

1992

Risk-sensitive foraging in the blue jay (*Cyanocitta cristata*).

Kevin C. Clements
University of Massachusetts Amherst

Follow this and additional works at: <https://scholarworks.umass.edu/theses>

Clements, Kevin C., "Risk-sensitive foraging in the blue jay (*Cyanocitta cristata*).\" (1992). *Masters Theses 1911 - February 2014*. 2209.
<https://doi.org/10.7275/7675905>

This thesis is brought to you for free and open access by ScholarWorks@UMass Amherst. It has been accepted for inclusion in Masters Theses 1911 - February 2014 by an authorized administrator of ScholarWorks@UMass Amherst. For more information, please contact scholarworks@library.umass.edu.



312066013694184

RISK-SENSITIVE FORAGING IN THE BLUE JAY (*CYANOCITTA CRYSTATA*)

A Thesis Presented

by

KEVIN C. CLEMENTS

Submitted to the Graduate School of the
University of Massachusetts in partial fulfillment
of the requirements for the degree of

MASTER OF SCIENCE

February 1992

Department of Psychology

© Copyright by Kevin Charles Clements 1992

All Rights Reserved

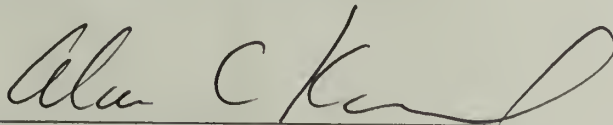
RISK-SENSITIVE FORAGING IN THE BLUE JAY (*CYANOCITTA CRYSTATA*)

A Thesis Presented

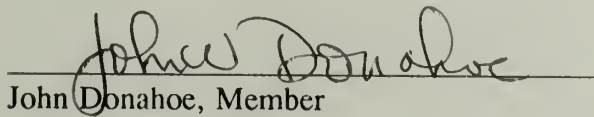
by

KEVIN C. CLEMENTS

Approved as to style and content by:



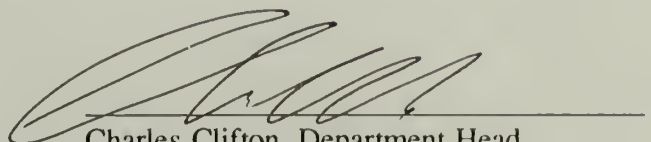
Alan C. Kamil, Chair



John Donahoe, Member



Paul W. Ewald, Member



Charles Clifton, Department Head
Department of Psychology

For my Grandmother, Norene Hesterley. May her kind and loving spirit live forever.

ACKNOWLEDGEMENTS

Thanks are due a number of people who helped me to complete this study. Betsy Frederick, Deb Olson, and Paul Rivest helped collect data and care for the animals. My committee members, Paul Ewald and John Donahoe, provided insightful comments during the writing of the thesis. Finally, thanks are most especially due my advisor, Al Kamil, without whom I would never have been in a position to even begin a study such as the one reported here. The research reported here was supported in part by a research grant from the National Science Foundation (BNS-8418721). The blue jays were captured and maintained under federal scientific collecting permit PRT 2-737 and Massachusetts permit 745.

ABSTRACT

RISK-SENSITIVE FORAGING IN THE BLUE JAY (*CYANOCITTA CRYSTATA*)

FEBRUARY 1992

KEVIN CHARLES CLEMENTS, B.S., JAMES MADISON UNIVERSITY

M.S., UNIVERSITY OF MASSACHUSETTS

Directed by: Professor Alan C. Kamil

The study of risk-sensitive foraging asks how foragers respond to variation in the distribution of food within patches. Models have been proposed that explain risk-sensitive behavior in terms of the relation between energy need and energy intake. Tests of these models have used very small, specialized foragers with high energetic demands. During the present study, five blue jays (*Cyanocitta cristata*), medium-sized, omnivorous foragers, were tested under an operant simulation of foraging. The jays were allowed to choose from two patches over several trials per session (choices early in the session were forced while those later in the session were free). One patch provided a constant number of rewards per trial while the other provided a variable number. Both provided the same mean value during a session. Six reward conditions were used. Some unique procedural precautions were taken to minimize the biasing effects position preferences can exert on tests of patch choice. The jays developed a strong and consistent preference for the constant patch (risk aversion). This preference was not affected by reward conditions. Choice latencies during forced-choice trials indicated that patch preferences shown during free-choice developed at least partially during the preceding forced-choice phase, an assumption heretofore untested. The jays' strong and consistent risk aversion may be due to their relatively large size or to the fact that they are omnivorous foragers.

TABLE OF CONTENTS

	<u>Page</u>
ACKNOWLEDGEMENTS	v
ABSTRACT	vi
LIST OF TABLES	viii
LIST OF FIGURES	ix
Chapter	
1. INTRODUCTION	1
Risk sensitivity and traditional foraging models	2
History of risk-sensitive foraging studies	3
Models of risk sensitivity	9
The present study	13
Predictions of the models	14
2. MATERIALS AND METHODS	18
Subjects	18
Apparatus	18
Procedure	19
Color/Patch assignment	22
Data Analysis	22
3. RESULTS	26
Overall patch choice	26
Patch preference on the nonpreferred side	27
Effect of reward conditions	27
Forced phase choice latencies	28
4. DISCUSSION	39
APPENDIX: DERIVATION OF z WITH RESPECT TO σ	56
BIBLIOGRAPHY	57

LIST OF TABLES

Table	Page
1. Patch choice, both overall and on the nonpreferred side, averaged over all 30 sessions for each subject.	29

LIST OF FIGURES

Figure	Page
1. Utility functions for different relationships between energy requirements and available resources.	15
2. Probability distribution of net energy gained after one day of foraging.	16
3. Each figure shows two superimposed, cumulative (normal) probability distributions.	17
4. A schematic diagram of the operant conditioning chamber and its fittings.	24
5. A flow diagram of typical within-trial events.	25
6. Development of risk-aversion over the entire experiment as shown by the proportion of choices in the constant patch on both sides of the chamber.	30
7. Development of side preferences over the entire experiment as shown by the proportion of choices made on one side of the chamber.	31
8. Scatter plot for subject 65 showing the relationship between the strength of the side preference and the strength of the overall patch preference.	32
9. Development of risk-aversion over the entire experiment as shown by the proportion of constant patch choices on the nonpreferred side of the chamber.	33
10. Effect of reward condition on patch preference in the last (fifth) replication.	34
11. Effect of reward condition on mean patch preference averaged over replications 2 through 5.	35
12. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to reward condition, trial block (1-8 and 9-16), and patch type (■ = constant, □ = variable).	36
13. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to reward condition and patch type (■ = constant, □ = variable).	37

14. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to trial block (1-8 and 9-16) and patch type (■ = constant, □ = variable).	38
15. Pure color preferences, averaged over subjects, from each of four color preference tests.	55

CHAPTER 1

INTRODUCTION

Some animals find their food distributed continuously throughout their habitat. A grazing wildebeest may have no difficulty finding food on the vast plains of Africa during the wet season. But for many other animals, the task of finding food is a bit more complicated. These animals may find their food distributed in discrete areas. For example, a titmouse might forage for moth larvae in pine cones. Or a kangaroo mouse might find seeds in the sand surrounding the desert plant that produced them. Biologists typically refer to clumped resources, especially food resources, as patches. In the example of the titmouse above, the patch in question might be a pine cone, or a tree containing pine cones, depending on the specificity of one's definition. There is an entire body of theory which addresses the ways in which animals exploit patchily distributed food resources. This body of theory assumes that animals have been prepared through natural selection to respond optimally when searching for food. As such, we refer to this body of theory as optimal foraging theory. Optimal foraging models of patch use assume that, when foraging in patchy environments, animals should optimally allocate their time between foraging within and travelling between patches (see Stephens and Krebs 1986 for a recent and thorough review of foraging theory). Given this assumption, the optimal forager is faced with at least two important decisions concerning patch use.

First, while between patches, the forager must decide which patch is best to exploit at a given point in time. Second, once foraging within a patch, the forager must determine when it is best to leave and hunt elsewhere. The first decision was the subject of this thesis.

Risk sensitivity and traditional foraging models

Foraging models have traditionally assumed that, when making the above decisions, the quantity to be maximized is the rate of energy intake per unit time of foraging. If rate of intake is the only quantity to consider in deciding where to hunt, then a successful forager need only concern itself with choosing patches which will maximize the expected (mean) value of the food distribution. Such a strategy implies that variability within food distributions may be safely ignored by the forager. But is a pure mean-maximizing strategy always the best option available to the animal which may actually perish if bad patches are chosen? Consider the following example:

An animal has enough time remaining in its foraging period to visit only one more patch. The animal requires at least two more units of energy to survive the night. There is a choice between only two patches. The first patch is known by the animal to provide zero or 10 units of energy per visit, each with a probability of 0.5 (giving an expected value of five and variance of 25). The second patch provides two or four units of energy per visit, each with a probability of 0.5 (giving an expected value of three and a variance of one). It is clear that a strict expected value maximizer would choose the first patch and have only a 50% chance of surviving. But an animal that is sensitive to both the mean *and* variance of the two distributions, and is willing to trade a higher expected value with higher variance for a lower expected value with lower variance has a 100% chance of surviving.

The above example demonstrates that animals should be sensitive to both the mean *and* variance of a reward distribution. A number of researchers have recently shown that some animals possess such sensitivity and use it to determine which patches are best to exploit. This phenomenon has been called "risk sensitivity" because of the riskiness

involved in making choices based on variability. The study of risk asks how foragers can choose between probability distributions to optimize patch choice and thereby gain the highest chance of survival possible. The most widely recognized of these studies are those done by Caraco and his colleagues (Caraco 1981, 1982, 1983, and Caraco, Martindale & Whittam 1980).

History of risk-sensitive foraging studies

In the first rigorous demonstration of this phenomenon, Caraco et al. (1980) offered food deprived yellow-eyed juncos (*Junco phaeotus*) a choice between two feeding stations at opposite ends of a small aviary. One station provided a constant number of seeds per visit while the other station provided a variable number of seeds per visit, based on some probability distribution. The variable station had the same expected value as the constant station. So, for example, the constant station might provide three seeds per visit while the variable station would provide zero or six seeds per visit, each with a probability of 0.5. Expected values for the stations were changed daily by varying the average number of seeds available. After an initial period of forced choice to acquaint them with the daily conditions, the juncos were allowed to freely choose one of the two feeding stations over a number of trials. This two-phase sequence of forced choice followed by free choice between two patch types is the general approach taken by most studies of risk-sensitive foraging.

Caraco et al. (1980) found that feeding preferences depended on the relationship between present energy reserves, expected energy expenditure, and expected energy intake (determined by severity of food deprivation and reward conditions). If expected energy intake exceeded expected expenditure the juncos preferred the constant station (risk

aversion). If expected intake was less than expected expenditure the juncos preferred the variable station (risk proneness). So, "taking a risk" means choosing variability in the hope of gaining *more* than the expected number of rewards on any given patch visit. Use of this risky strategy is presumably motivated only when the expected values of the available patches will not meet energetic requirements.

Caraco et al. (1980) predicted the above results from their expected energy budget hypothesis which is based on the concepts of utility theory. If environmental conditions are such that available energy resources can be expected to meet or exceed an animal's energetic requirements the animal is said to have a positive expected energy budget. If the expected energy budget is positive the utility function is a negatively accelerated, concave curve (Figure 1A). With a concave utility function each additional reward adds less utility than the last. Because the expected value is enough to meet requirements, additional energy is treated as an unnecessary bonus. Since it need not seek more than the expected value, the animal with a positive energy budget should prefer the certain expectation over the expected value based on a probability distribution. This is called risk-averse (or risk-avoiding) behavior; the animal avoids risky patch choices, those with high variability, because they are not necessary for survival. The concave utility curve becomes more 'bowed' as the (positive) difference between expected intake and requirements becomes larger (larger positive energy budget, Figure 1B). As this difference increases, any specific food value will loose utility, which should make the animal become even more risk averse.

If available energy resources cannot be expected to meet an animal's energetic requirements the animal has a negative expected energy budget. If the expected energy budget is negative the utility function is a positively accelerated, convex curve (Figure

1C). When the utility function is convex each reward adds more utility (is valued more) than the previous one. Because the expected value is *not* enough to meet requirements, additional units of reward (beyond the mean) are required to survive. Since it needs to seek more than the expected value, the animal with a negative energy budget should prefer the variable food distribution over the certain expectation. This is called risk-prone (or risk-taking) behavior; the animal is given to risky patch choices because these provide the only chance for survival. The convex utility function becomes more 'bowed' as the (negative) difference between expected intake and requirements becomes more negative (larger negative energy budget, Figure 1D). As this difference becomes more negative, additional units of reward gain more and more utility which should make the forager become even more risk prone.

Caraco has also demonstrated both risk-prone and risk-averse behavior in dark-eyed juncos, *Junco hyemalis* (1981), and white-crowned sparrows, *Zonotrichia leucophrys* (1982, 1983), which are somewhat larger than juncos. However, these studies were partially flawed by a methodological oversight. In order to manipulate energy budgets, biasing the birds toward a negative or positive expectation over time, Caraco and his colleagues controlled the severity of food deprivation between sessions as well as rates of energy intake (energy divided by foraging time) within sessions. The rates of intake were controlled by systematically varying the expected value of each patch along with the delay between trials (intertrial interval or ITI). Since a bird could only feed during a trial, the ITI in Caraco's procedure represented the amount of time a bird had to wait before the next opportunity to feed. The ITI was always at least 30 seconds and there were 30 seconds added for every seed eaten during the preceding trial.

This variation in ITI presents both theoretical and analytical difficulties. Two assumptions which most theories of risk share is that opportunities to find food are randomly and independently distributed. The assumption of independence is violated when the events of one trial affect the the events within, or the onset of, the next. Caraco's variation of ITI constituted such a violation; the onset of a trial was delayed, via the ITI, by 30 seconds for every seed eaten during the previous trial. As a result, constant patches produced constant delays before the next opportunity to feed while variable patches produced variable delays before the next such opportunity. (It is important to note that the term 'delay', as used in the operant psychology literature, typically means the time elapsed between a criterion response and its consequence. I use the term more generally here to mean any span of time imposed by the experimenter which interrupts the ongoing behavior of the animal.)

Even if varying ITI in studies of risk does not violate assumptions of the model being tested, varying ITI on a one-to-one basis with reward conditions makes accurate interpretation of data analyses very difficult. Since the ITI in Caraco's studies was perfectly correlated with reward distributions between conditions, it is impossible to separate the two as independent variables in the data analyses. Were the effects due to the variation in the ITI or to the variation in the number of seeds available at the feeding stations? Or, since both play a role in the rate of energy intake, was there an interaction between the two variables, one enhancing or hindering the effect of the other?

A rate is simply an amount divided by a span of time. To change a rate one could alter one of the two quantities while holding the other constant. The quantities could also be changed together, either in opposite directions or in the same direction by disproportionate amounts, but this method is obviously less procedurally sound. To vary

the rate of energy intake in a discrete trials procedure, one could either change the amount of reward per trial while keeping trial duration constant, or change trial duration while keeping reward constant. The duration of a trial could be altered by varying the length of any delay the animal encounters within trials. Such delays should ideally occur while the animal is working for, or handling, its rewards. In studies of foraging, these delays might include the amount of time needed to travel to a chosen patch, to search a patch for prey, or to capture and handle prey. Delays between trials (ITI) could also be changed as long as the change is not correlated, and therefore confounded, with events within trials.

Reports in the operant psychology literature, dating more than three decades ago, indicate that both the magnitude of reward and the delay between response and reward are important variables in experiments that are strikingly similar to today's studies of risk. Leventhal, Morrell, Morgan, and Perkins (1958) allowed rats to choose between the two outer arms of an E-shaped maze. One arm provided a constant amount of food while the other provided one of two equally probable amounts of food. Both arms yielded the same mean amount. With the exception of variability introduced by the rats in navigating the maze, the duration of each trial was constant. The rats preferred the variable arm (risk proneness) but this preference declined as the mean amount of food available in the arms was increased.

Pubols (1962) gave rats a choice between the arms of a Y- maze. But rather than varying the amount of reward, one arm imposed a fixed delay between response and reward while the other imposed one of two equally probable delays. Both arms provided the same mean delay and the amount of reward available in each arm was the same. The

rats preferred variable delays (risk proneness), but as the average delay was decreased the strength of their preference decreased.

Taken together, the Pubols (1962) and Leventhal et al. (1958) experiments showed that, no matter which defining quantity of the rate of energy intake was varied, as the rate increased the strength of risk proneness declined toward indifference. This is exactly what Caraco's energy budget rule predicts; as environmental resources improve on average, with energy requirements held constant, the possible benefits associated with resource variability decrease so that risk proneness becomes less useful.

Though they produced interesting effects, neither the Leventhal et al. nor the Pubols studies manipulated the animal's actual energy budgets. Manipulating rate of energy intake is a necessary part of manipulating energy budgets but is not sufficient in itself to bias the animal toward a positive or negative expectation over time. One must, in addition, determine the energy intake rate required to meet the animal's needs for a specified period of time. Then, the rate of intake should be experimentally increased or decreased relative to the requirement without compromising the assumptions of the model or the procedural integrity of the experimental design.

Recently, Barnard and Brown (1985) have demonstrated risk sensitivity in another mammalian predator, the common shrew (*Sorex araneus*). As is typical, the shrews were offered a choice between two feeding stations which provided a constant or a variable number of mealworm segments, each station with the same expected value. By manipulating daily energy budgets Barnard and Brown (1985) found that the shrews, as predicted, were risk averse under positive expected energy budgets and risk prone under negative expected energy budgets. This is the only reported study, other than those of Caraco and his colleagues, that has directly measured and manipulated energy budgets.

However, only one reward condition was used. Also, energy intake rates were decreased by "making stations temporarily unavailable between visits" (p. 162). In other words, Barnard and Brown (1985) also varied the ITI. Whether this variation was dependent on within-trial events and in violation of the assumption of trial independence cannot be ascertained since the method by which stations were made "temporarily unavailable" was not specified. But it seems reasonable to assume that the method used was similar to that used by Caraco given the similar purposes of the manipulation.

Models of risk sensitivity

Results such as those above have been described in terms of minimizing the probability of an energetic shortfall. Stephens (1981, see also Stephens and Krebs 1986) and Stephens and Charnov (1982) have formulated the z-score model of risk-sensitive foraging. The model assumes that an animal obtains its energy throughout the day in small units of food which are randomly and independently distributed. The sum of the pay-offs from these units of food should, according to the central limit theorem, be normally distributed. Thus, the animal's energy supply at the end of the day (S_0) should be normally distributed. The forager should be able to exercise some behavioral control over the mean (μ) and variance (σ^2) of this distribution by deciding where to feed. The model assumes that the animal has a fixed daily energy requirement (R) which must be met to survive the night. The model seeks to maximize:

$$P(\text{surviving the night}) = P(S_0 \geq R) \quad (1)$$

This is equivalent to minimizing the likelihood of an energetic shortfall. Since S_0 is normally distributed, the probability of survival can be determined by first converting R to a standard normal deviate (a z-score).

$$Z_R = \frac{R - \mu}{\sigma} \quad (2)$$

For any requirement (R) the probability of starvation is $\Phi(Z_R)$ where Φ represents the cumulative probability of the normal distribution. Thus, the probability of survival becomes:

$$P(S_0 > R) = 1 - \Phi(Z_R) \quad (3)$$

Since it is cumulative Φ increases with Z_R (see Figure 2). Thus, the probability of survival decreases as Z increases. Following this logic, the optimal forager should minimize the value of Z_R .

Consider once more the animal faced with a choice between two patches with equal means but different variability. Under what conditions should high or low variance be preferred when μ is fixed? Following the development of the model by Stephens and Krebs (1986), the effect of σ on Z can be described by finding the first derivative of Z with respect to σ :

$$\frac{dz}{d\sigma} = \frac{-(R - \mu)}{\sigma^2} = \frac{\mu - R}{\sigma^2} \quad (4)$$

(See Appendix 1 for derivation). Since the sign of equation (4) is determined by $\mu - R$:

$$\frac{dz}{d\sigma} > 0 \text{ if } \mu > R$$

$$\frac{dz}{d\sigma} < 0 \text{ if } \mu < R$$

Thus, a z-minimizer's preference for variance depends only on the relative sizes of μ and R (Figure 3). When $R - \mu$ is negative Z decreases as σ decreases. Stated another way, when the expected value (μ) exceeds the requirement (R) the probability of starvation ($\Phi(Z_R)$) is minimized by choosing the patch with the lowest variance (Figure 3A). Conversely, when $R - \mu$ is positive Z decreases as σ increases. In other words, when the expected value is less than the requirement the probability of starvation is minimized by choosing the patch with the highest variance (Figure 3B).

In general, the z-score model predicts indifference when $\mu = R$, risk aversion when $\mu < R$, and risk proneness when $\mu > R$. These generalities simply restate the predictions of the expected energy budget rule. However, the z-score model is more precise in its predictions by virtue of its higher level of quantification. When μ is fixed between patches and $\mu > R$ the z-score model predicts that the optimal choice is that with the *lowest* variation in reward. Conversely, when $\mu < R$ the optimal choice is that with the *highest* variation in reward.

There have been other studies of risk which did not manipulate energy budgets. Many of these studies examined the response of nectarivorous foragers to risky situations. Wunderle and O'Brien (1985) offered foraging bananaquits (*Coereba flaveola*), a small tropical bird, a choice between constant and variable nectar rewards in artificial flowers. The bananaquits preferred the constant rewards (were risk averse) over all reward

conditions. This was hardly surprising, however, since the birds always had *ad lib* access to food in nonexperimental conditions. Real (1981) and Real et al. (1982), and Waddington et al. (1981) also found consistent risk aversion when offering bumblebees and wasps choices between constant and variable nectar reward in artificial flowers.

Real (1980) has proposed a descriptive model of feeding preferences in foraging situations involving variability in the food distributions. The model seeks to maximize:

$$\mu - k\sigma^2 \quad (5)$$

where μ and σ^2 are the mean and variance of the reward distribution and k is a constant which corresponds to the undesirability of uncertainty (variation). According to the model, a forager should be willing to forfeit k units of mean reward for a unit reduction in variance. Thus, as variance becomes less desirable k becomes larger regardless of the mean's size. The result is a model which maximizes the expected value discounted by a certain proportion of the variance. Accordingly, Real calls this the "variance discounting" model of risk-sensitive foraging. The model predicts that foragers should always avoid uncertainty (be risk averse) even at the cost of lowering the value of the certain alternatives. In other words, the animal should always choose the lowest variability even if it means choosing patches with lower means than those provided by more variable patches.

Unlike the z-score model, the variance discounting model is an *a posteriori*, descriptive model. The constant, k , is a "fitting-variable" that fits the equation to an animal's performance after the fact. One must first observe and quantify an animal's performance in an uncertain situation to know what value of k the animal used. Thus, the variance discounting model finds its only utility as a descriptive model; it has no real

all situations. This lack of predictive utility in the variance discounting model is especially evident for species which have yet to be tested in the context of risk. The z-score model, however, allows predictions of responses to variability before the animal forages if the animal's energy requirements and the available resources are known. These predictions are possible even if the species has never before been tested in risk-sensitive foraging experiments.

The present study

The experiment reported here sought to determine some of the responses to risk in the blue jay (*Cyanocitta cristata*), a species not yet tested in the context of risk. Most studies of risk use very small, often highly specialized, foragers with high daily energy requirements that can easily lead to starvation if not met. The blue jay, however, is a larger, omnivorous species with less stringent energy requirements.

The jays were trained to choose between two patches during an operant simulation of a foraging bout. One patch provided a constant food distribution and the other a variable distribution, but both always had the same expected value. Energy budgets were not manipulated but the daily energy intake was held consistently below that which is normal for a free-feeding blue jay. The expected values and variability of reward distributions in the two patches were changed daily.

Many of the procedures used in this study were typical of risk-sensitive foraging experiments. However, there were two important methodological differences between this study and previous studies of risk which used discrete trials. The first difference involved the ways in which patches were represented experimentally, and the consequences of these representations on the analysis of patch choice data. The second methodological

difference between this and previous studies of risk involved the types of data collected. The single measure taken in all studies of risk is the number of times each patch type is chosen while all patch types are freely available. This measure was used in the present study along with two additional measures never before used: patch choice on the less frequently chosen side of the chamber and patch choice latencies. The latency measure is theoretically applicable to behavior during the forced phase of foraging sessions—behavior not tested in previous studies. The procedural and analytical differences between this and previous studies of risk will be discussed in more detail later.

Predictions of the models

The variance discounting model always predicts risk aversion over all reward conditions. But the z-score model predicts decreasing risk aversion as the mean increases while variance remains the same. Five of the six reward conditions in the study allowed for comparison between changing means when the variance was held constant. Specific predictions from the energy budget model could not be made in the absence of more specific information concerning the energetic state of the animals.




Figure 1. Utility functions for different relationships between energy requirements and available resources. Cumulative food intake is on the abscissa and utility is shown on the ordinate. (A) Utility function is concave when expected resources meet or exceed energy requirements; indicative of positive expected energy budgets and risk-aversion. (B) The more expectation exceeds requirements, the more bowed the utility function, giving a value of 10 units of food less utility in B than in A. The animal should be more risk-averse in B than in A. (C) Utility function is convex when requirements exceed expected resources; indicative of negative expected energy budgets and risk-proneness. (D) The more requirements exceed expectation, the more bowed the utility function, giving a value of 10 units of food more utility in D than in C. The animal should be more risk-prone in D than in C.




Figure 2. Probability distribution of net energy gained after one day of foraging. Net energy is shown on the abscissa and M represents the expected (mean) value. The probability of survival is one minus the cumulative probability of R or, $1 - \Phi(Z_R)$; this is the upper tail of the curve. When the daily energy requirement, R , is equal to R_1 , the probability of survival equals the total area under the curve from point A to point C . When R equals R_2 , the probability of survival equals the area under the curve from point B to point C . Thus, the probability of survival decreases as the energy requirement increases.




Figure 3. Each figure shows two superimposed, cumulative (normal) probability distributions. Net energy intake is on the abscissa where M represents the expected (mean) gain. The probability of starvation is on the ordinate. Both distributions have the same expected gain but the distribution labeled V1 has a lower variance than the distribution labeled V2. (A) When the energy requirement, R, is less than M, the probability of starvation, $P(V)$, is lowest for the distribution with the lowest variance, V1. (B) When R exceeds M, the probability of starvation, $P(V)$, is lowest for the distribution with the highest variance, V2.

CHAPTER 2

MATERIALS AND METHODS

Subjects

Five experimentally naive blue jays (*Cyanocitta cristata*) of unknown sex and approximately one year old served as subjects. The jays were taken from their nests near Amherst, MA at 10-14 days of age and hand-raised in the laboratory. Throughout the study, the jays were weighed daily and maintained at approximately 80% of their free-feeding weights by controlled daily feedings of turkey starter and mynah pellets. The blue jay colony was kept at 22-27 degrees Celsius with a constant 14/10 hr light-dark cycle.

Apparatus

Sessions were conducted in an operant conditioning chamber, 122 cm long, 41 cm high and 39 cm wide at the front tapering to 25 cm wide at the back. (See Figure 4 for a schematic diagram of the chamber.) Four round pecking keys, 2.5 cm in diameter, were arranged horizontally across the front panel, 23 cm from the chamber floor, with the center of the two outer keys 5 cm from the side walls. The two inner keys were not used in this study. Another round key was located in the center of the rear stimulus panel 12 cm from the chamber floor. A small IEE stimulus projector was mounted behind each key. There were three perches, 1.2 cm in diameter, mounted parallel to the front panel 12 cm below the pecking keys. Two of these perches, each 7.5 cm long, were located below the two outer keys and the third, 12.5 cm long extended below both inner keys. A fourth perch, 12.5 cm long, was located below the key on the back panel. In their

resting positions (unoccupied), all perches were 8.5 cm from the panel they served and 11 cm from the floor of the chamber. When occupied by a bird, each perch dropped about 0.8 cm (closing a microswitch) so that a perching subject's eye was approximately level with the center of a key. The keys could not be reached from the floor and a bird had to be on the perch in front of a key for any pecks at that key to be recorded.

A 5-cm wide food cup was located in the center of the front panel, 15 cm from the floor of the chamber. Reinforcers, half pieces of small *Tenebrio molitor* larvae averaging 1.0 cm in length, were delivered to the food cup by a Davis UF-100 universal feeder mounted outside the chamber's front panel. The cup was illuminated by a 24-V bulb during reinforcement periods. White noise played through a small speaker mounted on the outside of the right wall, 18 cm from the chamber floor and 36 cm from the front panel, provided masking noise throughout each session. A 24-V white houselight, positioned on the right wall 15 cm above the speaker, remained lit throughout each session. All stimulus events and contingencies, as well as all data collection, were controlled by an NCR PC4 personal computer.

Procedure

Each daily session consisted of 16 forced trials followed by 24 free-choice trials. (See Figure 5 for a flow chart of within-trial events.) Every trial began at the back of the chamber when the back key was illuminated with a white light. A single peck at the white key caused it to become dark and inoperative while simultaneously causing the illumination of one or both of the front keys. The remainder of the trial took place at the front of the chamber. If the trial was forced then only one of the front keys was lit but

if the trial was free-choice both keys were lit. A front key could be either red or green on any given trial.

During a forced trial a single front key was lit with one of four possible color-side combinations: red-left, red-right, green-left or green-right. A single peck at the illuminated key initiated a 5-s delay, simulating a handling time requirement. Any pecks at the key during this delay had no effect. The first peck after completion of the handling time delay simultaneously caused the darkening of the key and the onset of a 12-s reward period. The four color-side combinations were presented equally often and in random order over every block of eight forced trials. No more than three successive forced trials could occur with the same color-side combination.

During a free-choice trial both front keys were lit, one red and the other green. To continue the trial once the front keys became lit, the bird chose a patch (color-side combination) by landing on one of the two outer perches on the front panel. Landing on one of these perches caused the illuminated key above the unchosen perch to become dark and inoperative while the key above the chosen perch remained lit. A single peck at the lit key initiated a 5-s delay, simulating a handling time requirement. Any pecks at the key during this delay had no effect. The first peck after completion of the handling time delay simultaneously caused the darkening of the key and the onset of a 12-s reward period. The two stimulus combinations possible on a free-choice trial (red/left-green/right, green/left-red/right) were presented equally often and in random order over every block of eight free-choice trials. No more than four of the same combination could occur over successive trials.

A 10-s ITI, during which all keys were dark and inoperative, followed the reward period of every trial. Thus, the timing of events under the control of the experimenter

remained constant on every trial. Variability in the duration of a session could only be introduced by a bird not responding at certain points within trials (e.g. withholding the peck that starts the trial).

The two possible colors for the front keys represented two patches in which the jays could choose to forage. These patches differed only in the number of rewards each provided when chosen. One patch, designated the "constant" patch, yielded the same number of rewards each time it was selected. The other patch, designated the "variable" patch, provided one of two possible reward values when selected, each with a probability of 0.5. The mean of the two reward values available in the variable patch always equaled the reward value of the constant patch. Thus, the *expected* values were always the same for both patches.

Six reward conditions were used. These conditions were designated as: 0-2(1), 1-3(2), 0-4(2), 2-4(3), 1-5(3), and 0-6(3). For each condition, the hyphenated values represent the two equally probable reward values used for the variable patch while the single value (in parentheses) represents the number of rewards that were available in the constant patch. The six reward conditions were presented in random order over every block of six sessions. No single reward condition was presented twice in succession. (A block of six sessions will henceforth be referred to as a replication.) Each of the two possible variable patch reward values within a session were used equally often and in an unpredictable order over every block of eight variable patch choices. A single variable patch reward value could occur no more than four times over successive trials.

The reward period always lasted for 12 seconds, during which the food cup was illuminated. If rewards were to be delivered during a given trial, the feeder would begin operating at the onset of the reward period. A piece of a mealworm was delivered every

1.5 seconds until the number of rewards designated for a trial had been reached. There were always at least three seconds of illumination following the last reward to ensure that all rewards delivered during a trial could be eaten within the trial.

Color/Patch assignment

On the day prior to their first foraging session, each jay was tested for a preexisting color preference. The color preference test was conducted with the same discrete trials procedure described above except that each patch always provided three rewards per visit. There were no reliable color preferences in evidence. Assignments of patch colors were made as randomly as possible while trying to counterbalance across subjects. Three jays were assigned red as the constant patch and green as the variable patch while the remaining two jays were assigned the reverse. These patch-color assignments remained in force throughout the study.

Data Analysis

Three measures of patch preference were analyzed. The primary measure was overall patch choice, as indicated by the proportion of choices made for the constant patch, during free-choice trials. The second measure was the patch preference on the nonpreferred side of the chamber. Both patch types occurred equally often on both sides of the chamber. Since free-choice trials allowed the birds to choose one patch as well as one side of the chamber, there were often fewer choices on one side. The side chosen less often in a session was designated the nonpreferred side. The proportion of constant patch choices on the nonpreferred side served as an index of patch preference on that side. This measure was useful in determining patch preference for sessions in which a

subject showed a strong positioni preference. The third measure of patch preference was patch choice latencies for all trials, forced as well as free-choice. Choice latency was defined as the time elapsed between leaving the back perch and landing on one of the two outer front perches.




Figure 4. A schematic diagram of the operant conditioning chamber and its fittings. The perspective is from above the chamber looking down.




Figure 5. A flow diagram of typical within-trial events. The trial in this figure is a free-choice trial. The jay starts the trial by pecking the white key on the back panel (box 1). In this example the jay chooses the red patch on the left side of the front panel by landing on the perch under the red key (box 2). The final peck of the trial (box 3) turns the chosen key off (box 4) and initiates the reward period. The reward period is followed by an ITI, after which the next trial begins.

CHAPTER 3

RESULTS

Overall patch choice

There were 30 foraging sessions, five for every reward condition, for each bird. According to the primary measure of patch preference, the overall proportion of choices for the constant patch, all of the jays began the experiment close to indifference (see Figure 6). The mean overall patch preference of each subject, averaged over all 30 sessions, was tested against indifference using single-sample t-tests. The initial indifference quickly changed into a strong and consistent preference for the constant patch for Jays 44, 59, 87, and 104 (see Figure 6, and see Table 1 for statistics). The preference for the constant patch was often exclusive by the end of the experiment for these four jays, indicting a high level of risk aversion. However, the primary measure of patch preference yielded unclear results for the remaining subject, 65, especially toward the end of the experiment (Figure 6C).

Subject 65 developed a strong and reliable preference for the right side of the chamber (see Figure 7C; $t(29) = 8.28$, $p < 0.005$). Because both red and green, and therefore both patch types, appeared equally often on both sides of the chamber, a significant side preference would necessarily bias patch preference toward indifference. For subject 65, there was a significant negative relationship between the strength of the side preference and the strength of the constant patch preference as assessed by the primary measure (Pearson's product moment correlation: $r(28) = -0.3117$, $p < 0.0005$, 1-tailed). As the strength of the side preference increased from 0.5 to 1.0, the primary measure of patch preference decreased from 1.0 to 0.5 (see Figure 8).

Patch preference on the nonpreferred side

Side preferences, even with subject 65, were seldom exclusive; a few responses usually occurred on the nonpreferred side. It was therefore possible to calculate the proportion of times each patch type was chosen on the nonpreferred side. The mean patch preference on the nonpreferred side, averaged over all sessions, was tested against indifference for each subject using single-sample t-tests. Sessions 2 - 4 and session 6 for subject 87 were excluded from these analyses because the side preference was 1.0 during these sessions (see Figure 7D), invalidating the measure of patch preference on the nonpreferred side.

Subject 65 had a very strong preference for the constant patch on the nonpreferred side throughout the experiment. (see Figure 7C; $t(29) = 45.04$, $p < .005$). For this subject, risk aversion on the nonpreferred side was exclusive (1.0) for 23 of the 30 foraging sessions, continuously so for the last 17 sessions (Figure 9C).

Side preferences were far less pronounced in subjects 44, 59, 87, and 104, but were occasionally strong. By the end of the experiment, side preferences were negligible for these jays (see Figure 7). Nevertheless, the analysis of patch preference on the nonpreferred side confirmed the strong tendency toward risk aversion reported above for these four jays (Table 1, Figure 9).

Effect of reward conditions

The effect of reward condition on patch preference was analyzed with one-way, repeated measures analyses of variance (ANOVA). The overall patch preference and the preference on the nonpreferred side were analyzed for each of the last four replications. The first replication was excluded from these analyses because of the variability

introduced by the initial acquisition of a response strategy. The jays were consistently risk averse across conditions (see Figures 10 and 11). Reward condition had no effect on patch preference, for either measure, in any of the replications (all values of $p > 0.05$).

Forced phase choice latencies

Analysis of the choice latency data was restricted to the asymptotic performance obtained in the fifth replication. Since the number of data points in the free-choice phase was highly skewed toward the constant patch, only the latencies from the forced phase of the foraging sessions were analyzed. There were 16 forced trials per session. The jays were forced to each of the two patch types four times in the first eight trials and four times in the last eight trials. Therefore, choice latencies were averaged over two blocks of trials, 1-8 and 9-16, to represent early and late responses in the forced phase. Forced phase choice latencies were analyzed with a three-way, repeated measures ANOVA using reward condition, patch type, and trial block as factors.

Choice latencies were not affected by reward condition (see Figures 12 and 13; $F(5,20) = 2.09$, $p > 0.10$). However, the jays consistently took longer to go to the front of the chamber when forced to the variable patch ($F(1,4) = 10.36$, $p < 0.05$), especially during the last half of the forced phase. The latencies for forced variable patch visits were an average of 1.6 seconds longer than latencies for forced constant patch visits. This effect is illustrated in Figure 12 (right panel) and, more clearly, in Figure 13. There was also a significant trial block by patch type interaction: as the forced phase neared completion, the difference between variable and constant patch latencies increased (see Figures 12 and 14; $F(1,4) = 11.32$, $p < 0.05$).

Table 1. Patch choice, both overall and on the nonpreferred side, averaged over all 30 sessions for each subject. Each of the two measures is accompanied by the corresponding value of *t* for each subject. All *df* = 29 except where noted.

Subject	Overall		Nonpreferred Side	
	Constant	<i>t</i>	Constant	<i>t</i>
44	0.88	18.51 *	0.95	25.88 *
59	0.89	16.73 *	0.93	20.15 *
65	0.78	12.38 *	0.98	45.04 *
87	0.73	9.14 *	0.81	33.39 * ⁺
104	0.81	9.44 *	0.88	9.31 *

* $p < 0.001$. ⁺ *df* = 25. The side preference for this subject was exclusive (1.0) for four sessions. These sessions were therefore excluded from the calculation of preference on the nonpreferred side.




Figure 6. Development of risk-aversion over the entire experiment as shown by the proportion of choices in the constant patch on both sides of the chamber. Preferences for all five subjects are shown individually in panels A - E (A=44, B=59, C=65, D=87, E=104). Panel F shows means across subjects 44, 59, 87, and 104.




Figure 7. Development of side preferences over the entire experiment as shown by the proportion of choices made on one side of the chamber. Preferences for all five subjects are shown individually in panels A - E (A=44, B=59, C=65, D=87, E=104).




Figure 8. Scatter plot for subject 65 showing the relationship between the strength of the side preference and the strength of the overall patch preference. The overall patch preference (on the ordinate) was measured by the proportion of constant patch choices on both sides of the chamber.




Figure 9. Development of risk-aversion over the entire experiment as shown by the proportion of constant patch choices on the nonpreferred side of the chamber. Preferences for all five subjects are shown individually in panels A - E (A=44, B=59, C=65, D=87, E=104). Panel F shows means across all subjects.




Figure 10. Effect of reward condition on patch preference in the last (fifth) replication. Two measures of patch preference are shown; the overall measure (■, see text) and patch preference on the nonpreferred side (□). It can be easily seen in this figure that, as variance increases with the mean held constant (conditions 2-4(3), 1-5(3), 0-6(3)), the level of risk-aversion does not change.




Figure 11. Effect of reward condition on mean patch preference averaged over replications 2 through 5. Two measures of patch preference are shown; the overall measure (■, see text) and patch preference on the nonpreferred side (□). It can be easily seen in this figure that, as variance increases with the mean held constant (conditions 2-4(3), 1-5(3), 0-6(3)), the level of risk-aversion does not change.




Figure 12. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to reward condition, trial block (1-8 and 9-16), and patch type (■ = constant, □ = variable). Reward condition is identical to foraging session, on the abscissa, since each condition occurs only once in every replication.




Figure 13. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to reward condition and patch type (■ = constant, □ = variable). The effect of patch type can be seen best in this figure.




Figure 14. Mean choice latencies, averaged over subjects, from the forced phase of the last replication arranged according to trial block (1-8 and 9-16) and patch type (■ = constant, □ = variable). The interaction between patch type and trial block can be seen best in this figure.

CHAPTER 4

DISCUSSION

The jays began this experiments with no definite preference for either the variable or the constant patch. But all of the jays quickly developed a robust preference for the constant patch. All three of the dependent measures used in this study indicated a strong level of risk aversion which did not vary across reward conditions. The overall patch choice during the free-choice phase showed a strong level of risk aversion which did not change over sessions once performance reached asymptote. The strength of risk aversion was unaffected by reward condition according to this measure; risk aversion seemed consistently strong across conditions. But, as demonstrated in the case of subject 65, when combined with a strong side preference this traditional measure of patch preference becomes useless.

This experiment used a discrete trials procedure which differed from the procedures used by other researchers who have studied risk in the lab. Patch assignments in this study were made based on color, where the colors appeared equally often on both sides of the chamber. All previous studies of risk which used a discrete trials procedure assigned patches based on the side of the apparatus. For example, the variable patch might be the left side of an aviary while the constant patch would be the right side. It has been argued (Caraco, personal communication) that assigning patch type by side is the better method because patches are spatially discrete places in the animal's natural environment. A titmouse, for example, finds pine cones in and around coniferous trees. But one could equally well argue that if food sources are rapidly depleted, or change spatially from day to day (e.g. mobile prey), some other attribute of the prey which

compose the patch, such as color, may be a good indicator of patch type or quality. For highly visual animals, as many birds are, a visual indicator of patch type may even add the advantage of being able to identify the quality of the patch while still some distance from it. A hummingbird, for example, may have learned through prior experience that a patch is likely to contain flowers of a highly variable quality just by seeing the color of the flowers.

The two methods of patch assignment, by side or by color, though seemingly not very dissimilar, have very different consequences with respect to analyzing patch choice and preference. These consequences are associated with a common problem in studies which allow animals a choice between two or more concurrently available alternatives; position preference. When patches are assigned by side, a side preference has a potentially devastating effect. Because side and patch cannot be separated, the correlation between the strength of the side preference and the strength of the patch preference must be 1.0, a perfect positive relationship, for any given session. In other words, the side preference cannot be separated from the patch preference in a single foraging session. However, the relationship between side and patch, and the relationship between side preference and patch preference are not necessarily the same thing.

To avoid the confounding of side and patch, and to separate side preference from patch preference, requires that side and patch be independent on some level. One way to accomplish this, given that patches are assigned by side, is to counterbalance the side/patch combinations over days. In effect, this technique makes side and patch independent between sessions but not within sessions. This method has been used successfully by some researchers (e.g. Caraco 1981, 1982, 1983), but still others have failed to take even this basic procedural precaution (e.g. Barnard & Brown 1985).

However, the success of this technique requires that no side preference exists. Since patch would still be inseparable from side within sessions, this precaution would have limited usefulness when confronted with a strong side preference. When confronted with a strong side preference, this type of counterbalancing would only serve to confirm the existence of the side preference rather than minimizing its effect. The confound would still be in effect but the direction of the effect over days would change.

Given the existence of a side preference, if the side/patch combinations were equally counterbalanced, so that each patch type occurred equally often on both sides over days, then the perfect positive relationship between side preference and patch preference within sessions would be reversed between sessions. As the strength of the side preference increased, the average strength of patch preference would decrease toward indifference over sessions. Thus, the positive relationship within sessions becomes a negative relationship between sessions when a side preference affects choice. This relationship would make a side preference obvious over sessions without specifically testing for one. In the absence of counterbalancing, the perfect positive relationship between side preference and patch preference within sessions would remain the same between sessions since side and patch would never be independent. One simply could not know if the animal was choosing the patch or the side without specifically testing for a side preference (which would undoubtedly develop over time if the animal's patch preference was consistently for one general patch type as was the case in this study).

Though counterbalancing is certainly preferable to not controlling the confounding of side and patch at all, it is only of limited usefulness. This technique is only useful if the animal has no significant side preference over days. This is especially true if patch preferences switch from the variable to the constant patch from day to day. It is also

possible that an animal may have a side preference on some days but not on others (as with the last 10 sessions of subject 65 in the present study; see Figure 7C). If counterbalancing were used in this situation, patch preferences might be real on some days but confounded with side preferences on other days.

All of the above problems can be avoided by assigning patch type based on colors that can change position rather than on position alone. When patches are assigned by color there is no mandatory relationship between side and patch; they are independent both within and between sessions. However, if a side preference exists, the correlation between the strength of the side preference and the strength of the overall patch preference could conceivably approach -1.0, a perfect negative relationship, within any given session. This is because patch type (color) is perfectly counterbalanced with side so that each patch occurs equally often on both sides. Thus, even a slight side preference would bias overall patch choice toward indifference. The stronger the side preference the stronger the bias toward indifference within and between sessions. But since patch and side are not the same thing in this procedure, the strengths of patch and side preferences can be analyzed separately. The proportion of constant patch choices on the less frequently chosen, or nonpreferred, side of the chamber is unaffected by a side preference as long as the side preference is not exclusive (100%), because, by definition, only one side is involved in the measure.

The advantages of assigning patch types by color are clear. The chief benefit is the independence of side and patch both within and between sessions. This allows for the analysis of pure patch preference on only one side of the chamber. Such analyses can be performed on both raw and averaged scores since patch and side are independent both within and between sessions. Moreover, since side and patch are always independent, this

measure is not adversely affected by position preferences. The nature of this procedure is such that side preferences become immediately apparent within, and therefore across, sessions without having to specifically test for them. If a side preference is evident the patch preference can be accurately assessed by analyzing the data from the nonpreferred side. For one subject in this study, number 65, the side preference was so strong that accurately assessing patch preference would have been impossible without this measure. Patch choice on the nonpreferred side showed a very strong constant patch preference for this bird. It would seem that the risk aversion was often strong enough to overshadow the side preference, thereby attracting responses to the nonpreferred side. Thus, measuring patch choice on the nonpreferred side in the present study may have been a more sensitive index of patch preference than the more traditional measure. Indeed, as shown in Figures 6, 9, 10 and 11, this measure consistently indicates a higher level of risk aversion than the measure which takes both sides of the chamber into account.

Both of the above measures of patch preference indicated a strong preference for the constant patch that was unaffected by reward condition in this study. However, these measures speak only of the behavior observed when the jays could choose freely between concurrently available patches. Because of the experimental design, measuring actual patch preference in the free-choice phase cannot address the possible development of patch preference during the initial stages of a foraging session.

In most studies of risk (as in this one), free choice between concurrently available patches is only allowed after an initial period in which the animal is forced to experience all patch types. This period of forced trials is designed to mimic the sampling period of a foraging bout in the real world (Krebs, Kacelnik & Taylor 1978). By sampling the available patches, the forager supposedly gains experience with which to decide what

patch or patches will best meet current energetic needs. In laboratory simulations of risk, the forced phase of a foraging session provides the animal with equal exposure to all patches. It is presumably in this phase that the animal develops the patch preference shown in the following free-choice phase. This assumption is implicit in all previous studies of risk whether they used a discrete trials procedure or not. Yet the only evidence in support of this assumption is indirect; when Caraco's birds face the same patch conditions under two different energetic states and show two different preferences, we assume that the preferences must have developed independently during the forced phases of the two sessions. This assumption has never been directly tested in the context of risk by gathering and analyzing data from the forced phase.

Given that a patch preference exists during the free-choice phase, there are, of course, three possibilities concerning preference during the forced phase of a foraging session. First, as normally assumed, the animal may use the forced phase as an opportunity to learn the daily reward distributions and develop the preference it will ultimately show in the next (free-choice) phase. This implies that the animal carries no preference between sessions. If this is the case then it is reasonable to expect at least a weak preference toward the end of the forced phase (but not the beginning). The second possibility is that the preference shown in the free-choice phase does not develop at all in the preceding forced phase, implying that the animal's preference is preserved between sessions. In this case one would expect the preference to be evident throughout the forced phase (assuming that the preference does not suddenly develop, as if by magic, in the single step which separates the forced and free-choice phases). The third possibility is a combination of the preceding two: the animal carries some patch preference between sessions which is strengthened in the forced phase of each session. In either of these

three possibilities, one could expect a patch preference to guide behavior at some point during the forced phase.

If a patch preference guides a jay's behavior during the forced phase, when only one patch is available per trial, the preference should become manifest in the jay's choice latencies. A jay starts a trial at the back panel, then leaves the back perch to go to the front panel where the trial will be completed. While moving toward the front panel the jay is confronted with only one available patch. If the available patch happens to be the preferred one the jay should go to it immediately. But if the patch is of the nonpreferred variety, the jay might hesitate long enough to make certain that the preferred patch is, in fact, unavailable. Once the jay is certain that the preferred patch will not become available, it should continue the trial. This logic applies whether the preference shown on forced trials develops during the forced phase or not.

Recording choice latencies in the forced phase of foraging sessions made it possible to test the assumption that the animal develops its preference anew from the start of every session. In a discrete trials procedure, trials represent a discrete measure of time. By definition, it takes time to develop a patch preference from the beginning of a session. Thus, if the assumption being tested is true, one would expect to see an effect of trials in the forced phase of a foraging session. Choice latencies should be similar for both patch types early in the phase, when no preference is evident, and should show a preference develops. The magnitude of this difference should increase as the preference develops, exerting a stronger influence on behavior (an interaction between patch type and trials).

Conversely, if the assumption is false, one would expect to see no effect of trials in the forced phase. This lack of effect could be achieved by two means. First, choice

latencies could be the same for both patch types throughout the forced phase, indicating that any patch preference seen in the free-choice phase was developed while the bird had free choice. This is not likely since the development of a preference in the free-choice phase would involve choosing both patches, on alternate trials, until the preference takes control of patch choice. Since there are only 24 free-choice trials in a session, the preference would have to develop extraordinarily fast for a statistically significant proportion of trials to be attributed to one patch.

There is a second, more likely way in which trials could have no effect on choice latencies. The latencies could differ between patch types by the same amount throughout the forced phase. This case would indicate that the bird retained its patch preference from the previous foraging session (or sessions). But no matter the reason, finding no effect of trials on choice latencies, and no interaction between patch type and trials, would refute the assumption that patch preferences develop anew from the start of every foraging session.

As it turned out, there was an effect of patch type on choice latency in the force phase. The latency to go to the variable patch was an average of 1.6 seconds longer than the latency to go to the constant patch when forced. This hesitation indicated a preference for the constant patch (or perhaps a functionally equivalent aversion for the variable patch). There was no effect of reward condition on these latencies. As with the other two measures of patch preference, the latency data show risk aversion which is consistent across reward conditions. There was no effect of trial block because as variable patch latencies increased, constant patch latencies decreased (though not to the same extent). Thus, the latencies averaged between patch types did not change appreciably over trials.

The most important finding with respect to the forced choice latencies was the interaction between trial block and patch type. The jays always tended to delay patch choice longer when forced to the variable patch, but the difference between variable and constant patch latencies increased from the first to the second half of the forced phase (Figure 14). This is exactly what would be expected if patch preference were developing, or being strengthened during the forced phase. As such, the choice latency measure seems to be the first measure actually recorded during forced trials which directly supports the assumption that patch preferences shown in studies of risk develop, at least partially, during the forced phase.

This effect may seem a bit surprising at first glance, given that the two colors which represented patch types remained fixed from day to day. If we were to assume a bias toward consistent risk aversion, it may have been a more efficient strategy to simply start each session choosing the color that always represented the constant patch rather than wasting time sampling. However, this strategy may be less available in the real world where one might assume that general patch types are not represented as consistently from day to day. In other words, patch type "cues" may be more variable in the real world. But even if general patch type cues are available in the real world (e.g. location), specific patch qualities (e.g. expected values, variance) may change over days because of local competition, seasonal changes, depletion-repletion cycles, prey mobility, or any number of other factors. In this experiment both the expected value and the variability associated with each patch type were changed unpredictably over days. Both of these variables would need to be accurately tracked by the successful risk-sensitive forager. This is especially true if there are more than two patches, as would presumably be the case in the real world. If the animal relies solely upon past experience with general

indicators of patch type, the indicators may lose their validity as true patch qualities change; making a strategy based on patch type cues alone less efficient. Conversely, if the animal relies solely upon sampling to determine patch type, it may spend more time than necessary gathering information if there is some cue available that can partially indicate general patch type. Thus, the best way to track changing patch qualities might be to rely partially on past experience with available patch type cues, and regularly update this experience through moderate sampling of the currently available patches.

Since the choice latencies often differed (to a lesser degree) between patches at the beginning of the forced phase, we cannot assume that the patch preference developed in its entirety from the start of every session. It is possible, even likely, that the jays carried some tendency toward risk aversion between sessions. This would mean that the jays based their preference not only on currently sampled information but on past experience with patch type cues as well. To test this hypothesis the jays were returned to the apparatus for a color preference test. This test was performed approximately 60 days after the last day of data collection. The method by which the color preference was determined was the same as that described in the patch assignment section in the description of Methods above. Briefly, the jays received the same sequence of forced and free-choice trials as usual, but each patch provided the same number of rewards (three) regardless of color. This procedure was conducted until each jay had completed four sessions (a range of 3 to 6 days).

All five jays showed a strong preference for the color which had previously been associated with the constant patch (see Figure 15). The strength of this color preference was even more impressive given the passage of time between the end of the experiment and the color test. However, the choices of color were not as exclusive as they had been

when the colors had indicated actual patch differences. This would be expected if real patch differences acted in combination with patch color to determine the strength of the final preference. Unfortunately, it was not possible to determine whether this discrepancy was caused by the lack of real patch differences or the passage of time between the end of the experiment and the color test. But the fact remains that if the jays had not carried any tendency for a patch preference between sessions (i.e. had always developed their preference completely within each session) then a preference for color between sessions should not have been learned.

In summary, the data demonstrate that blue jays are sensitive to variability in food resources and tend to avoid such variability. The actual patch choices, both overall and on the nonpreferred side, indicate strong, often exclusive risk aversion. The speed with which this risk aversion developed over the course of the experiment, combined with the lack of effect of reward condition and the result of the color preference test, may indicate that these jays have a natural tendency toward risk aversion and may rely at least partially on general cues of patch type. The forced phase choice latencies seem to show that this tendency toward risk aversion is strengthened from the start of every session so that some learning through sampling of the patches is involved in patch preference. In short, the strategy the jays used in choosing patches was to rely on general patch type cues, updating the validity of these cues through forced sampling early in the session.

The fact that the jays were always risk averse and that the strength of risk aversion did not change over conditions is a point in favor of the variance discounting model. As mentioned earlier, the only a priori capacity of the variance discounting model was to predict constant risk aversion under all conditions. This was certainly the case. The z-score model, on the other hand, predicted changing levels of risk aversion as patch

means changed and variance was constant. The lack of a reward condition effect runs counter to this prediction.

A die-hard proponent of energetic shortfall avoidance models, like the energy budget rule or the z-score model, would attribute their lack of support in this experiment to the absence of specific knowledge concerning the jays' daily energetic requirements and how these requirements were met by the reward conditions. One might argue that the unwavering risk aversion was caused by daily positive expected energy budgets. Energy budgets are typically said to have positive or negative expectations based on the rate of energy intake, relative to the rate required for survival, while the animal is foraging. In this way, the animal develops its patch preference while foraging, with no weight placed on any food it might encounter later (say, during an afternoon feeding). Thus, to say that these jays had a positive expected energy budget would mean that the rate of energy intake while they were in the chamber was equal to or greater than the rate necessary to sustain them through the night.

Energy budgets are based on two variables: energy need and resource availability (both in terms of time and amount). In this study, body weight might be a general index of energy need. Weights were consistent throughout the study, but they were consistently low. A blue jay at 80% of its free-feeding weight in the wild would certainly face a real possibility of starvation, a situation often demanding risk proneness (at least theoretically). The resources available in this study were tiny pieces of tiny mealworms, largely composed of chitin, and not very energetically potent. Expected patch values were also occasionally low (one reward per visit). Hence, it is possible that the jays were sometimes operating under a negative expected energy budget while in the chamber (but

they did not perish over night because they were fed in the afternoon). If nothing else, one could have reasonably expected varying levels of risk aversion across conditions.

Why then, were all five jays consistently, and persistently risk averse? Three possibilities spring to mind. First, energetic shortfall avoidance models traditionally assume that energetic needs should be calculated for, and projected over, 24- hour intervals of time (normally called days). These models have always been tested on small animals (e.g. sparrows, juncos, hummingbirds, shrews) with high energetic needs that can lead to death if not met over a short time span. For these animals, it might be advantageous to avoid, at all costs, energetic shortfalls over fixed, short spans of time. Blue jays are larger animals and have the ability to store more energy. It may be that larger animals have longer time windows over which energetic needs are specified. The time window for an animal that can store energy may also be less specific, increasing in size when stores are good, decreasing when stores are poor. Also, an increased ability to store energy, as long as the surplus need not remain constant (which it should not always do given the purpose of energy stores) may mean that absolute energy requirements are less fixed during the time window.

The second possible explanation for the proposed natural tendency toward risk aversion in blue jays involves their normal mode of foraging. Blue jays are omnivorous creatures, accused of everything from dominating bird feeders to stealing hatchlings from their nest. Certainly much of this is born from myth but, just as certainly, blue jays eat many things including seeds, insects, flora, and fauna (personal observation). In effect, being omnivorous may increase the number of low-variability patches from which the forager can choose. In turn, instances of not being able to find rich patches with low variability would be less numerous. Thus, it might make adaptive sense for an omnivore

to be naturally risk averse. While foraging in patches of one food type, say moth larvae in pine cones, if an omnivore discovers that the expected value of patches does not meet its energetic needs, it need not opt for higher variation in patches of the same food type. Rather, an omnivore has at least two options: (1) Switching to a different type of food in a similar patch, say pine seeds instead of moth larvae in the same pine cones, or (2) switching to a different type of food in a different type of patch, say *Catocala relict* moths on birch trees. This option has the versatility of still allowing risk-averse choices of patches with low variability and expected values that meet the energetic needs of the animal. Therefore, in the evolution of an omnivorous species, risk aversion would be the foraging mode of choice, the one through which the most members of the species survived. Risk aversion, though always preferable, is not as available to a specialized forager, like a sparrow, which may only be able to switch from patches of one seed type to patches of another seed type (that may be just as variable due to local competition and seasonal limitations).

Finally, the third possibility of constant risk aversion in these five blue jays involves a simpler rule governing response to risk. As a rule, energetic shortfall avoidance rules are complicated. Models like the energy budget rule and the z-score model are built upon precise mathematical quantification as well as a number of statistical assumptions. It can be argued that these models may even attribute abilities to the animal beyond those which are necessary to describe, predict, or explain behavior; abilities like projecting energy intake rate into the future relative to present requirements. It is far more in keeping with Morgan's Cannon, far more parsimonious, to choose a variable which is conceivably accessible to the animal on an immediate and continuous basis. Yet, since we are dealing with an energetic phenomenon, this variable should have some

immediate relationship to the animal's energetic state. The animal's body weight is such a variable. It does not require any great stretch of the imagination to assume that the size or weight of an animal's own body can become the subject of perception, or that changes in the body weight can be detected over certain spans of time.

Under many circumstances, a variable such as body weight might serve as input for a simple rule (a "rule of thumb") governing response to risk. For example, regardless of absolute weight, if any given weight were life-sustaining and remained constant or increased over the time window then risk aversion would be in order. Why take a gamble if you are surviving already? But if body weight decreased over a critical time window, a trend with serious repercussions, then risk proneness would become necessary. This drop in body weight should trigger the shift to risk proneness no matter the absolute weight from which the decrease began. But there could be a cushion before the shift in strategy is triggered to guard against unnecessary risk proneness in response to drops in weight well within normal variation. The cushion could be defined in terms of a variable time window related to the absolute body weight. For example, a 10-g drop in weight over a span of 20 hours from an initial weight of 450 g is not as threatening as a 10-g drop in weight over 20 hours from an initial weight of 300 g. In the former case, the trend might be allowed to continue over a longer time window before the animal becomes risk prone. In the latter case risk proneness might be invoked sooner to avoid the nasty and inevitable consequences

The rule proposed above is consistent with the data collected for this thesis. Simply put: even though the jays may have sometimes foraged under a negative expected energy budget while in the chamber, they were still fed later in the day. This need not imply that the jays were somehow aware that they would be fed every afternoon (though

increased activity among different groups of jays when their feeding time approaches suggests some temporal effect). The important point with respect to the proposed rule is that the jays' diet was supplemented every day to maintain them at a constant weight. These weights were low, but they were life-sustaining and fluctuated little over days. According to the proposed rule, if the weight remains constant over days, and the animal is alive, then risk aversion is the appropriate strategy to follow.

This rule should be effective in avoiding death by starvation, but at the same time it attributes no unreasonable abilities to the animal. This rule is simple, independent of fixed values and statistical assumptions, relying instead on changes in only one value over time, giving the rule the flexibility needed to adapt to a changing environment. The rule does not assume that the animal can project its needs into the future -- only that the animal is aware of recent changes in one, accessible value over a time window which can easily vary according to the animal's size, abilities, and mode of foraging. It should be noted that this rule still uses the logic of energetic shortfall avoidance, but in a simpler way than previous models. By using this rule, animals may be able to closely approximate the precise behavioral optima specified by more complex models of risk sensitivity.




Figure 15. Pure color preferences, averaged over subjects, from each of four color preference tests. Preferences were measured by the proportion of choices made for the color that had previously indicated the constant patch. This measure was calculated for both sides of the chamber combined (■) as well as for the nonpreferred side of the chamber alone (□).

APPENDIX

DERIVATION OF z WITH RESPECT TO σ

$$Z_R = \frac{R - \mu}{\sigma} = R \frac{1}{\sigma} - \mu \frac{1}{\sigma} = R\sigma^{-1} - \mu\sigma^{-1}$$

R is a constant fixed by the animal.

μ is assumed to be a constant across patches.

$$\frac{dz}{d\sigma} = -R\sigma^{-2} - (-\mu\sigma^{-2})$$

$$\frac{dz}{d\sigma} = -R \frac{1}{\sigma^2} + \mu \frac{1}{\sigma^2}$$

$$\frac{dz}{d\sigma} = \frac{-R}{\sigma^2} + \frac{\mu}{\sigma^2}$$

$$\frac{dz}{d\sigma} = \frac{-R + \mu}{\sigma^2}$$

$$\frac{dz}{d\sigma} = \frac{\mu - R}{\sigma^2}$$

BIBLIOGRAPHY

- Barnard, C. J., & Brown, A. (1985). Risk-sensitive foraging in common shrews (*Sorex araneus* L.). *Behavioral Ecology and Sociobiology*, 16, 161-164.
- Caraco, T. (1981). Energy budgets, risk and foraging preferences in dark-eyed juncos (*Junco hyemalis*). *Behavioral Ecology and Sociobiology*, 8, 213-217.
- Caraco, T. (1982). Aspects of risk-aversion in foraging white-crowned sparrows. *Animal Behaviour*, 30, 719-727.
- Caraco, T. (1983). White-crowned sparrows (*Zonotrichia leucophrys*): foraging preferences in a risky environment. *Behavioral Ecology and Sociobiology*, 12, 63-69.
- Caraco, T., Martindale, S., & Whitham, S. (1980). An empirical demonstration of risk-sensitive foraging preferences. *Animal Behaviour*, 28, 820-830.
- Krebs, J. R., Kacelnik, A. & Taylor, P. (1978). Test of optimal sampling by foraging great tits. *Nature*, 275, 27-31.
- Leventhal, A. M., Morrell, R. F., Morgan, E. F., & Perkins, C. C., Jr. (1959). The relation between mean reward and mean reinforcement. *Journal of Experimental Psychology*, 57, 284-287.
- Pubols, B. H., Jr. (1962). Constant versus variable delay of reinforcement. *Journal of Comparative and Physiological Psychology*, 55, 52-56.
- Real, L. A. (1980). Fitness, uncertainty, and the role of diversification in evolution and behavior. *American Naturalist*, 115, 623-638.
- Real, L. A. (1981). Uncertainty and plant-pollinator interactions: the foraging behavior of bees and wasps on artificial flowers. *Ecology*, 62, 20-26.
- Real, L. A., Ott, J. & Silverfine, E. (1982). On the tradeoff between mean and variance in foraging: an experimental analysis with bumblebees. *Ecology*, 63, 1617-1623.
- Stephens, D. W. (1981). The logic of risk-sensitive foraging preferences. *Animal Behaviour*, 29, 628-629.
- Stephens, D. W. & Charnov, E. L. (1982). Optimal foraging: some simple stochastic models. *Behavioral Ecology and Sociobiology*, 10, 251-263.

- Stephens, D. W. & Krebs, J. R. (1986). *Foraging Theory*. New Jersey: Princeton University Press, Monographs in Behavioral Ecology.
- Waddington, K. D., Allen, T., & Heinrich, B. (1981). Floral preferences of bumblebees (*Bombus edwardsii*) in relation to intermittent versus continuous reward. *Animal Behaviour*, 29, 779-784.
- Wunderle, J. M. & O'Brien, T. G. 1985. Risk aversion in hand-reared bananaquits. *Behavioral Ecology and Sociobiology*, 17, 371-380.

