


SCANNED BY


**OXFORD UNIVERSITY LIBRARIES
IMAGING SERVICE**

FROM THE COLLECTIONS IN

**THE BODLEIAN LIBRARY
UNIVERSITY OF OXFORD**



Centimeter



Inches

STOCHASTICITY IN FORAGING THEORY:

Risk and Information


David William Stephens

The Queen's College, Oxford

and

Edward Grey Institute

Department of Zoology, Oxford



Thesis submitted for the Degree of

Doctor of Philosophy

July 1982, Trinity Term

To my parents
for whom I wish I had a greater gift.

Preface

This thesis is largely a theoretical and mathematical work. I take this opportunity to point out how I have approached the theoretical problems discussed. Many early workers in optimal foraging theory have discussed the sweeping generality of their results. This search for generality has lead to two types of disappointment. The fieldworker who attempts to test the "general" theory is disappointed because he discovers that the theory is too recondite for him to apply. The theoretician is disappointed because it's usually easy to find small modifications which make the supposed generality evaporate.

Despite this lack of generality optimal foraging models have clearly been helpful guides to our intuition, and from this point of view simple models are perhaps more helpful than ponderous generality. The approach I adopt here is to study simple, analytical and (I hope) illustrative special cases. Except where I say so explicitly, I do not claim any generality. I prefer to let the reader decide how special the special cases presented are.

There are two views about mathematics in biological writing. There are those who favour mathematical appendices (the closet approach), and those who prefer to put the expressions straight in the text. I have adopted the second view, at the cost of some "turgidity," (in John Krebs's description). However, I have tried to write (and illustrate) bearing in mind the type of biological reader who skips equations. I am not sure I have always been successful; in cases where I have failed I hope that the reader who has a right to expect some biology may find it by working through the equations.

Acknowledgements

I am grateful to my parents, D.R. Stephens and B.H. Stephens, who have supported me and this work in every possible way.

I thank John Krebs, Alasdair Houston, and Anne Sorensen for reading all of the manuscript. In addition, I thank John and Alasdair for their helpful and constructive supervision throughout my time in Oxford, and their friendship.

I thank a number of people who have contributed to this work by reading and commenting on earlier manuscripts: T. Caraco, E.L. Charnov, T. Getty, R.F. Green, R.E. Hegner, J. Horsfall, A.I. Houston, J.R. Krebs, D.W. Lendrem, J.M. McNamara, N.E. Pearson, H.R. Pulliam, W.J. Sutherland, and R.C. Ydenberg. I have benefitted from correspondence and exchanging manuscripts with T. Caraco, R.F. Green, A.I. Houston, J. McNair, J.M. McNamara, N.E. Pearson, and G.H. Pyke. My thoughts on risk have been profoundly influenced by Tom Caraco, who has unselfishly shared his ideas with me. My thoughts on the value of sampling (chapter ten) originated with a conversation with John McNamara and Alasdair Houston. I will always be grateful for the encouragement and sound advice given to me by Ric Charnov. Special thanks to Robin McCleery for help of many kinds.

I am grateful to C.M. Perrins for tolerating me as a member of the Edward Grey Institute, and to A.S. Richford for tolerating me in the Alexander Library in the pre-coffee ban days. I am grateful to Allen Stevens, who was very nearly the first person I met in Oxford, for his continued friendship.

The figures were drawn by Denise Blagden. I thank an anonymous typist. Hairstyles by René. I cede all the rights and privileges of the Arbie Award to René.

Finally, I thank Anne Sorensen, who ought (at the minimum) to be second author on this thesis, for all help of all sorts. I thank Anne for making my life more pleasant, and for tolerating me as I was writing this.

Table of Contents

Abstract.....	i
Glossary of Symbols.....	ii
Chapter 1. Introduction to Stochasticity in Optimal Foraging Theory.....	1
1.1 Optimality in evolution and behaviour....	1
1.2 The state of the art in optimal foraging theory.....	2
1.3 Stochasticity.....	7
1.4 Introduction to risk.....	9
1.5 Introduction to information.....	10
PART I: RISK.	
Chapter 2: Risk: Theoretical Background.....	12
2.1 Introduction.....	12
2.2 Utility theory: Risk-Sensitivity.....	12
2.3 The biological theory or risk.....	18
2.4 Summary.....	20
Chapter 3: Are Animals Risk-Sensitive? A Review of the Evidence...21	
3.1 Introduction.....	21
3.2 The operant psychologists' evidence.....	22
3.3 The behavioural ecologists' evidence.....	22
3.4 Caraco et al's and Caraco's evidence.....	24
3.5 Summary and discussion.....	25
Chapter 4: Z-score Models of Risk.....	27
4.1 Introduction.....	27
4.2 Z-score models: the basic result.....	28
4.3 The generalized Z-score.....	31
4.4 Patch-use and the risk of starvation: a worked example.....	35

4.5 Discussion and summary.....	39
Chapter 5. Objections to the Theory.....	42
5.1 Introduction.....	42
5.2 The Lazy-L: death in the afternoon.....	44
5.3 The value of changing your mind.....	54
5.4 Carry-over.....	59
5.5 Summary and discussion- objections to the theory.....	71
Chapter 6. Risk: Conclusions and Discussion.....	74
PART II: INFORMATION	
Chapter 7. Information: Introduction to Part II.....	77
7.1 Introduction.....	77
7.2 Elements of the problem.....	79
7.3 The value of recognition.....	82
7.4 Patch sampling.....	85
7.5 Tracking the environment by sampling.....	87
7.6 Some unstudied problems.....	89
7.7 Habitat sampling and rules for learning..	91
7.8 Summary, discussion, and outline of part II.....	93
Chapter 8. Tracking a Changing Environment: Some Simple Theory....	96
8.1 Introduction.....	96
8.2 Optimal treatment of a run of bad luck...	97
8.3 The Estabrook-Jespersion-Arnold model.....	101
8.4 Discussion and summary.....	106
Chapter 9. Tracking a Changing Environment: A test.....	109
9.1 Introduction.....	109
9.2 Materials and methods.....	110
9.3 Results.....	115
9.4 Discussion and summary.....	119

Chapter 10. The Value of Patch Sampling.....	122
10.1 Introduction.....	122
10.2 A definition of the value of sampling...	123
10.3 The self-moderation problem.....	125
10.4 The importance of alternatives.....	129
10.5 Patch depression and patch sampling.....	133
10.6 Partial patch recognition.....	137
10.7 Discussion and summary.....	142
Chapter 11. Information: Conclusions and Discussion--Ignorance is Cheap.....	146
11.1 Conclusions.....	146
11.2 Ignorance is cheap.....	148
11.3 Complication.....	149
11.4 The importance of risk.....	150
Summary of Parts I and II.....	151
Literature Cited.....	154
Appendix I: Stephens and Ydenberg--Materials and Methods.....	168
Appendix II: Stephens 1981.....	170
Appendix II: Stephens and Charnov in press.....	172

i
ABSTRACT

Stochasticity in foraging theory: risk and information

(Submitted for the degree of D.Phil in Trinity Term 1982 by
David William Stephens, The Queen's College, Oxford).

This thesis considers the importance of adding stochasticity to models of optimal foraging behaviour. The problem is divided into two elements, risk and information, which are treated separately.

Part One: Risk

The theoretical and empirical results concerning animal preferences in risky situations are reviewed. Animals are known to show both risk-averse and risk-prone preferences over food reward. It is shown, theoretically, that a simple optimality model minimizing the probability of death due to starvation accounts for at least some of the observed patterns of preference. The model is generalized to consider preference from arbitrary combinations of mean and variance. Three limitations of the model are treated in detail, that is the importance of starvation by "ruin," mind-changing about risk preferences, and energetic carry-over are discussed. The implications and limitations of these models are outlined.

Part Two : Information

The theoretical literature is reviewed, and the problem of information is divided into three elements.

A simple model of environmental tracking is studied. The model suggests that there is a trade-off between sensitivity to change and the costs of sampling. The model is tested using great tits (Parus major) foraging in an aviary. The trend in sampling was as predicted, but the birds were less sensitive to change than predicted.

The problem of patch sampling is critically discussed. The value of sampling is defined, and this definition is used to compare the assumptions of previous models. Three such problems are treated: the importance of variance in the mixing distribution of patch sub-types; the importance of alternative and unambiguous patch types; and the importance of patch depression. It is concluded that previous models have often over-valued sampling. A simple and natural model of partial patch recognition is considered, and is shown to have empirical support. Implications and limitations of these models of information are discussed.

Glossary of Symbols and Conventions

The following symbols and conventions have been used:

Means and variances:

Generally:

$E(X)$ - the expectation of a random variable X .

$V(X)$ - the variance of a random variable X .

Specifically:

μ - the mean of a random variable.

usually the mean of energy gain.

σ^2 - the variance of a random variable.

usually the variance of energy gain.

Other terms and conventions:

$P(\text{event})$ - indicates the probability of an event.

$P(\text{event-1}|\text{event-2})$ - indicates the conditional probability of event-1 given event-2.

$\max$ - maximum, used as a subscript, or to denote the operation of finding the maximum.

$\min$ - minimum, used as a subscript, or to denote the operation of finding the minimum.

$\phi(x)$ - the probability distribution function of a standard normal ($\mu = 0$, $\sigma = 1$) distribution. (x is an arbitrarily chosen variable.)

$\Phi(x)$ - the cumulative distribution function of a standard normal distribution; i.e. If Y has a standard normal distribution, then $\Phi(x) = P(Y < x)$.

ϵ - used to mean "is contained in," (see e.g. below).

Glossary of symbols and conventions (page2)

$[a,b]$ - a closed interval, (see e.g. below).

(a,b) - an open interval, (see e.g. below).

$\{a,b\}$ - the set consisting of a and b , (see e.g. below).

Therefore:

$x \in [a,b]$ means $a \leq x \leq b$.

$x \in (a,b)$ means $a < x < b$.

$x \in \{a,b\}$ means x may be a or b .

$*$ - is used to denote a critical or an optimal value.

λ - is used only to mean an encounter rate.

CHAPTER ONE

Introduction to Stochasticity in Optimal Foraging Theory

1.1 Optimality in evolution and behaviour

In the last twenty years mathematical models of optimality and stability have become important tools in the study adaptation by natural selection (see Maynard Smith 1978 for a review). This has been especially true for the study of the functional value of animal behaviour, (McFarland 1977). The central idea is that natural selection can be viewed as an optimizing process. This view has been controversial (Gould & Lewontin 1978), and its limitations have been constructively reviewed by McCleery (1978) and Oster and Wilson (1978). Despite the controversy and limitations, there can be little doubt that optimality has been a profoundly influential idea in the study of animal behaviour as an adaptation.

Currencies and Constraints. All optimality models have two elements in common. The modeller must specify a currency of maximization and the constraints of behaviour which limit this maximization (Schoener 1971), although McFarland (e.g. McFarland 1977) takes a somewhat different view; that is, he uses observed behaviour to specify the currency of maximization, but this is a difference in the application of the models not in the elements of the models themselves.

Many problems have been studied by optimality techniques. For example, the theory of optimal life history strategies (see Horn 1978) often assumes that natural selection has acted to maximize the number of offspring an individual produces over its lifetime, taking into account that the animal's productivity is constrained by an inverse relationship between annual survival and annual reproduction. In addition to life history theory, there have been theoretical and

experimental optimality or stability analyses of territorial behaviour (Brown 1964, Brown & Orrians 1970, Pyke 1979), animal conflicts (Maynard Smith and Price 1973), breeding systems and behaviour (Orrians 1969), behavioural sequences (Sibly & McFarland 1976), flock size (Caraco 1980, Pulliam 1973), insect sociality (Oster and Wilson 1978), the allocation of sexual resources (Charnov 1982) and foraging behaviour (Pyke et al. 1977, Krebs 1978), Werner and Mittelbach 1981). Among these analyses, optimality models of foraging behaviour have certainly been one of the most influential applications of optimality in biology. Thus the first paper on optimal foraging was published in 1966, (MacArthur & Pianka 1966) representing only .05% of the literature in 1966 contained in American Naturalist, Animal Behaviour, Ecology, and The Journal of Animal Ecology; by 1974 the proportion of papers in these journals indexed as "optimal foraging" was 2%, and 8% in 1981 (source Bodeleian Library computerized search). This is exponential growth!

1.2 The state of the art in Optimal Foraging theory

The theory of optimal foraging makes very simple assumptions about currency and constraints. It has normally assumed that foraging behaviour has been "designed" by natural selection to maximize the net rate of energy gained by a forager, an hypothesis that I shall call "rate maximization." The most fundamental constraint in optimal foraging is the assumption that animals cannot search for food and eat simultaneously. I call this the search/handling constraint. These two ideas, rate maximization and the search/handling constraint, follow from independent considerations of natural history. It is argued (Schoener 1971, Pyke et al. 1977) that the rate of energy gain should be maximized because animals will be limited by energy (the

more energy an animal gets per day, the better off it is) or by time (the sooner it satisfies its daily needs, the more time it has for other activities). Rate maximization is a sensible strategy in either case, (Schoener 1971, Pyke et al 1977, Krebs 1978). The considerations which suggest the search/handling time constraint need little explanation; in most (but not all; McNair in press) cases a forager will only be able to do one thing at a time.

The basic results of optimal foraging deal with the implications of putting these two ideas together, within a framework of other assumptions about encounter rates, prey or patch values, etc. Most, if not all, of the resulting models hinge on an idea which is the natural result of combining rate maximization with the search/handling constraint which I call the "principle of lost opportunity." For example, in standard models of optimal diet (Charnov 1976b, Maynard Smith 1974, McNair 1979, Pearson 1974, Pulliam 1974, Schoener 1969, Werner and Hall 1975) it is shown that the decision about whether a particular prey type should be included in the diet depends only on the abundance of better types; that is the decision is based on whether taking this item also requires an unacceptable loss in opportunity to take better items which might be encountered in the time it takes to handle the item at hand.

The principle of lost opportunity arises in a similar way in simple models of patch-use, (Charnov 1976a, Orians and Pearson 1979, Stephens and Charnov in press, McNair in press). A forager should leave a patch before the patch is empty because the value (the marginal value) of the patch will normally decline as the forager searches the patch. This means that there will be a patch residence time beyond which staying in the patch would involve a loss of opportunity (to leave and do better elsewhere).

Both rate maximization and the search/handling constraint have been challenged theoretically, (e.g. Caraco 1980, McNair in press, Stephens 1981). For example, the first five chapters of this thesis represent an argument against rate maximization in some cases. McNair (in press) considers a patch model in which it is possible to encounter new patches while foraging within a patch, a violation of the search/handling constraint. McNair's result is an exception which proves the rule for the principle of lost opportunity; e.g. if the rate at which new patches are encountered while hunting within a patch is the same as the rate of encounter when hunting for new patches, then there is no lost opportunity, and the forager should stay until the patch is empty. Despite these exceptions the ideas of rate maximization, the search/handling constraint, and their corollary the principle of lost opportunity are of general importance in foraging theory.

The conventional theory has made a large number of other assumptions. It has assumed that prey and patch types can be instantaneously recognized without ambiguity, (an assumption I spend chapters seven to eleven relaxing and analysing). It has assumed that encounters are independent of one another; this is what I call (at slight variance to previous usage) the fine-grained assumption. The assumption which is the starting point for the work presented in this thesis, is that most (but not all) conventional models have used deterministic arguments to derive their results.

Many types of models, all based largely on the ideas discussed above, have been presented. I list four types of models here. (A) Diet Models (Charnov 1976b, Maynard Smith 1974, McNair 1979, Pearson 1974, Pulliam 1974, Schoener 1969) predict what types of items should be

included in the diet. The basic result is that the ratio of net energy gained per encounter with a type to the handling time (E/h) should be used to rank diet items, (E/h is often called the profitability, Krebs 1978). (B) Patch-Use Models (Charnov 1976a; the marginal value theorem, mvt) consider the case where patches "depress" (the terminology of Charnov et al 1976); that is patches get worse, in terms of rate, as they are exploited. The models predict that foragers should stay longer in patches when the rate of gain elsewhere is lower (a direct lost opportunity result). (C) Central-Place Models (Andersson 1978, Orians and Pearson 1979) consider how animals which are constrained to return food to a central-place (say a parent bird returning food to its young) should forage. The models make predictions of increased selectivity and patch residence time at greater distances from the central place. (D) Other Models: more recently more complex problems have received attention. For example, the trade-off between foraging and territory defence has been studied by Davies and Houston (1981).

Theory and Experiment. I have recently reviewed the literature of empirical tests of foraging theory elsewhere (together with J.R. Krebs and W.J.Sutherland; Krebs et al. in press). In table 1.1 I reproduce Krebs et al's table 1, which summarizes the empirical evidence bearing on optimal foraging theory to date. The table follows the classification of the previous paragraph: (A) Diet models, (B) Patch-use models, (C) Central-place foraging models and (D) Other models. The table shows that the balance between theory and evidence in optimal foraging has remained good. This is in marked contrast to some other aspects of theoretical ecology, for example, niche and competition theory. The table also shows that many

TABLE 1.1

A. PREY/DIET

Reference	Species	Result	Lab/Field	Agreement with model
Belovsky 1978	Moose	Maximize rate of energy intake subject to sodium constraint.	F	Quantitative agreement with nutrient constraint model.
Barnard & Stephens 1981	Lapwing	Select profitable size classes of earthworm; more selective at higher prey density	F	Qualitative agreement with diet model
Barnard and Brown 1981	Shrews	Prefer larger of two sizes of mealworm even though E/h smaller; greater preference at higher encounter rates	L	Preference based on size not profitability
Davidson 1978	Harvester ants	Specialization increases with seed density	F	Consistent with diet model.
Davies 1977 a&b	(a)Pied & Yellow wagtails (b) Spotted flycatcher	Select more profitable size classes of insects; more selective when profitable prey more abundant; not effect of abundance of less profitable sizes	F	Qualitative agreement with diet model.
Ebersole & Wilson 1980	Mice	No change in diet diversity with seed density; or greater handling time increases specialization.	L	Relevance to OFT cannot be judged; no measure of E/h presented
Elner & Hughes 1978	Shore Crab	Choose close to most profitable size mussels; take less profitable prey after run of bad luck; selectivity greater at higher densities; density of less profitable sizes affects selectivity	L	Qualitative agreement with recognition time model; but partial preferences.

TABLE 1.1

A. PREY/DIET

Reference	Species	Result	Lab/Field	Agreement with model
Erichsen et al 1980 (see also Rechten et al. 1981.)	Great tit	Choose more profitable of two prey sizes of mealworm, including effects of crypsis; selectivity changes with encounter rates.	L	Quantitative agreement with recognition time model; but partial preferences
Gardner 1981	Bluegill sunfish	Rejection of small <u>Daphnia</u> not due to apparent size effects.	L	Qualitative agreement with diet model.
Gibson 1980	3-spined stickleback	Rejection of small <u>Daphnia</u> not due to apparent size effects.	L	Qualitative agreement with diet model.
Goss-Custard 1977a	Redshank	Prefer most profitable sizes of polychaetes; selectivity increases with encounter rate with profitable prey; no effect of less profitable sizes.	F	Quantitative agreement; but partial preferences.
Goss-Custard 1977b	Redshank	<u>Corphium</u> preferred to polychaetes, although lower E/h.	F	Not consistent with energy maximizing.
Hames & Vickers 1982	Amazonian Indians	Lower ranking prey added to diet as prey stocks depleted.	F	Consistent with diet model.
Houston et al. 1980 (see also Rechten et al. 1981)	Great tits	Test effects of making prey hard to discriminate.	L	Qualitative agreement in some treatments and individuals but not others.
Hughes & Elner 1979	Shorecrabs	No selectivity in relation to profitability.	L	Not in agreement with diet model.
Hughes and Seed 1981	Shorecrab	Choose smallest mussels which also have highest E/h.	L	Agreement with preference based on size or E/h.

TABLE 1.1

A. PREY/DIET

Reference	Species	Result	Lab/Field	Agreement with model
Jaeger & Barnard 1981	Salamander	Increased specialization in large <u>Drosophila</u> with increasing density; little effect of altering density of small <u>Drosophila</u>	L	Energy gain would have been higher with no selection.
Kaufman & Collier 1981	Rats	Prefer husked seed over unhusked seeds when both provided ad lib.	L	Consistent with diet model.
Kislaliogu & Gibson 1976	15-spined stickleback	Prefer most profitable (E/h) sizes of prey.	L	Consistent with diet model.
Krebs et al 1977 (see also Rechten et al. 1981)	Great tits	Selection for more profitable sizes of mealworms increases with increasing encounter rate; little effect of encounter rate with less profitable prey.	L	Quantitative agreement for some individuals; but partial preferences.
Lea 1979	Pigeons	Pigeons in Skinner box prefer schedule with higher E/h preference increases with decreasing "search time;" immediate small reward preferred to larger reward after delay, even if E/h is the same for both.	L	Qualitative agreement with diet model; but partial preferences; failure to show "self control" not consistent with model.
Lewis 1980	Grey Squirrel	Prefer profitable (E/h) nut species within a patch.	F	Consistent with diet model.
Lobel & Ogden 1981	Parrotfish	Prefer algal species with high net E/h except for one toxic species.	L/F	Consistent with prey choice influenced by energy, toxins and nutrients; not consistent with simple energy maximizing.

TABLE 1.1

A. PREY/DIET

Reference	Species	Result	Lab/Field	Agreement with model
Mittelbach 1981 (see also Werner & Mittelbach 1981).	Bluegill sunfish	Preference for more profitable sizes of <i>Daphnia</i> increases with prey density.	L/F	Quantitative agreement; but partial preferences.
Milton 1979	Howler monkeys	Preference for young leaves with high protein/fibre ratio.	L/F	Consistent with quality rank- ing; not energy maximizing.
Owen-Smith & Novellie 1982	Kudu	Selectivity not maximizing intake.	F	Inconsistent with protein maximizing; no information of energy maximizing.
Palmer 1979	Dogwhelk	Prefer prey species/sizes with highest E/h (and yielding highest growth rate). More selective at high prey densities.	F	Quantitative agreement with diet model.
Pastorak 1980	<u>Chaeoborus</u>	More selectivity when less hungry.	L	Consistent with diet model.
Pleasants 1981	Bumblebees	Preference for flowers/patches with high profitability (E/h).	F	Consistent with diet model.
Pulliam 1980	Chipping sparrow	Preference for high profitability (E/h) seeds with some notable exceptions; more selective when density high	F	Partial agreement with diet energy maximizing.
Rapport 1980	<u>Stentor</u>	Prefer mixed diet. Fitness higher when eating mixed diet.	L	Consistent with complementary resources model; but not with energy maximizing

TABLE 1.1

A. PREY/DIET

Reference	Species	Result	Lab/Field	Agreement with model
Ringler 1979	Trout	Prefer most profitable (E/h) sizes of prey.	L	Apparent agreement with diet model.
Stein 1977	Smallmouth bass	Select crayfish with highest E/h (h includes wasted pursuit time in unsuccessful attempts.)	L	Consistent with diet model.
Tinbergen 1981	Starling	Sometimes prefer prey species which does not have highest E/h.	F	Energetic value of prey less important than prey quality.
Waddington & Holden 1979	Honeybees	Preference for more profitable (E/h) depends on relative abundance.	L	Quantitative agreement with simultaneous encounter model in average results; individuals vary
Werner & Hall 1974	Bluegill sunfish	Preference for more profitable (E/h) size classer of <u>Daphnia</u> increases with density.	L	Qualitative agreement with diet model.
Winterhalder 1981	Cree Indians	Selectivity for profitable (E/h) prey species increased with increasing encounter rate (due to acquisition of snowmobiles).	F	Qualitative agreement with diet model.
Vadas 1977	Sea Urchins	Preference for species of algae that yield highest reproductive success and have highest assimilation efficiency, but not highest calories.	F	Consistent with maximizing assimilation efficiency, but not energy intake.

TABLE 1.1

B. PATCH USE

Reference	Species	Result	Lab/Field	Agreement with model
Best & Bierzychudek 1982 (See also Pyke 1978a)	Bumblebees on foxglove	Within inflorescence diminishing returns from bottom to top; departure rule consistent with energy maximizing.	F	Quantitative agreement with patch-depression model.
Bond 1980, 1981	Green lacewing	Giving-up time increases with hunger level	L	Relevance to patch depression model not clear (patches apparently non-depressing).
Cook & Cockrell 1978	Notonecta & <u>Ladybird</u> larvae	Prey as patches; food extracted per prey declines with increasing prey density.	L	Qualitative agreement with marginal-value theorem (mvt) but curve relating patch time to travel time wrong shape.
Corbet et al 1981 (see also Waddington & Heinrich 1979	Bumblebees on <u>Scrophularia</u>	Bee move upwards on vertical inflorescence even though nectar rewards show no clear pattern of vertical distribution	L	Vertical movement rule not related to maximizing energetic gain.
Cowie 1977	Great tit	Patch time changes when travel time altered.	L	Quantitative agreement with mvt when energy costs considered.
Giller 1980	<u>Notonecta</u>	As for Cook & Cockrell, but time per prey related to previous intercatch interval.	L	Supports gut-filling model rather than mvt

TABLE 1.1

B. PATCH USE

Reference	Species	Result	Lab/Field	Agreement with model
Hartling & Plowright 1979	Bumblebee	More persistence within a patch with longer travel distances between patches.	L	Qualitative agreement with mvt.
Heinrich 1979	Bumblebee on (a) <u>Trifolium</u> (b) <u>Aconitum</u>	(a) Visit more flowers per head at low head densities. (b) Move from bottom to top even though nectar rewards higher at top	F	(a) Relevance to mvt not clear since patch depression not measured. (b) Bee's movement rule does not maximize energetic gain; may reduce revisiting.
Hodges & Wolf 1981	Bumblebees on <u>Delphinium</u>	Nectar extracted from individual flowers related to standing crop: more nectar left behind in rich sites.	F	Quantitative agreement with mvt
Hodges 1981	Bumblebees	Departure rule from <u>Delphinium</u> plants based on rate of intake	F	Mechanism produces results in agreement with energy-maximizing.
Howell & Hartl 1980	Glossphagine bat	Nectar left behind in artificial inflorescences consistent with rule that maximizes benefit/cost per patch	L	Results not consistent with mvt.
Hubbard & Cook 1978	<u>Nemeritis</u> (<u>Ichneumonid</u> wasp).	Allocation of time to host patches of different density similar to that which would maximize encounter rate with unparasitized hosts.	L	Quantitative agreement with patch use model.

TABLE 1.1

B. PATCH USE

Reference	Species	Result	Lab/Field	Agreement with model
Larkin 1981	Barbary dove	On progressive interval schedules with reset after switching and variable change over delay, birds maximize food or water intake. Energetic costs influenced switching as predicted.	L	Quantitative agreement with mvt
Lewis 1980	Grey squirrel	Prefer to feed in patches with highest intake per unit time.	F	Qualitative agreement with patch use model.
Parker 1978 (see also Parker & Stuart 1976).	Male dungfly	Duration of copula and guarding consistent with maximizing fertilization rate.	F	Qualitative agreement
Pyke 1978a (see also Pyke 1979)	Bumblebee on <u>Delphinium</u>	Nectar rewards diminish from bottom to top; bee starts at bottom and works upwards; departs before top.	F	Qualitatively consistent with mvt
Pyke 1978b	Hummingbirds on Scarlet gilia	Number of flowers visited per inflorescence increases with number available; consistent with departure rule based on expected gain.	F	Qualitatively consistent with depressing patch model.
Sih 1980	<u>Notonecta</u>	Similar to Cook & Cockrell	L	Qualitative agreement with mvt.

TABLE 1.1

B. PATCH USE

Reference	Species	Result	Lab/Field	Agreement with model
Townsend & Hildrew 1980	Caddis fly larvae	Giving-up time equal to average capture rate in field; no variation in GUT between habitats.	L/F	Relevance to patch depletion model not clear; patches do not deplete.
Whitham 1977	Bumblebee on	Amount of nectar left behind Desert willow in flowers decreases as general level of availability in habitat decreases.	F	Qualitative agreement with mvt
Waage 1979	Nemeritis (Ichneumonid wasp).	Patch-leaving rule is a result of interaction between habituation to host odour and incremental effect of successful oviposition.	L	Simple rule produces results qualitatively consistent with OFT
Waddington & Heinrich 1979	Bumblebee	Bees move upwards on vertical artificial inflorescences regardless of vertical distribution of nectar rewards.	L	Rule of thumb used by bee not sensitive to variation in nectar rewards.
Williams 1982	White rat	On progressive ratio schedules with reset key to simulate travel time "switch between patches" in a way predicted from OFT.	L	Quantitative agreement with mvt

TABLE 1.1

A. PATCH USE

Reference	Species	Result	Lab/Field	Agreement with model
Ydenberg in prep.	Great tit	Similar to Williams	L	Quantitative agreement with mvt
Zimmerman 1981	Bumblebee on <u>Polemonium</u>	When travel time between inflorescences increases new flower species included in diet; no clear increase in number of flowers visited per inflorescence.	F	Significance for mvt not clear since gain curves not measured.

TABLE 1.1

C. CENTRAL PLACE FORAGING

Reference	Species	Result	Lab/Field	Agreement with model
Anderson 1981	Whinchat	Search effort inversely related to distance from nest; adding food to territory caused search effort to shift towards nest.	F	Qualitative agreement with model of allocation search time.
Carlson & Moreno 1981	Wheatear	Evidence for loading curves; load size increases with distance; time allocation different to patches at some curves not as predicted.	F	Qualitative agreement with load-size predictions.
Davidson 1978	Harvester ants	Take larger prey at greater distances from nest.	F	Consistent with single-prey loader model.
Evans 1982	Blackbilled gull	Intensity of use of feeding site inversely related to distance from central place	F	Qualitative agreement with model of allocation of search time.
Giraldeau & Kramer in press	Eastern chipmunk	Load size increases with distance.	F	Qualitative but not quantitative fit with multiple-prey loader model.
Hegner 1982	White-fronted bee-eater	Selectivity of parents for large prey to nest increases with distance from nest to feeding site.	F	Qualitative agreement with single-prey loader

TABLE 1.1

C. CENTRAL PLACE FORAGING

Reference	Species	Result	Lab/Field	Agreement with model
Heithaus & Fleming 1978	Frugivorous bats	Search intensity decreases as a function of distance from roost.	F	See Evans
Jenkins 1980	Beaver	Select smaller trees at greater distances	F	Appears to conflict with single-prey loader model, but travel cost and "pursuit time" in relation to tree size not measured.
Kacelnik in prep.	Starling	Load size increases with distance from nest in experimental patches with progressive interval schedules.	F	Quantitative agreement with multiple-prey loader models.
Kramer & Nowell 1982	Eastern chipmunk	Load size increases with distance.	F	Qualitative agreement multiple-prey loader model.
Krebs & Avery in press	European bee-eater	Selectivity of parents (for large prey) bringing prey to nest increases with travel distance; not due to changes in prey availability.	F	Quantitative agreement with single-prey loader model.
Killen et al. 1981	Norway rat	Rats accumulate larger "load" of pellets behind a door when "travel time" increased (delay before access to food).	L	Apparent agreement with mvt but not clear whether there is a depression effect.

TABLE 1.1

C. CENTRAL PLACE FORAGING

Reference	Species	Result	Lab/Field	Agreement with model
Tinbergen 1981	Starling	Demonstration of loading curve. Load size of parents increases with distance to nest.	F	Qualitative agreement with multiple prey-loader loader model.

TABLE 1.1

D. OTHER MODELS

Reference	Species	Result	Lab/Field	Agreement with model
(1) Tradeoffs:				
Davies & Houston 1981	Pied Wagtail	Exclusion or tolerance of of satellites on feed- ing territories varies with prey density as predicted by intake maximiz- ing model.	F	Quantitative agreement with intake maximizing subject to constraints of defence time and competition for food.
Dunstone & O'Connor 1979	Mink	Animals maximize capture rate subject to constraint of oxygen reserves while hunting underwater	L	Qualitative agreements in predictions of pursuit time, giving-up time and dive frequency.
Heller & Milinski 1979 (see also Milinski & Heller 1978	3-spined stickleback	At high hunger levels fish choose to attack high densities of <u>Daphnia</u> ; at low hunger levels they prefer lower prey densities; increased predation risk causes hungry animals to behave like satiated ones.	L	Qualitative agreement with dynamic optimization tradeoff between feeding and vigilance.
Kacelnik et al. 1981	Great tit	Birds sacrifice foraging efficiency for territorial vigilance.	L	Qualitative agreement with tradeoff between feeding and vigilance.
Krebs 1980	Great tit	Handling time increases with cumulative food intake, due to increased vigilance while handling prey.	L	Qualitative agreement with dynamic optimization trade- off between hunger and vigilance.

TABLE 1.1

D. OTHER MODELS

Reference	Species	Result	Lab/Field	Agreement with model
Martindale 1982	Gila wood-pecker	Parents make shorter trips (smaller loads) closer to nest delivering food to young after presentation of stuffed intruder near nest.	F	Qualitative support for tradeoff between foraging and nest defence.
Mittelbach 1981	Bluegill sunfish	Small fish (but not large ones) avoid feeding in profitable sites because of predation risk.	F/L	Qualitatively consistent with trade-off between feeding and predation risk.
Sih 1979	<u>Notonecta</u>	Small individuals (but not large ones) avoid profitable feeding areas because of predation risk.	F/L	See Mittelbach
<u>(ii) Energetic costs of foraging and morphological constraints</u>				
DeBendictis et al. 1978	Hummingbirds	Meal size can be interpreted as tradeoff between cost of carrying nectar in crop and cost of travelling to flowers.	F/L	Quantitative agreement
Pyke 1981	Hummingbirds & Honeyeaters	Difference in feeding made (hovering vs. perching) can be explained in terms of energetic cost and movement speed.	F/L	Qualitative account of interspecific difference.

TABLE 1.1

D. OTHER MODELS

Reference	Species	Result	Lab/Field	Agreement with model
Wilson 1981a,b	Leafcutting ants	Size at which workers become leaf harvesters maximizes energy gain/cost for colony.	L/F	Consistent with colony energy-efficiency model.
Zach 1979	Northwestern crow	Height from which crow drops mollusc shells to break them minimizes cost (in upward flight) per shell.	F	Quantitative agreement with cost minimizing.

studies have found empirical support for the theory, although there are notable failures. Perhaps the most common conclusion has been to find qualitative support for the theory (e.g. the trend is as predicted) and quantitative failure (e.g. the details of the behaviour are not exactly as predicted). This suggests that the theory is on the right track, but in need of some refinement. Although many refinements have been suggested, perhaps the most frequently mentioned difficulties are violations to the assumption that the world is deterministic (cf. Oaten 1977, Mittlebach & Werner 1981).

Limitations and Criticisms. There have been two types of criticism leveled against optimal foraging theory. Firstly, there have been criticisms about whether it is a reasonable approach at all, (cf. Gould and Lewontin 1978). Secondly, there have been criticisms about the techniques and assumptions made in foraging theory. I deal with the first type of criticism first.

Specifically Gould and Lewontin (1978) have criticised optimality in general on two counts, that is it not testable and that it ignores constraints. Their argument about testability is that the hypothesis that a particular behaviour is "optimal" cannot be falsified. This is true, but a particular model, which represents a potentially complex set of assumptions about currencies and constraints, can be shown either to be a sensible description of nature or not. Their criticism of optimality's treatment of constraints leads to a more general question. Their criticism cannot be that optimality ignores all constraints because, as I have pointed out above, constraints are a fundamental part of all optimality models. Their criticism can only be interpreted as meaning that optimality should consider more constraints, in particular constraints of morphology and phylogeny.

This leads to the question of whether the ideal model is a perfect representation of reality or only captures a few essential features; how much complexity should be included? Optimal foraging theory comes down of the side of simple and illustrative models, and this is the approach I adopt in this thesis. The empirical and intuitive (in the sense of stimulating much experiment and discussion) success of optimal foraging theory argues in favour of this approach.

There are two important technical criticisms. Firstly, all the simple models discussed above find the optimal behaviour using nothing more complicated than elementary differentiation. This means that none of these models allow for the possibility that behaviour may change as a function of the forager's state. To account for this, techniques from the mathematic theory of dynamic programming are required, (e.g. McFarland and Houston 1981). In this thesis I follow the traditional "static" approach of optimal foraging, and this is an acknowledged difficulty, (in fact it arises in chapter five). Secondly, as I pointed out earlier, the majority of earlier models of optimal foraging used deterministic arguments. Foraging animals search for, encounter and handle their prey according to rules which are most accurately described by statements of probability, that is, by stochastic models. In this thesis I consider the implications of stochasticity for optimal foraging theory.

1.3 Stochasticity

Definition. If the outcome of some process can be determined exactly and without error it is deterministic. However, there are many processes in nature where the best description that can be made must be in terms of probabilities, such processes are stochastic. I adopt here the definition of Coleman (1977), that "stochastic is jargon for

random." However, the value of using the word "stochastic" instead of "random" is that "random" is often used to mean uniformly random. The outcome of a coin toss is said to occur "at random," this means that heads and tails are equally likely. If heads were two times more likely than tails then in common usage this would not be "at random" but it would be stochastic.

The Stochastic Problem. The problem of stochasticity in optimal foraging is actually two problems, risk and information.

Risk.

In a stochastic environment a forager's behaviour will not yield, for example, ten calories with certainty; instead a forager will, at best, be able to choose a probability distribution of calories. Choosing from a known set of probability distributions and choosing from a set of certain alternatives are clearly different problems. The first stochastic problem in foraging theory is how should foragers change their preferences when they are faced with known types of stochastic variability. Micro-economists, whose models closely parallel optimal foraging, call this the problem of pure risk (Hey 1979).

Information

The second stochastic problem is more complicated and less direct. Imagine that the state of nature is represented by one of two possible probability distributions, A and B. Assume that the forager knows what the alternatives are, but cannot distinguish between them before deciding what tactic to use. Further imagine that if the forager knew it was faced with distribution A it should adopt a different tactic than if faced with B. The tactic actually chosen should then reflect information about the relative likelihoods of A

vs. B, gained while foraging. That is, the forager has a problem of information gathering and use. In the terminology of micro-economics this is a problem in "pure uncertainty" (Hey 1979). However, for clarity, I follow the usage of Stephens and Charnov (in press), and call this the problem of information.

The separation of the stochastic problem into these two elements (risk and information) is logically important, not because animals actually face pure information or pure risk problems in nature, but because it suggests a sensible way for the study of stochasticity to proceed. In this thesis I consider risk and information separately because, at least as a first step, this approach leads to the simplest and most understandable models. In the next two sections of this chapter I introduce the topics of risk and information separately.

1.4 Introduction to Risk

The problem of risk is how should foragers choose from a set of variable alternatives. In other words, how should a forager value, or prefer, variable alternatives. This leads naturally to considerations of how risk should change the currencies of optimization used in foraging theory. I therefore treat risk as a "currency problem." My approach is not to try to specify general currencies that will always hold, but to consider the implications of simple and special types of currencies of optimization for risk. Specifically I consider models based on minimizing the probability of death due to starvation.

The problem of risk is examined in five chapters. In chapter two I review the mathematical and biological theories which have been put forward to account for risk. In chapter three, I review the empirical evidence regarding animal preferences over variability. In chapter four, I present a detailed account of the risk of starvation currency

and its implications. In chapter five, I consider three difficulties with the theory of chapter four. Finally, in chapter six, I discuss my results as a whole.

1.5 Introduction to Information

The problem of information is how should a forager treat aspects of its foraging environment (e.g. patches, prey items, etc.) which it knows to be ambiguous. For example, a small bird may know that all red beetles are either tasty or noxious, but may not know whether a given red beetle is either tasty or noxious.

If risk is a currency problem, then information is a problem of constraint; i.e., what is optimal behaviour if a forager is unable to know which state is which? Information is not as straightforward a problem as risk; because there are many ways a forager can be ignorant (e.g. a forager may recognize nothing about beetles, or recognize only their size, but not their flavour) and many ways it can gain information (e.g. it may directly sample or "pay" recognition costs). Secondly, it is not as direct an implication of stochasticity as risk. There is variability in any stochastic situation, but there need not be a multiplicity of indistinguishable states. I do not present general models of information; instead I have tried to present simple and understandable models which illustrate what I think are important points.

I consider information in five chapters. In chapter seven, I review the theoretical literature of information. In chapter eight I consider a simple model of optimal sampling to track a changing environment; I test this model in chapter nine. In chapter ten I present a series of simple models which criticize the conventional approach of the patch sampling literature. In chapter eleven I

summarize and discuss my work on information.

Parts I and II.

In this thesis I consider stochasticity in two separate parts. Part I, Risk (chapters two, three, four, five, and six) and Part II, Information (chapters seven, eight nine, ten, and eleven). Each part is self-sufficient. In each part I review the relevant theoretical and empirical literature. In chapters six and eleven I provide separate discussions of risk and information.

PART I: RISK

CHAPTER TWO

Risk: Theoretical Background

2.1 Introduction

The problem of risk is how should a forager choose between known but probabilistic outcomes. To study this question it is necessary to use ideas from mathematical statistics and the microeconomic theory of decision-making. In this chapter, I review the basic theory required. In addition, I review biological ideas about animal preference for and against variability.

Mean Maximization. The problem of risk has usually been treated glibly in optimal foraging theory. The conventional view is that if (in a deterministic model) animals are assumed to maximize their rate of energy gain, then in a stochastic case they should maximize their mean rate of energy intake (see Pyke et al., 1977). For convenience, I call this the mean-maximization hypothesis. The mean-maximization hypothesis has a troublesome implication. It requires that foragers be indifferent to variability in food reward. Suppose a forager has a choice between two distributions of energy gain; the first has mean 10 kcals. and variance 100; the second has mean 9.99 kcals and variance 1. Mean-maximization requires that the forager choose the first, although many people would clearly prefer the second (if the reward were money rather than calories). The mean-maximization hypothesis and its implication of variance-indifference are at the centre of my treatment of risk (Part I) ; throughout Part-I ,I will be considering when mean-maximization and variance indifference are valid, and when they are not.

2.2 Utility Theory: Risk-Sensitivity. The most general solution to the problem of risk is provided by utility theory. Utility theory was

first presented by Von Neuman and Morgenstern (1947) in their famous monograph "Theory of Games and Economic Behavior". In developing a theory of economic behaviour Von Neumann and Morgenstern faced the problem of finding a way to compare two entirely different probability distributions; that is they faced the problem of risk. In outline the procedure they proposed is simple. Assume that there exists a function defined over all possible outcomes of some decision problem (say over the range of possible daily energy gains). This function is called a utility function, and it specifies which outcomes are preferred and how much a given outcome is preferred over another. Figure 2.1 shows three simple utility functions: a. a linear utility function, b. a concave-down utility and c. a concave-up utility. The principal differences between these three functions is how their rates of change change (i.e. their second derivatives). In other words, the increase in utility due to adding a single unit of reward to the forager's present states is: (a) the same for all levels of reward for a linear function, (b) decreasing with increasing reward for a concave-down function and (c) increasing with increasing reward for a concave-up function.

Utility theory considers what happens if the reward in these graphs is the outcome of some stochastic process, i.e. described by a probability distribution. Now, given a function $[U(X)]$ and a probability distribution of the variable X $[f(x)]$ it is a straightforward mathematical operation to find the expected (or mean) value of the function

$$E[U(X)|f(x)] = \int_{-\infty}^{\infty} U(x)f(x)dx \quad (2.1)$$

Von Neumann and Morgenstern proposed that probability distributions could be put in rank order of preference by their expected-utilities.

Figure 2.1




Figure 2.1 The figure shows three utility functions with different concavities: (a) is a linear utility function, (b) is a concave-down utility function and (c) is a concave-up utility function.

That is $f_1(x)$ is preferred to $f_2(x)$ if $E[U(X)|f_1(x)] > E[U(X)|f_2(x)]$. In more conventional optimal foraging terms, expected-utility can be used as a currency of maximization when the gains due to foraging are stochastic.

Risk-Sensitivity, Risk-Aversion and Risk-Proneess. For present purposes, the most important result of utility theory is that the three functions of figure 2.1 act very differently ^{under} the expectation operation. To see this, consider a forager offered a choice between a lottery consisting of two possible rewards r_1 and r_2 each with probability one-half, and a certain reward with value $(r_1 + r_2)/2$. This is a choice between a lottery and its mean reward with certainty. A forager with a linear utility function will be indifferent to those two choices; a forager with a concave-down utility function will prefer the certain choice, and a forager with a concave-up utility function will prefer the lottery. These results are illustrated in figure 2.2. Caraco et al. (1980) coined the term "risk-sensitivity" for these types of preferences. In the terminology of economics, a forager preferring a probability distribution to its mean with certainty is said to be risk-prone, and if the certain choice is preferred the forager is said to be risk-averse. The definitions of risk-aversion and risk-proneess are quite specific; they deal only with preference in the case of a single probability distribution vs. its mean with certainty. In later chapters these definitions will prove to be too limiting; then I will usually consider preferences over a range of variances. When rewards are normally distributed, variance is a perfectly adequate "measure" of risk for present purposes, but risk and variance are not equivalent, (this problem has its own body of literature in micro-economics, see Hey 1979, chapter six).

Figure 2.2a




Figure 2.2b




Figure 2.2 This illustrates why concave-up and concave-down utility functions "act" differently under risky conditions. In both (a) and (b) the forager is offered a fair lottery where the same number of units of reward may be lost as may be won, and where the winning and losing both have probability one-half. In the concave-down case the forager might lose more utility than it might win, so a forager with a concave-down utility function would prefer not to gamble. In the concave-up case (b) the forager might win more than it loses, so it prefers to take the gamble.

The Expected-Utility Hypothesis

The most important principle in applying utility theory to the prediction of behaviour is the expected-utility hypothesis (Raiffa 1968). The expected-utility hypothesis (perhaps better described as an assumption) is the supposition that there exists a utility function which ranks probability distributions in order of preference. At its most hollow, this is just the tautology that utility theory works if there exists a utility function which describes the observed behaviour. However, the importance of the expected-utility hypothesis is that conditions can be found for it to hold. Two conditions must hold for the expected utility hypothesis to be true: substitutability and transitivity, (I follow the terminology of Raiffa 1968, to whom the reader is referred for a readable account of the expected-utility hypothesis). Substitutability is the idea that if outcome A is preferred to outcome B, but the forager is indifferent between outcome B and outcome C, then A must be preferred to C. If the forager is indifferent between two outcomes, then it must be possible to substitute them for each other in every type of preference relationship. Transitivity is a similar idea. If A is preferred to B, and B is preferred to C then A must be preferred to C for the expected utility hypothesis to hold. In other words substitutability and transitivity may be thought of as the prerequisites of the expected-utility hypothesis, and consequently the application of utility theory to a given problem.

In economics there are famous examples where humans show preferences inconsistent with transitivity (Allais 1953) and hence inconsistent with the expected-utility hypothesis. Utility theorists reply to this failure by arguing that utility theory is prescriptive

and not descriptive of human behaviour. Utility theory, they argue, tells us how rational human decisions should be made, not how human decisions are made. Foraging theorists clearly do not want to instruct animals in how to make better decisions (although analyses from this point of view might be constructive); instead the goal of foraging theory is to describe the decisions animals make. This means that the postulates of the expected-utility hypothesis deserve some consideration when applying utility theory to foraging behaviour (see Stephens & Charnov 1982). In my discussion of risk (Part I) I adopt an approach which is (indirectly) based on utility theory. Therefore, I acknowledge here that the assumptions of the expected-utility hypothesis are part of the models which I will present.

Currency vs. Expected Utility? Conventionally optimal foraging begins by making an assumption about the currency of optimality (the terminology of Schoener 1971), e.g. maximize the rate of energy or nutrient intake. These currencies have normally been guesses, based on the modeller's knowledge of natural history. These a priori currencies, with their connection to natural history, are one of the strengths of optimal foraging theory. The conventional approach of utility theorist (cf. Raiffa 1968) is to allow the subject to specify its own utility function by giving it a series of preference tests. Both approaches can be useful in the study of foraging behaviour (for example McFarland's, 1977, inverse optimality or in a stochastic case Caraco et al. 1980). However, allowing the forager to specify its own utility function is, in a sense, the least satisfactory of two alternatives, because we want to know not only what the forager's preferences are, but why it has these preferences. Utility theory, per se, sheds no light on this question, only on the form the answer

should take, i.e., a utility function. The approach I adopt is to guess at reasonable utility functions, or stochastic currencies, because I am interested in the question why should a forager change its behaviour in risky situations.


The Deterministic vs. the Risky Case. To help make the distinctions clear consider the form of a deterministic foraging problem. The steps in a deterministic analysis are : (1) choose a currency and (2) specify cost-benefit functions which relate the currency to the behaviour (Schoener 1971). Expected-utility is the same as Schoener's currency because it is the quantity we maximize, but in a risk problem cost-benefit functions do not relate behaviour to expected-utility. They relate behaviour to the distributions of some intermediate reward (e.g. amount of food or nutrients). The expectation of the utility function, given this distribution, provides the final link to the currency.

A utility analysis of foraging behaviour would take the following form. (1) Choose a utility, for example, the probability of survival. All the considerations which apply to the choice in a deterministic currency apply to this decision, e.g. when is the probability of survival the most important attribute of fitness? (2) Specify the relationship between the proposed utility and food. (3) Specify the relationship between behaviour (e.g. patch residence time) and the distribution of food reward. Here, the effect of stochasticity is to require a second type of function; that is, not only must behaviour be related to food gains (Schoener's cost-benefit functions) , but we must relate food gains to preference (utility functions). I illustrate these distinctions in table 2.1.


Although I seldom call upon all of the formal and ponderous machinery of utility theory in my treatment of risk, it is an implicit

TABLE 2.1 SCHEMATIC REPRESENTATION OF
THE STOCHASTIC VS. THE DETERMINISTIC CASES


DETERMINISTIC MODELS:




FOR EXAMPLE:



STOCHASTIC MODELS:



FOR EXAMPLE:



Solid Arrows indicate how behaviour influences reward/currency.
Dashed Arrows the action of selection on behaviour via currency
or expected utility.

part of most of the ideas I present here.

2.3 The Biological Theory of Risk

There have been three approaches to the problem of risk in the literature of behavioural ecology. Firstly, several authors (Oster and Wilson 1978; Caraco 1980, Real 1980 a,b) have proposed that mean maximization may be replaced by maximizing the mean discounted by a certain amount for variability. I call this idea "variance discounting." Secondly, Caraco and his colleagues (Caraco 1980,1981; Caraco et al. 1980) have taken the traditional utility theorist's approach of measuring utility functions. As a part of this approach Caraco (1980) has considered the theoretical implications of arbitrary types of utility functions for optimal patch exploitation behaviour. Finally, many authors have simply proposed that minimizing the probability of death due to starvation may replace mean maximization in some cases (Caraco 1980, McNamara & Houston in press, Stephens 1981 and Stephens & Charnov 1982). The idea that this minimization of the probability of starvation may explain many observed risk phenomena is the hypothesis developed in the next four chapters. In this chapter, I discuss only variance discounting in more detail.

Variance Discounting. If food reward is normally distributed the formal machinery of utility theory may not be necessary. When normality occurs the intuitive idea of a trade-off between a measure of central tendency (mean) and a measure of variability (variance) can be justified by rigorous mathematics. This is because the mean and variance completely specify a normal distribution. This leaves the problem of finding some general way to combine mean and the variance. In the following chapters I propose that the Z-score of statistics may be useful for this purpose. However, several authors have proposed a

straightforward linear combination of mean and standard deviation (Oster & Wilson 1978, Caraco 1980, Real 1980a,b). This currency has the form

$$\max\{\mu - k\sigma\}$$

where μ is the mean and σ is the standard deviation of the reward distribution, and k is a constant measuring the forager's risk-sensitivity. The mean is maximized "discounted" by a certain amount for "undesirable" variability (because k is usually positive). This formulation is recommended by its simplicity, and its claim to generality may be strong as well. Using arguments from utility theory, it has been shown that this formula is approximately correct under very general conditions using a series expansion argument (Oster & Wilson 1978, Real 1980a,b). It is exactly correct if the underlying distribution is normal and the utility is a simple exponential function (Oster and Wilson 1978, Caraco 1980).

The drawback of this currency is that a forager adopting such a policy must always be risk-averse, or always be risk-prone. As I will show in the next chapter some foragers show both types of risk-sensitivity. Although, it may be reasonable to argue that k should change during the course of the forager's life history. This shows that variance discounting, like the arbitrary utility function approach of Caraco (1980), is liable to the same criticism that I have made of utility theory; that is, it lacks the connection to natural history which has been one of the strengths of conventional optimal foraging theory. Specifically, there is no reason to suppose that k is positive, except that (according to Real 1980a,b) animals ought to be risk-averse.

2.4 Summary.

In this chapter I review theoretical ideas about risk. Conventional optimal foraging models have assumed that, in the stochastic case, mean values should be maximized. This implies that foragers should also be indifferent to variance.

Mathematical utility theory deals with preferences over probability distributions. The central idea of utility theory is that a function exists such that probability distributions can be ranked by their expected utilities. Different types of utility functions imply different preferences over variability: a linear utility function implies risk indifference, a concave-down utility function implies risk-aversion, and a concave-up utility function implies risk-proneness. From the point-of-view of optimal foraging theory utility is often applied backwards; because utilities are measured rather than assumed on a priori grounds. However, this is not a necessary part of utility theory, and I do not use it in this way in my study of risk. The use of utility theory pre-supposes that the behaviour in question has certain attributes of rationality.

There have been three approaches to risk in the literature of behavioural ecology. Firstly, it has been proposed that a combination of mean and standard deviation should replace the mean as a currency of maximization (Oster and Wilson 1978, Caraco 1980, Real 1980a,b). I call this variance discounting. Secondly, Caraco (1980) has taken the traditional approach of utility theory by considering the effects of arbitrary utility functions in foraging theory. Finally, it has been proposed that minimizing the probability of death due to starvation may explain observed patterns of risk-sensitivity (Caraco 1980, McNamara and Houston in press, Stephens 1981, Stephens and Charnov in press). This is the hypothesis which I discuss in the remaining chapters of part I.

CHAPTER THREE

Are Animals Risk-Sensitive? A Review of the Evidence

3.1 Introduction

In this chapter I review the evidence concerning animal preferences over risk. This problem has interested workers from two schools of thought, operant psychology and behavioural ecology. The patterns of risk-sensitivity found by these two groups of workers are precisely opposite to one another. Operant psychologists have tended to find that laboratory rats and pigeons are risk-prone, while behavioural ecologists have found that bees, wasps, and great tits are risk-averse.

Despite this difference, all of the work I will discuss has employed the same simple experimental design, which I call the standard risk experiment. The subject is trained to make some measurable dichotomous choice, by using equivalent behaviours (for example choices of keys to peck, or sides of an aviary in which to feed). The animal is then taught that one behaviour yields a reward called C, the certain choice and that the other behaviour yields a reward called L, the lottery choice. The animal learns that if it chooses C then it obtains C units of food with certainty. If it chooses L, it obtains L_1 units with probability one half and L_2 units with probability one half. The magnitudes of C, L_1 and L_2 are always chosen so that $C = (L_1 + L_2)/2$. In the terms of utility theory a forager consistently choosing C can be called risk-averse, and if L is consistently chosen the forager can be called risk-prone; this is because the choice is between a probability distribution, L, and its mean, C with certainty. I illustrate the form of this standard experiment with histograms in figure 3.1. In the following sections I discuss the evidence of the operant psychologists and behavioural

Figure 3.1

STANDARD RISK SENSITIVITY EXPERIMENT




Figure 3.1 This shows the standard design of risk-sensitivity experiments. The forager (subject) is offered a choice between (A) a certain (or constant) alternative and (B) a probabilistically determined alternative with the same mean.

ecologists in turn. Finally, I discuss the crucial evidence for variable risk sensitivity presented by Caraco et al. (1980) and Caraco (1981).

3.2 The operant psychologist's evidence.

Several workers in operant psychology have been interested in risk-sensitivity (e.g. Leventhal et al. 1959, Pubols 1962, Herrnstein 1964, Davison 1969), although they have been interested in risk for very different reasons than those presented here (see Mackintosh 1974 for review and discussion). The results of these experiments are summarized in table 3.1. The data for these experiments have been presented in one of two forms. Herrnstein and Davison have presented their data as matching diagrams (see Mackintosh chapter four) so that risk-proneness is shown because the pigeons spend less time feeding from the C-choice than if both choices had the same variance. Of more interest here, because it is more direct, are the results of Leventhal et al. and Pubols. They presented their results as "learning curves" which show the proportion of response to the lottery choice, L, as a function of individual experience with the choice problem (measured in number of trials). A typical curve from Pubols (1962) is shown in Figure 3.2. Both studies show a general tendency to fix on the lottery choice.

The mean-effect. Leventhal et al and Pubols both found, and were puzzled by, effects of changing the mean $[C=(L_1 + L_2)/2]$. Pubols (1962) varied delay time to get a fixed reward, and Leventhal et al (1959) varied the amount of food reward. Leventhal et al. found that when mean reward size was small the tendency to become risk-prone was unequivocal, but when mean reward was increased they found that risk-proneness developed more slowly, and at the highest mean tested it

TABLE 3.1 SUMMARY OF EMPIRICAL RESULTS ON RISK-SENSITIVITY

STUDY	SPECIES	QUANTITY VARIED	RESULTS
<u>OPERANT PSYCHOLOGISTS</u>			
Leventhal et al. 1958	Lab. Rats	No. of Food Pellets	Risk-prone, but tendency decreased as mean amount was increased.
Pubols 1962	Lab. Rats	Delay Time Before Reward	Risk-prone, but tendency decreased as mean delay was decreased.
Herrnstein 1964	Pigeons	Delay Time	Risk-prone.
Davison 1969	Pigeons	Delay Time	Risk-prone.
<u>BEHAVIOURAL ECOLOGISTS</u>			
Real 1981	Bumblebees & Wasps	Nectar Reward in Artificial Flowers	Risk-averse.
Waddington et al 1981	Bumblebees	Nectar Reward in Artificial Flowers	Risk-averse.
Stephens & Ydenberg Unpub.	Great tits	No. of Fly Pupae	Risk-averse.
Caraco et al. 1980	Yellow-eyed juncos	No. of Millet Seeds	Risk-averse when energy budgets were positive, risk-prone when energy budgets were negative.
Caraco 1981	Dark-eyed juncos	No. of Millet Seeds	As above.

Figure 3.2




Figure 3.2 This figure reproduces Pubols's (1962) Fig. 1. It illustrates a typical result from the operant psychologists' literature of risk. As the animal gains more experience with the choice problem (measured in experimental days on the abscissa), it increasingly chooses the lottery or variable side on the ordinate.

did not develop at all, (figure 3.3a). Pubols found a similar effect, but in the opposite direction; decreasing the mean delay time decreased the tendency to risk proneness (figure 3.3b).

3.3 The behavioural ecologist's evidence.

Very recently behavioural ecologists have become interested in risk-sensitivity, and there have been a number of experimental studies of risk preference (Real 1981, Waddington et al. 1981, Caraco et al. 1980, Caraco 1981, Stephens & Ydenberg unpub. data). These results are summarized in table 3.1. They show (with the exception of Caraco's work which I discuss in a separate section) a clear tendency to risk-aversion. Of course this raises the question of why bumblebees, wasps, and great tits should be so different from rats and pigeons. I will return to this problem in the discussion.

I discuss the results of Stephens and Ydenberg in more detail, because they have not been published elsewhere, and because they are fairly representative. The experimental materials and methods are described in appendix one. We taught each of six great tits to perform a choice experiment by teaching them that food could be obtained from each of two trays in the experimental cage. Initially, they were given equal amounts of food (numbers of fly pupae) on each tray, and tested to see if they had a preference, by giving them simultaneous presentations of the trays. All but one of the birds showed some preference (i.e. made most responses on one side). We then made the preferred side a lottery [yielding seven (L_1) or one (L_2) fly pupae, each with equal probability], and we made the unpreferred side a certain reward yielding four fly pupae. We taught the birds the new schedule by presenting only one tray at a time. After this training the birds were given a second preference test. The results for both

Figure 3.3a




Figure 3.3b




Figure 3.3 This figure represents (a) Leventhal et al's (1959) Fig. 7 and (b) Pubol's (1962) Fig. 2. The figure illustrates what I have called the "mean effect" in the text. In both cases the mean reward offered changes the response to risk. (a) When mean reward size is increased from 0.25 g. food pellets to 1 g. food pellets the risk-proneness does not develop. (b) When mean delay to gain a constant reward size is reduced risk-proneness develops more slowly.

TABLE 3.2 THE RESULTS OF STEPHENS AND YDENBERG

PROPORTION CHOICE OF THE LOTTERY TRAY

Parus major

BIRD NO.	1st PREFERENCE TEST	2nd PREFERENCE TEST
1	.50 (20)	0 (20)
2	.85 (20)	.08 (12)
3	1.0 (20)	.24 (17)
4	.55 (20)	.05 (20)
5	.95 (20)	.30 (20)
6	.55 (20)	.33 (9)

The sample size (no. of choice trials) on which the proportion is based is shown in parenthesis. Significance of a sign test for equality of the preferences test is 0.015.

preference tests are reported in table 3.2. The table shows that every bird reduced its preference for the lottery tray, (its formerly preferred tray). A sign tested conducted on these data showed that the probability that this should have occurred if the birds had not changed their preference is miniscule ($p \approx 0.015$). We concluded that the great tits showed pronounced risk-aversion.

3.4 Caraco et al's & Caracos' Evidence

The recent work of Caraco and his colleagues (Caraco et al. 1980, Caraco 1981) stands out as the most important single piece of empirical evidence in the study of risk. For my purposes the results of Caraco et al. (1980) and Caraco (1981) are identical except that the tests were performed on different species of Junco. I will discuss only Caraco et al's experiment at any length.

Caraco et al. presented yellow-eyed juncos (Junco phaeonotus, a species of granivorous sparrow) with the standard risk experiment. The rewards in the experiment were millet seeds. Most importantly Caraco et al. varied the juncos' energy budgets by manipulating food deprivation and the rate of experimental feeding. They used the standard techniques of "certainty equivalents" to measure the juncos' utility functions (see for example, Raiffa 1968). In figure 3.4a,b I reproduce Caraco et al.'s figures 1 and 2. Figure 3.4a shows a typical utility function for a junco with a positive energy budget, because this utility function is concave-down we can conclude that this bird showed risk-aversion. Similarly, figure 3.4b shows a utility function for a junco with a negative energy budget, which is concave-up and therefore indicates risk-proneness. That is, Caraco and his colleagues have demonstrated the following simple empirical rule: be risk-averse, if your expected daily energy budget is positive; be risk-prone, if your expected daily

Figure 3.4a




Figure 3.4b




Figure 3.4 This figure reproduces Caraco et al.'s (1980) Figs. 1 and 2. These figure show utility functions estimated for the same individual of Junco phaeotus. Reward is the number of millet seeds offered. (a) When the bird's expected energy budget was positive its estimated utility function was concave-down indicating risk-aversion, but (b) When its expected energy budget rule was negative its estimated utility function was concave-up indicating risk-proneness. This result is called the expected energy budget rule.

energy budget is negative. Throughout this thesis I call this the expected energy budget rule.

3.5 Summary and Discussion

The clear empirical result is that foraging animals show both risk-proneness and risk-aversion. It is therefore untenable to hold as Real (1980a,b) does that risk-aversion is universal. I conclude that a reasonable theoretical approach to risk must at least allow for both possibilities.

More importantly Caraco et al.'s expected-energy budget rule shows that (at least in some cases) an individual's risk-sensitivity can be altered by manipulating its energy budget, which suggests another pre-requisite for a theory of risk. However, it may be asked, can the expected energy-budget rule account for all the risk preferences in table 3.1. The short answer is that no conclusion can be reached until more energy budgets have been manipulated. If the expected energy budget rule is taken to be general, it means that risk-aversion should be the norm; since animals must usually (but not always) have positive energy budgets or they would die. This means that the risk-proneness found by the operant psychologists is the most troubling. Could these pigeons and rats have had negative energy budgets? The evidence is suggestive but by no means convincing. Firstly, psychologists confess (Shettleworth, pers. comm.) that it is normal procedure to choose reward sizes to maximize the number of responses per session; this means a lot of work for little reward. Secondly, all of the rats and pigeons were food deprived to 80% of ad. lib. weight. Finally, the mean effect found by Pubols (1962) and Leventhal et al. (1959) is entirely consistent with the expected energy budget rule, because as reward goes up (Leventhal et al. 1959) or mean delay goes down (Pubols 1962) the

energy budget should become more positive. I return to this problem in chapter five.

The starting point of my theoretical work (in the next two chapters) is Caraco's (1980) suggestion that an hypothesis based on minimization of the probability of starvation might account for some aspects of what is known about risk-sensitivity foraging preferences. In the next two chapters I consider this hypothesis in detail, and I show that Caraco's conjecture is right.

CHAPTER FOUR

Z-Score Models of Risk

4.1 Introduction

In this chapter I present simple mathematical models of risk-sensitive foraging behaviour. These models represent the first formal analysis of the hypothesis that animals sometimes forage in a manner so as to minimize their probability of death due to starvation. I show that this hypothesis is sufficient to account for the existence as well as some aspects of the nature of risk-sensitivity reported in chapter three. Specifically, both risk-proneness and risk-aversion are predicted, as is the empirically demonstrated expected energy budget rule

Caraco (1980) was the first to suggest for intuitive reasons that minimizing the probability of starvation may predict risk-sensitivity. The motivation for this suggestion can be seen by considering the simple choice problem illustrated in figure 4.1. The figure shows two lotteries (as histograms), A and B. If the forager chooses lottery A, it gets seven units of food with certainty. If it chooses lottery B it may get three units and it may get 9, each with probability one-half, but it expects to get six units, $[(3+9)/2=6]$. Imagine further that the forager must obtain at least eight units to survive. A conventional (mean-maximizing) optimal forager would choose lottery A (because it has the highest expected value), but this means that it would die with certainty. A non-mean-maximizing forager which chose lottery B would at least have a one-half probability of surviving. The ideas developed here follow more or less directly from this simple example. The conclusion is that when starvation is defined by some threshold food requirement, mean-maximization may not be a very sensible strategy.

Figure 4.1




Figure 4.1 The figure shows a simple example where mean-maximization fails. Imagine that a forager is offered a choice between two lotteries A and B. A always yields seven units of food and B yields three units with probability one-half and nine units with probability one-half. A mean maximizer would choose the certain alternative A. Imagine that at least eight units must be acquired if the forager is to survive. A mean maximizer choosing A would die with certainty, but a forager taking the risky alternative B would have one-half probability of survival. The conclusion is that where there are energetic thresholds of survival, mean maximization may not be a good strategy.

4.2 Z-score models: the basic result.

To illustrate the derivation I consider a small bird foraging in winter. Imagine that at least R units of food reserve are required to survive the night. Let S_0 be the random variable which describes the bird's food reserves when it becomes too dark to continue foraging. The probability of starvation is equivalent to $P(S_0 < R)$. (This assumes that starvation only occurs at night, which may not be true; I relax this assumption in chapter five). Assume that the day's foraging is divided into n intervals. At the beginning of the i th interval the forager makes some decision about how to forage and receives X_i units of food. The X_i 's are random variables, and are assumed to be independently distributed with $E(X_i) = \mu_i$ and $\text{Var}(X_i) = \sigma_i^2$. Assume that exactly r units of food are expended during each interval (this assumes that r is independent of the foraging decision, which cannot be strictly true). Let S_n be the bird's internal food supply when there are n decisions left to make before nightfall. (This amounts to counting time backwards; n is the time left, not the time that has gone by; I adopt this convention throughout Part I). I express the reserves at the end of the day as

$$S_0 = S_n + \sum_{i=1}^n X_i - nr \quad (4.1)$$

the probability of starvation is

$$P(S_0 < R) = P\left(\sum_{i=1}^n X_i < R - S_n + nr\right) \quad (4.2)$$

The central limit theorem for independent random variables states that for n sufficiently large, the random variable

$$T = \frac{\sum_{i=1}^n X_i - \sum_{i=1}^n \mu_i}{\left(\sum_{i=1}^n \sigma_i^2\right)^{1/2}}$$

has a normal distribution with mean zero and variance one. For

convenience I define the following variables

$$\bar{\mu} = \frac{1}{n} \sum_{i=1}^n \mu_i, \text{ the average expected value,}$$

$$\bar{\sigma}^2 = \frac{1}{n} \sum_{i=1}^n \sigma_i^2, \text{ and,}$$

$$\bar{\sigma} = \sqrt{\bar{\sigma}^2}$$

the average variance and its square root. Also I define

$$z = \frac{(R - S_n) - n(\bar{\mu} - r)}{\bar{\sigma} \sqrt{n}} \quad (4.3)$$

It follows from the central limit theorem and appropriate substitutions in equation (4.2) that

$$P(\text{Starving}) = P(S_0 < R) = P(T < z) = \Phi(z) \quad (4.4)$$

where $\Phi(z)$ is the cumulative distribution function (cdf) of a standard normal distribution, $N(0,1)$. By definition of the cdf of any continuous random variable, $\Phi(z)$ increases in a strictly monotonic way with z . This means that as z increases, the probability of starvation increases; minimizing z will minimize the probability of starvation. The effect of variance (a measure of risk) can be seen by differentiating z with respect to σ

$$\frac{\partial z}{\partial \sigma} = \frac{1}{\sigma^2} \frac{(n(\mu - r) - (R - S_n))}{\sqrt{n}} \quad (4.5)$$

the sign of $\partial z / \partial \sigma$ is the same as the sign of the term $n(\mu - r) - (R - S_n)$.

that is,

$$\partial z / \partial \sigma > 0 \text{ where } \bar{\mu} - r > (R - S_n)/n \quad (4.6a)$$

$$\partial z / \partial \sigma < 0 \text{ where } \bar{\mu} - r < (R - S_n)/n \quad (4.6b)$$

These are asymptotic (large n) conditions for risk-sensitivity.