choice was measured with an adjusting-delay procedure, which provided estimates of indifference points, or pairs of alternatives that a subject chose about equally often. The subjects were usually more responsive to differences in amount than to differences in probability, leading to predictions of risk-proneness for choices between two probability-amount combinations. The predictions were confirmed in almost every case. As the number of food pellets delivered by the two alternatives was increased while maintaining a 2:1 difference between them, the tendency toward risk-proneness declined. These results suggest an explanation of the inconsistent findings obtained in previous experiments on risk-taking by rats.

Animals in the wild must frequently make choices when the consequences of different courses of action are uncertain. The same is true, of course, for many of the choices people must make in everyday life. Perhaps because of the ubiquity of decision-making under conditions of uncertainty, many recent studies have examined risk-taking by animal subjects. Much of this research has involved choices between two food reinforcers—a small reinforcer that is certain to occur and a larger reinforcer than occurs with some probability. The probabilities and amounts are usually chosen so that the two alternatives have equal expected values—they would be expected to deliver the same amount of food per trial in the long run. Thus, if the subject's goal is simply to maximize the amount of food received over some long period of time, it should be indifferent between the two alternatives. In practice, however, subjects frequently exhibit a clear preference for one alternative or the other. A subject is said to demonstrate *risk-aversion* if it shows a preference for the small, certain reinforcer and *risk-proneness* if it shows a preference for the larger, uncertain reinforcer.

This choice situation has been studied with species as diverse as shrews (e.g., Barnard & Brown, 1985a), juncos (e.g., Caraco, 1981), and bees (Real, Ott, & Silverfine, 1982; Waddington, Allen, & Heinrich, 1981). Much of this research has examined factors that determine when animals will show risk-aversion versus risk-proneness. For instance, Caraco and his colleagues have found that juncos and sparrows tend to be risk-prone when food is

scarce and risk-averse when food is abundant (e.g., Caraco, 1983; Caraco & Chasin, 1984; Caraco, Martindale, & Whittam, 1980).

In this work on animal risk-taking, there has been little discussion of how a subject's choices might depend on its sensitivity to the two separate dimensions that distinguish risky and certain alternatives—probability and amount of reinforcement. A tendency toward risk-proneness could occur, for instance, because a subject is relatively insensitive to differences in reinforcer probability. Just because the objective probabilities of two reinforcers differ by a factor of two, their subjective probabilities do not necessarily differ by a factor of two. Indeed, research with human subjects has frequently found systematic discrepancies between objective and subjective probabilities (e.g., Kahneman & Tversky, 1979), and there may well be similar discrepancies with animal subjects (see Mazur, 1985). Similarly, if two reinforcers differ in amount by a factor of two, this does not necessarily imply a twofold difference in their *subjective utilities* (a term commonly used in utility theory to denote the psychological effect of a reinforcer of a given objective size).

Imagine a situation in which a subject shows a preference for two food pellets delivered with a probability of 0.5 over one pellet delivered with a probability of 1.0. From the standpoint of utility theory (see Lee, 1971, chap. 4), there are several possible reasons for this preference. One is that, for this subject, a twofold difference in objective probabilities results in less than a twofold difference in subjective probabilities. Another possibility is that a twofold difference in amount results in more than a twofold difference in subjective utilities. A third possibility is that both of the above are true. In some of the previous studies on animal risk-taking, these differ-

This research was supported by Grant MH 38357 from the National Institute of Mental Health. I thank David Coe for running the animals in this experiment. Correspondence concerning this article should be sent to James E. Mazur, now at Psychology Department, Southern Connecticut State University, New Haven, CT 06515.

ent possibilities were not discussed (e.g., Barnard & Brown, 1985b; Waddington et al., 1981). In others, it was assumed from the outset that one of these possibilities was correct. For instance, Caraco et al. (1980) used their data to derive functions relating objective amount of food and subjective utility, and they obtained different functions for conditions of food scarcity versus food abundance. However, these functions were based on the assumption that subjective and objective probabilities were equal in all conditions. It is certainly possible that what differed between the scarcity and abundance conditions in the Caraco et al. experiment was not the utility function for different reinforcer amounts, but the subjective probability function. The more general point, however, is that a better understanding of the determinants of risk-proneness and risk-aversion may result from research on the two distinct dimensions of reinforcer probability and reinforcer amount.

The purpose of the present experiment was to determine if it is possible to make predictions of risk-proneness or risk-aversion on the basis of subjects' responses to reinforcer probabilities and amounts when these two variables are examined separately. The subjects' choices were measured with an adjusting-delay procedure that has been used in a number of previous experiments in which reinforcer probabilities and amounts were studied (e.g., Mazur, 1984, 1985, 1987). In this procedure, the subjects choose between a *standard alternative*, which had a constant delay between a choice response and reinforcement, and an *adjusting alternative*, for which the reinforcer delay was increased or decreased many times a session. The purpose of these adjustments was to estimate an indifference point—a delay at which the two alternatives would be chosen equally often.

The following example shows how this procedure was used to make predictions of risk-proneness or risk-aversion. Suppose that in the first condition, the standard alternative is one pellet delivered with a probability of 0.5 after a delay of 5 sec. The other alternative delivers one pellet with a probability of 1.0 after an adjusting delay, and the subject exhibits indifference when the adjusting delay is 10 sec. This result suggests that a twofold difference in delays compensates for a twofold difference in probability. In a second condition, the alternatives differ in amount. The standard alternative delivers one pellet with a probability of 1.0 after a 5-sec delay. The adjusting alternative delivers two pellets with a probability of 1.0, and the adjusting delay at the indifference point is 15 sec. This result suggests that a threefold difference in delays is needed to compensate for a twofold difference in amount. Taken together, these two results suggest that the subject is less sensitive to the twofold difference in probabilities than to the twofold difference in amount.

In the third condition, the standard alternative is one pellet after 5 sec with a probability of 1.0. The other alternative is two pellets with a probability of 0.5 after an adjusting delay. Because the results of the first two con-

ditions suggested that this subject is more sensitive to the differences in the two amounts than to the differences in probability, we can predict that the indifference point will be greater than 5 sec, indicative of a preference for the adjusting alternative. This would be called risk-proneness. (This discussion assumes that if the subject were indifferent to the two probability-amount combinations, the indifference point would be 5 sec, the same as the standard delay. Of course, the subject could have a bias for one of the two alternatives, based on a position preference or some other factor. If so, this bias would have to be taken into account in making predictions for the third condition.)

The present experiment included several such tests, to determine whether risk-proneness or risk-aversion could be predicted from a subject's separate responses to the reinforcer probabilities and amounts. A second purpose of the experiment was to determine how the subjects' choices would be affected when the amounts of reinforcement for the two alternatives were both increased by the same proportion. Across different conditions, the number of pellets delivered by the certain and risky alternatives was varied by a factor of four.

METHOD

Subjects

The subjects were three male hooded rats obtained from the Charles River Breeders. They were about 9 months old at the start of the experiment, and were maintained at about 80% of their free-feeding weights.

Apparatus

The experimental chamber was 23 cm long, 20 cm wide, and 19 cm high. Two response levers, each 5 cm long and protruding 1.5 cm into the chamber, were mounted in the front wall of the chamber. Each lever was 1 cm from the side wall and 9 cm above the metal grid floor, and required a force of approximately 0.1 N to operate. Two 2-W lights were mounted in the wall 5 cm above each lever. For the left lever, the left light was covered with a red plastic cap and the right light by a white cap. These colors were reversed for the lights above the right lever. A third, similar lever was mounted in the center of the rear wall of the chamber, 9 cm above the floor. There were also two 2-W lights above this lever, and both were covered with white plastic caps. Two 2-W red house-lights were mounted in the ceiling of the chamber. A pellet dispenser delivered 45-mg food pellets into a food tray located in a 4-cm square opening in the center of the front wall of the chamber, 1 cm above the floor.

The chamber was enclosed in a sound-attenuating box that contained an air blower for ventilation and a speaker producing continuous white noise to mask extraneous sounds. A PDP-8 computer in another room was programmed in SuperSKED to control the stimuli and record responses.

Procedure

The subjects had previous experience with the adjusting-delay procedure in another experiment conducted in a similar chamber, so no pretraining was necessary. Throughout the experiment, the left lever served as the standard alternative, and the right lever served as the adjusting alternative. The experiment consisted of 15 conditions, which differed in the numbers of pellets and the probabilities of reinforcement associated with these two alternatives. Every

session lasted for 64 trials or for 90 min, whichever came first. Within a session, each block of four consecutive trials consisted of two forced trials followed by two choice trials. The procedure on choice trials will be described first.

At the start of a choice trial, the lights above the rear lever were lit; a single response on this lever was required to begin the choice period. The rear-lever response requirement made it more likely that the animal would be an equal distance from both side levers when the choice period began. A response on the rear lever darkened the lights above this lever and illuminated the lights above the two side levers. A single response on the left lever darkened the lights above both levers and initiated the standard delay of 5 sec, during which the lights above the left lever flashed on and off two times per second (0.25 sec off, 0.25 sec on). At the end of the standard delay, the left lever lights stopped flashing and either (1) the number of pellets associated with the standard alternative were delivered, and then an intertrial interval (ITI) began, or (2) no pellets were delivered, and the ITI began. In either case, the next trial started 60 sec after the choice response of the current trial.

If the subject responded on the right (adjusting) lever during the choice period, the lights above both side levers were darkened and the adjusting delay began, during which the lights above the right lever flashed on and off two times per second. At the end of the adjusting delay, the lights stopped flashing and either (1) the number of pellets associated with the adjusting alternative were delivered, and the ITI began, or (2) no pellets were delivered, and the ITI began. As with the standard alternative, the next trial started 60 sec after the choice response.

The procedure on forced trials was the same as that on choice trials except that after the rear-lever response, only the lights above one side lever were lit; a response on this lever had the same effect as on a choice trial. A response on the opposite lever had no effect. Of every two forced trials, one involved the standard lever and the other, the adjusting lever. The temporal order of these two types of trials was varied randomly.

After every two choice trials, the delay for the adjusting alternative could change. If a subject chose the adjusting alternative on both choice trials, the adjusting delay was increased by 1 sec. If the subject chose the standard alternative on both trials, the adjusting delay was decreased by 1 sec unless it was already zero. If a subject chose each alternative once on the two choice trials, no change was made in the adjusting delay. In all three cases, this adjusting delay remained in effect for the next block of four trials. At the start of the first session of the experiment, the adjusting delay was set equal to 0 sec. For the start of every other session of the condition, the adjusting delay was determined by the above rules as if the first trial were a continuation of the preceding session.

For each condition, Table 1 shows the number of pellets and the probability of reinforcement for both the standard and adjusting alternatives. For all cases in which the probability of reinforcement was less than 1.0, reinforcer delivery was determined by pseudo-random sequences that ensured that the actual probability of reinforcement approximately equaled the nominal probability in each session. In Conditions 1 and 12, both alternatives delivered one pellet with a probability of 1.0. These conditions were used to measure possible position bias or bias for either the standard or adjusting delay. In five conditions (2, 3, 6, 13, and 15), the two alternatives delivered the same number of pellets, but the probability of reinforcement was 1.0 for the adjusting alternative and less than 1.0 for the standard alternative. In four conditions (4, 7, 11, and 14), both alternatives had reinforcement probabilities of 1.0, but the adjusting alternative delivered more pellets than the standard alternative. In the remaining four conditions (5, 8, 9, and 10), the adjusting alternative delivered more pellets, but its probability of reinforcement was lower, so the average number of pellets per trial was the same for the two alternatives. Note that a few conditions

Table 1
Reinforcement Probabilities (P), Number of Pellets per Reinforcer (N), and Obtained Adjusting Delays in Each Condition

Condition	Standard		Adjusting		Adjusting Delay (sec)		
	N	P	N	P	Subject 1	Subject 2	Subject 3
1	1	1.00	1	1.00	5.65	7.55	5.25
2	1	0.20	1	1.00	34.88	24.96	19.42
3	1	0.50	1	1.00	13.61	12.17	11.46
4	1	1.00	2	1.00	23.92	21.25	17.38
5	1	1.00	2	0.50	7.44	14.17	14.54
6	1	0.25	1	1.00	15.52	32.59	13.90
7	1	1.00	4	1.00	50.77	34.23	32.42
8	1	1.00	4	0.25	10.04	11.65	16.79
9	2	1.00	4	0.25	8.46	9.25	13.36
10	4	1.00	8	0.50	3.98	6.65	8.69
11	4	1.00	8	1.00	17.52	13.32	9.69
12	1	1.00	1	1.00	6.08	9.48	10.77
13	4	0.50	4	1.00	21.02	7.96	7.57
14	2	1.00	4	1.00	33.52	12.33	13.84
15	2	0.50	2	1.00	21.05	6.00	10.63

Note—The order of Conditions 6–8 was 8, 6, 7 for Subject 1 and 7, 8, 6 for Subject 3.

were run in a different order for the 3 subjects, as explained in Table 1.

For each of the first three conditions, there were a minimum of 20 sessions; for all other conditions, there were a minimum of 12 sessions. After the minimum number of sessions, a condition was terminated for each subject individually when several stability criteria were met. To assess stability, each session was divided into two 32-trial blocks, and the mean delay for the adjusting alternative in each block was calculated. The results from the first two sessions of a condition were not used, and a condition was terminated when the following three criteria were met, using the data from all subsequent sessions: (1) Neither the highest nor the lowest single-block means of a condition could occur in the last six blocks of the condition. (2) The mean adjusting delay across the last six blocks could not be the highest nor the lowest six-block mean of the condition. (3) The mean delay of the last six blocks could not differ from the mean of the preceding six blocks by more than 10% or by more than 1 sec (whichever was larger).

RESULTS

In conditions lasting a minimum of 20 sessions, the median number of sessions required to satisfy the stability criteria was 23 (ranging from 20 to 27). In conditions lasting a minimum of 12 sessions, the median number of sessions was 15 (ranging from 12 to 25). For each condition, the mean adjusting delay across the six half-session blocks that satisfied the stability criteria was treated as an estimate of the indifference point. These delays are presented for each subject in the right side of Table 1.

The results from Conditions 1 and 12, in which both alternatives delivered one pellet with a probability of 1.0, were used to test for possible bias. Table 1 shows that for every subject in both conditions, the adjusting delay was greater than 5 sec, indicating a bias for the adjusting alternative. This bias might have been the result of a position preference or a preference for the alternative that had adjusting rather than fixed delays. Averaged across

the two conditions, the mean indifference points were 5.87 sec, 8.52 sec, and 8.01 sec for the 3 subjects, respectively. Because the two probabilities and amounts were identical in these conditions, in the analysis of other conditions any indifference points greater than these means were treated as a preference for the adjusting alternative, and those smaller than these means were treated as a preference for the standard alternative.

For each subject, the experiment included four tests similar to the one described in the introduction, with every test consisting of three conditions. In one condition, the two alternatives differed only in reinforcer amount, and in a second, only in reinforcer probability. In a third condition, the standard alternative was a small certain reinforcer and the adjusting alternative was a large uncertain reinforcer. For example, Conditions 3, 4, and 5 constituted one test, and the parameters were identical to those used in the introduction's example. The results from these three conditions are presented in the left column of Figure 1. The striped bars show the indifference points from Condition 3, in which the two alternatives differed

only in reinforcer probability, and the black bars show those from Condition 4, in which the alternatives differed only in the number of pellets. The white bars show the indifference points from Condition 5, in which the two probabilities and amounts were combined, and the dotted horizontal lines show the mean indifference points from Conditions 1 and 12, which serve as a correction for bias. The vertical lines above each bar in Figure 1 are the standard deviations of the six half-session means that comprised each indifference point.

As the left column of Figure 1 shows, the indifference point for each subject was larger when only amount was varied than when only probability was varied, which suggests a greater sensitivity to the differences in amount. Risk-proneness is therefore predicted for Condition 5. Figure 1 shows that this prediction was confirmed: For each subject, the white bars extend above the dotted lines, suggesting a preference for the larger, probabilistic reinforcer.

The other three columns in Figure 1 show the results from the three other tests of this hypothesis. As can be




Figure 1. Each panel shows the indifference points from the three conditions that constituted one test of risk-taking. The vertical line above each bar is the standard deviation of the six half-session means that comprised the indifference point. For each subject, the dotted horizontal line shows the mean indifference point from Conditions 1 and 12, which serves as a correction for bias. Indifference points above these lines indicate a preference for the adjusting alternative, and those below the lines indicate a preference for the standard alternative.

seen, in some conditions there was substantial variability across subjects in the estimated indifference points. Nevertheless, there was sufficient orderliness in the results to yield successful predictions at the ordinal level. In 11 of the 12 tests shown in Figure 1, the indifference point from the amount-only condition was greater than the indifference point from the probability-only condition, leading to a prediction of risk-proneness. This prediction was supported in 10 of the 11 cases, with the only exception occurring for Subject 2 in the third test (third column). In the one case where risk-aversion was predicted because the indifference point was greater in the probability-only condition (Subject 1, third column), this prediction was supported. In summary, 11 of the 12 predictions were correct ($p < .01$, two-tailed sign test).

The accuracy of these predictions depends, of course, on the reliability of the estimates of bias, and there might be some reason to suspect the results obtained in the second test of bias, Condition 12. The substantial bias for the adjusting alternative in this condition might be in part a carry-over from the previous condition, in which the adjusting alternative delivered eight pellets on every trial. However, two additional analyses showed that the predictions were about equally successful if other estimates of bias were used. If only the results from Condition 1 are used to estimate bias, the predictions are successful in 11 of 12 cases (with the same single exception as before). If it is assumed that there is no bias (i.e., that the results from Condition 1 are slightly above 5 sec only because of measurement error), the predictions are successful in




Figure 2. Mean indifference points for the group from conditions in which both probability and amount varied are shown as a function of the percentage difference between indifference points in the corresponding conditions in which only probability or only amount varied. The percentages on the x-axis equal $100(A-P)/P$, where A and P are, respectively, group means of the indifference points from conditions in which only amount varied and in which only probability varied.




Figure 3. Mean indifference points for the group are shown as a function of the number of pellets delivered by the standard alternative. The filled circles are results from conditions in which the adjusting alternative delivered twice as many pellets as the standard. The other two functions are from corresponding conditions in which the standard and adjusting alternatives differed only in probability (open circles), and in both amount and probability (triangles).

all 12 cases. Thus, regardless of which of the three different measures of bias are used, the predictions of risk-proneness or risk-aversion are supported in almost every case.

Figure 2 presents another way to examine the relation between responsiveness to the two separate dimensions and risk-taking behavior. The x-axis plots the percentage change between mean indifference points from conditions in which only amount varied and those from conditions in which only probability varied. That is, these percentages equal $100(A-P)/P$, where A and P are, respectively, indifference points from the condition in which only amount varied and the condition in which only probability varied. The positive percentages indicate that the indifference points were always greater when only amount varied. The y-axis plots the mean indifference points from corresponding conditions in which both probability and amount varied. The monotonic relation in Figure 2 shows that as sensitivity to the difference in amounts increased relative to sensitivity to the difference in probabilities, so did preference for the adjusting alternative, which was composed of the smaller probability and the larger amount.

Figure 3 shows how the mean indifference points for the group varied as a function of the number of pellets delivered by the two alternatives. As the number of pellets

increased from one to four, indifference points from conditions in which the alternatives differed only in probability (open circles) showed no systematic change. (The probabilities were 1.0 and 0.5 for the standard and adjusting alternatives in these three conditions.) A one-way repeated-measures analysis of variance found no significant differences among these three conditions [$F(2,4) = 0.01$]. The filled circles in Figure 3 are the mean indifference points from the three conditions that differed only in amount, with the adjusting alternative delivering twice as many pellets. Despite the apparent decreasing trend shown in Figure 3, a repeated-measures analysis of variance revealed no significant differences among conditions [$F(2,4) = 1.71$]. The failure to reach statistical significance may have been due to the small number of subjects, but it should be noted that for each subject, the indifference point in Condition 11 (four vs. eight pellets) was more than 6 sec smaller than in Condition 4 (one vs. two pellets). The results from the three conditions in which both probability and amount varied by a factor of two (triangles) showed a decreasing trend similar to that of the amount-only conditions. In this case, an analysis of variance revealed a significant effect of the number of pellets [$F(2,4) = 10.28, p < .05$]. In short, as the number of pellets increased proportionately for both alternatives, the subjects exhibited a decrease in risk-proneness.

DISCUSSION

This experiment showed that it is possible to predict risk-taking behavior on the basis of separate measures of a subject's responses to alternatives that differ only in probability and only in amount of reinforcement. In almost every case, predictions made in this way were correct. This result suggests that it may be useful to analyze variables that affect risk-proneness in a similar manner, so as to determine whether any change in choice behavior is due to an altered sensitivity to a change in the subjective utilities of the two reinforcers or to a change in subjective probabilities. For instance, in this study it was found that risk-proneness decreased when the number of pellets delivered by the two alternatives increased, with their proportion held constant at 2:1. Although analyses of variance did not reveal a significant effect for either probability or amount when considered separately, the trends shown in Figure 3 suggest that the decrease in risk-proneness did not occur because of a change in subjective probabilities but rather because the difference between four and eight pellets did not affect the subjects as much as the difference between one and two pellets. This hypothesis is not especially surprising, since amount was the variable that was changed across these conditions. In other cases, however, it is not clear which of these two dimensions is affected by an experimental manipulation. For instance, in the research of Caraco et al. (1980), the main independent variables that produced a change in risk-taking were the number of hours of food deprivation and

the duration of the ITI. It is not obvious whether these variables altered the subjects' responsiveness to differences in amount, to differences in probability, or to both, but the type of analysis presented here could presumably answer this question.

The present results can be compared with those of previous studies on risk-taking that used rats as subjects. Logan (1965) found a preference for five certain pellets over an equal chance of getting either one or nine pellets, which in our terminology constitutes risk-aversion. Battalio, Kagel, and MacDonald (1985) also found evidence for risk-aversion in two experiments. On the other hand, Leventhal, Morrell, Morgan, and Perkins (1959) observed risk-proneness in two groups of subjects and no preference in another two groups. Finally, risk-proneness was the predominant result in the present study, except when large reinforcer amounts were used.

Although these studies make it clear that risk-proneness, risk-aversion, and no preference can all be obtained with rats, their results are not as contradictory as they may appear on the surface. All of these results are consistent with the finding of the present study that risk-proneness decreases with increasing reinforcer amounts. Recall that all subjects were risk-prone when the average number of pellets per trial was one or two, but when the average was four pellets, 2 subjects exhibited risk-aversion and the 3rd showed a shift in the direction of risk-aversion. Logan's (1965) evidence for risk-aversion with an average of five pellets per trial (each weighing 45 mg, as in the present study) is therefore consistent with the present results. The strong tendency for risk-aversion found by Battalio et al. (1985) was obtained with an average of 8 or 10 pellets per trial. In the experiment of Leventhal et al. (1959), two of the groups received food as a reinforcer and the other two received water. The amount of food or water per trial was small for one group and large for the other. The two groups with small reinforcers exhibited risk-proneness, whereas the two groups with larger reinforcers showed no clear preference. All of these results are consistent with the statement that rats' choices shift from risk-proneness to indifference and then to risk-aversion as reinforcer amount increases. The present study suggests that this shift occurs because the subjective difference between two reinforcers that differ objectively by a factor of two decreases as the reinforcers become larger (cf. Commons, Woodford, & Ducheny, 1982).

The degree to which these findings support Caraco's (1981) theory is questionable. He proposed that animals will be risk-prone when their current rate of food intake, as constrained by their environment, would not be sufficient for the subjects to survive in the long run. This situation is called one of negative energy balance. Conversely, he proposed that animals will be risk-averse when a positive energy balance is present (when they have more food than is necessary to maintain their survival). On the surface, the shifts away from risk-proneness with increasing reinforcer amounts that have just been described seem

consistent with Caraco's theory. That is, as reinforcer amounts increase, an animal will be more likely to achieve a positive energy balance. Nevertheless, it is probably not accurate to say that risk-taking in these studies was determined by the subjects' positive or negative energy balances. Logan (1965), Leventhal et al. (1959), and the present experiment all used procedures that kept subjects at 80% of their free-feeding weights, so it might be said that they experienced zero energy balances (neither gaining nor losing weight in the long run). Yet these experiments found a variety of effects ranging from risk-proneness to risk-aversion. Furthermore, in their second experiment, Battalio et al. (1985) explicitly tested Caraco's hypothesis by arranging a positive energy balance in some conditions and a negative energy balance in others. Their subjects were risk-averse in all conditions, and energy balance had no systematic effects on choice. Overall, these experiments thus provide little support for Caraco's hypothesis. It appears more accurate to say that rats' risk-taking tendencies changed as a function of the sizes of the reinforcers used rather than as a function of energy balance.

REFERENCES

- BARNARD C. J., & BROWN, C. A. J. (1985a). Competition affects risk-sensitivity in foraging shrews. *Behavioral Ecology & Sociobiology*, **16**, 379-382.
- BARNARD C. J., & BROWN, C. A. J. (1985b). Risk-sensitive foraging in common shrews (*Sorex araneus* L.). *Behavioral Ecology & Sociobiology*, **16**, 161-164.
- BATTALIO, R. C., KAGEL, J. H., & MACDONALD, D. N. (1985). Animals' choices over uncertain outcomes. *American Economic Review*, **75**, 597-613.
- CARACO, T. (1981). Energy budgets, risk and foraging preferences in dark-eyed juncos (*Junco hyemalis*). *Behavioral Ecology & Sociobiology*, **8**, 213-217.
- CARACO, T. (1983). White-crowned sparrows (*Zonotrichia leucophrys*): Foraging preferences in a risky environment. *Behavioral Ecology & Sociobiology*, **12**, 63-69.
- CARACO, T., & CHASIN, M. (1984). Foraging preferences: Response to reward skew. *Animal Behavior*, **32**, 76-85.
- CARACO, T., MARTINDALE, S., & WHITTAM, T. S. (1980). An empirical demonstration of risk-sensitive foraging preferences. *Animal Behavior*, **28**, 820-830.
- COMMONS, M. L., WOODFORD, M., & DUCHENY, J. R. (1982). How reinforcers are aggregated in reinforcement-density discrimination and preference experiments. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 25-78). Cambridge, MA: Ballinger.
- KAHNEMAN, D., & TVERSKY, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, **47**, 263-291.
- LEE, W. (1971). *Decision theory and human behavior*. New York: Wiley.
- LEVENTHAL, A. M., MORRELL, R. F., MORGAN, E. F., & PERKINS, C. C. (1959). The relation between mean reward and mean reinforcement. *Journal of Experimental Psychology*, **57**, 284-287.
- LOGAN, F. A. (1965). Decision making by rats: Uncertain outcome choices. *Journal of Comparative & Physiological Psychology*, **59**, 246-251.
- MAZUR, J. E. (1984). Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes*, **10**, 426-436.
- MAZUR, J. E. (1985). Probability and delay of reinforcement as factors in discrete-trial choice. *Journal of the Experimental Analysis of Behavior*, **43**, 341-351.
- MAZUR, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55-73). Hillsdale, NJ: Erlbaum.
- REAL, L., OTT, J., & SILVERFINE, E. (1982). On the tradeoff between the mean and the variance in foraging: Effect of spatial distribution and color preference. *Ecology*, **63**, 1617-1623.
- WADDINGTON, K. D., ALLEN, T., & HEINRICH, B. (1981). Floral preferences of bumblebees (*Bombus edwardsii*) in relation to intermittent versus continuous rewards. *Animal Behavior*, **29**, 779-784.

(Manuscript received June 30, 1987;
revision accepted for publication December 21, 1987.)