

Expression (4.6a) is the condition for risk-aversion, because the probability of starvation increases with increasing variance. Expression (4.6b) is the condition for risk-proneness, because the probability of starvation decreases with increasing variance. Most importantly this is an analytical derivation of the empirically demonstrated expected-energy budget rule. This result also shows the following for a forager minimizing its probability of death due to starvation. (a) Indifference to variance is unlikely, because it should occur only in the unlikely situation where expression (4.6) becomes an exact equality. (b) The direction of the change in the probability of starvation with variance is independent of the magnitude of the variance. Figure 4.2 presents a graphical argument which illustrates this derivation.

Expression (4.6) may be unnecessarily strict. If n is large,

$\lim_{n \rightarrow \infty} (R - S_n)/n = 0$, and these conditions become:

Choose low variance if $\pi > r$, (4.7a)

and choose high variance if $\pi < r$. (4.7b)

When many decisions are left to make, the forager may base its risk-sensitivity on a simpler rule-of-thumb. Expression (4.7) says that if environmental conditions are such that the forager's reserves are expected to increase, the forager should be risk averse. If reserves are expected to decrease, the forager should be risk-prone.

This analytical derivation is important, because it makes clear both the limitations and implications of the probability of starvation hypothesis of risk-sensitivity. Most importantly, the expected-energy budget rule holds because the total energy acquired is normally distributed. This is reasonable because it will usually be the sum of many random variables. However, if there are few random variables in

Figure 4.2

A. Risk Aversion Predicted

B. Risk Proneness Predicted




Figure 4.2 The figure illustrates the Z-score derivation of the extreme-variance rule. Both A and B show two normal distributions of daily energy gain with the same mean and different variances. LV= low variance; HV= high variance. $R-S_n$ is the net daily energy requirement. Consider A first. When the net energy required is less than the expected value (mean) then it can be seen from the shading under the curves (which represent the probability of starvation) that choosing low variance (LV) yields the lowest probability of starvation. In case B a complementary argument shows that when the net requirement is greater than the mean (expected value) then the variance choice (HV) yields the lowest probability of starvation.

the sum or it is non-normal for other reasons, this rule will be too simple. Secondly, a forager's risk-sensitivity may be influenced by the number of decisions left to make, n . Expression (4.6) suggests that a forager with only a few decisions to make will become risk-prone at a higher reserve level than a forager with many decisions left to make. That is, risk-sensitivity should have a sequential component (see chapter five). This is a characteristic of many procedures in statistical decision theory, (DeGroot 1970).

The extreme-variance rule.

This result is more than an analytical derivation of the expected energy budget rule. It is the derivation of a much more specific result which I call the extreme-variance rule. The extreme-variance rule considers the case where the forager is offered a range of variances at the same mean value. In this case, expression (4.6) says:

$$\begin{aligned} \text{Choose minimum variance if } \bar{\mu} - r &> (R - S_n)/n \\ \text{Choose maximum variance if } \bar{\mu} - r &< (R - S_n)/n \end{aligned} \quad (4.7)$$

In chapter two I emphasized that risk-proneness and risk-aversion considered only the choice between a probability distribution and its mean with certainty. It should be clear, therefore, that the risk-proneness and risk-aversion predicted by the expected energy budget rule are special cases of the extreme variance rule.

4.3 The generalized Z-score model.

Reconsider the critical definition of z (expression (4.3)) in the previous section:

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

Consider the elements of this definition separately. $R - S_n$ is the net amount of energy which the forager must gain in the remaining n

decisions to survive the night. I define $R_{\text{net}} = R - S_n$. $n(\mu - r)$ is the expected net energy gain in the remaining n intervals. I call this $\mu_\tau = n(\mu - r)$, where τ denotes the time remaining before nightfall. $\sigma/\sqrt{n}$ is the standard deviation of energy gain in time τ , $\sigma_\tau = \sigma/\sqrt{n}$. Using this new notation z is

$$z = \frac{R_{\text{net}} - \mu_\tau}{\sigma_\tau} \quad (4.9)$$

which is the familiar form of the z -score (or standard normal variate) of the statistics of the normal distribution. In this notation the extreme variance is simply:

Choose minimum variance if $\mu_\tau > R_{\text{net}}$

and choose maximum variance if $\mu_\tau < R_{\text{net}}$

This simpler form, shorn of its biological trappings, explains the name, Z -score models. It also turns out to be the most useful form for extending the model to the more general case.

The extreme variance rule assumes that the forager is allowed to choose between a range of variances at the same mean. While I know of one natural case in which such a choice may arise (Pulliam and Millikan in press), this is probably a fairly unnatural situation. In natural foraging problems the mean and variance will often be related.

The Z -score (4.8) can be used to solve this problem. The general approach is to specify sets of mean-standard deviation pairs which give the same probability of starvation. If z and R_{net} are fixed all (μ_τ, σ_τ) pairs which satisfy

$$z = \frac{R_{\text{net}} - \mu_\tau}{\sigma_\tau}$$

give the same probability of starvation. By re-arranging this expression, I find that all such points lie on the line

$$\mu_{\tau} = R_{\text{net}} - \sigma_{\tau} z \quad (4.10)$$

In other words, a forager minimizing its probability of starvation due to an energetic short-fall would be indifferent between any point on the line specified by equation 4.10, (where R_{net} and z are fixed). These lines have slope $-z$ and intercept R_{net} on the μ_{τ} axis.

This suggests the simple graphical method of solution illustrated in figure 4.3. In general the forager will be able to choose from a limited feasible set of mean-standard deviation pairs, which is illustrated by the hatched blob in the figure. Consider the family of lines, each of which intersects the μ_{τ} axis at R_{net} . The line of greatest slope which can be drawn through the feasible set gives the smallest possible probability of starvation, and the point, or points, $(\mu_{\tau}, \sigma_{\tau})$ of the feasible through which this line passes is (are) the optimal mean-standards deviation pair (or pairs). The line of greatest slope gives the solution, since when the slope $-z$, is largest, the value Z (which I want to minimize) is smallest.

In a particular optimality problem the modeller would have calculated the feasible set by considering the effects of varying some behavioural parameter, e.g. patch residence time, or the probability of pursuit upon encounter. These calculations would allow the modeller to work backward from the optimal $(\mu_{\tau}, \sigma_{\tau})$ pair to find the optimal behavioural parameter. This procedure is illustrated by a worked example in the next section.

Figure 4.4a shows that the extreme variance rule follows as a special case of graphical solution, because the feasible set is a horizontal line (i.e. μ_{τ} is constant). The figure shows that when $\mu_{\tau} > R_{\text{net}}$ the line of greatest slope passes through the smallest possible σ_{τ} . In the feasible set (condition(4.9a)). When

Figure 4.3




Figure 4.3 The figure illustrates the general graphical solution for optimal preferences from complex sets of mean and variance. The hatched shape, marked F, represents an arbitrary feasible set of mean-standard deviation pairs. The pair which minimizes the probability of death due to an energetic short-fall is found by constructing the line originating at R_{net} (the net energy requirement) on the mean axis and which passes through, or touches, the feasible set with the greatest slope possible.

Figure 4.4




Figure 4.4 A shows two horizontal feasible sets, i.e., sets in which the forager chooses from a range of variances (standard deviations in the graph) at a constant mean. This illustrates that the extreme variance rule, discussed in the previous section, is a special case of this graphical model. Crosses indicate the optimal choices. B shows an example of the general rules for linear feasible sets, i.e., when $\mu_T = m\sigma_T + b$ $\sigma_{min} < \sigma_T < \sigma_{max}$. The optimal preference depends only on the position of the intercept b relative to the net requirement, R_{net} . Note that in the lowest feasible set the extreme variance rule does not hold because all possible means are less than the net requirement, R_{net} , but the optimal policy is still to minimize variance (and consequently to maximize the mean).

$\mu_T < R_{\text{net}}$ the line of greatest slope passes through the largest possible μ_T in the feasible set, (condition (4.9b)).

If the feasible set is linear, but not horizontal, there is a simple extension of this rule. Assume that the feasible set is defined by the linear equation $\mu_T = m\sigma_T + b$, where $\sigma_{\min} \leq \sigma_T \leq \sigma_{\max}$. It is simple to show that the probability of starvation due to a short-fall is minimized by: choosing $\sigma_{\max}$ if $R_{\text{net}} > b$, and choosing $\sigma_{\min}$ if $R_{\text{net}} < b$. The direction of preference for variance is independent of the slope of the feasible set (m). Consequently, the optimal mean value may be either the maximum possible mean or the minimum possible mean in either of the two cases above (the case where $R_{\text{net}} < b$ is illustrated in Fig.4.4b). This figure also makes clear the importance of mean-variance relationships to the extreme-variance rule. In the lower feasible set in 4.4b, all possible means are less than the minimum survival requirement (R_{net}), but the optimal choice is the maximum mean and minimum variance. Linearity is only the simplest kind of non-trivial relationship which might exist between μ_T and σ_T . In the next section I show that natural μ_T/σ_T relationships may be much more complicated.

The importance of mean-variance relationships suggests that there will be some feasible sets for which mean-maximization is preserved. That is, cases in which an animal minimizing its probability of death due to a short-fall should act exactly like a mean-maximizer. In fact, it has been suggested that such feasible sets will be the most common type in nature (G.H. Pyke personal communication). I can offer no evidence on this point, but this formulation does allow me to specify the nature of such feasible sets. If $\mu_{\max}$ is the maximum mean, and this occurs at standard deviation $\sigma_{\mu_{\max}}$, then mean maximization

will be preserved as long as all remaining pairs in the feasible set lie below a line drawn between R_{net} on the μ_τ axis ($\mu_{\text{max}}, \sigma_{\mu_{\text{max}}}$), put algebraically, mean maximization is preserved as long as

$$\frac{\mu_\tau - R_{\text{net}}}{\sigma_\tau} < \frac{\mu_{\text{max}} - R_{\text{net}}}{\sigma_{\mu_{\text{max}}}} \quad (4.11)$$

for all μ_τ and σ_τ in the feasible set.

4.4 Patch use and the risk of starvation: a worked example

An important question raised in the last section is what mean-variance relationship (feasible sets) are most common and most general. In this section I apply the results of the previous sections to finding the optimal patch residence time for an animal minimizing its probability of starvation. This illustrates how the relationship between mean and variance may be found in a foraging problem.

The model Consider an animal foraging in discrete patches in a manner which minimizes its probability of starvation. Assume that time spent in a patch completely specifies the net amount of food acquired from the patch. Patches are identical, and numerous enough that overall environmental depletion due to the forager's behaviour is not significant; i.e. the forager is always entering undepleted patches. The only stochasticity is due to random variation in the time required to find a new patch.

The problem is to find a patch residence time, t , which minimizes the probability of starvation, as discussed above. Let T_s be the random variable which represents the time between leaving one patch and entering another. Assume that T_s is exponentially distributed with parameter λ ; that is patches are encountered at random. Let t be the patch residence time (which is the parameter over which I wish to

optimize), and let $g(t)$ be the negatively accelerated gain function (i. e., a deterministic function which allows patch depression) which relates the net amount of energy gained from a patch to the residence time. I now require equations for the mean and variance of such a foraging process. The required expressions have been presented recently by Stephens and Charnov (in press, appendix III). I use Stephens and Charnov's equations (1.1) and (1.2) to find the mean and variance of energy gain over a fixed long time τ as functions of patch residence time:

$$\mu_{\tau}(t) = \frac{g(t)\tau}{(1/\lambda) + t} \quad (4.12)$$

$$\sigma_{\tau}^2 = \frac{\tau(1/\lambda^2) [g^2(t)]}{[(1/\lambda) + t]} \quad (4.13)$$

It is a standard result of patch-use theory that μ_{τ} is maximized when

$$g'(t) = \frac{g(t)}{(1/\lambda) + t} \quad (4.14)$$

This result is a special case of Charnov's (1976) marginal value theorem. Consider the derivative of the standard deviation with respect to residence time,

$$\frac{\partial \sigma}{\partial t} = \sqrt{\tau} \left\{ \frac{g'(t)[(1/\lambda) + t]^{3/2} - (3/2)g(t)[(1/\lambda) + t]^{1/2}}{\lambda[(1/\lambda) + t]^3} \right\}$$

Since the term in the denominator is always positive the sign of the derivative is determined by the difference in the numerator. The standard deviation increases with t as long as

$$g'(t) > \frac{(3/2)g(t)}{(1/\lambda) + t} \quad (4.15)$$

and the standard deviation decreases with t when this inequality is reversed; it is therefore at a maximum when t is chosen such that

$$g'(t) = \frac{(3/2)g(t)}{(1/\lambda) + t} \quad (4.16)$$

Notice the similarity between (4.14) and (4.16). I deduce from this that the residence time at which the maximum standard deviation is attained is always less than the residence time at which the mean intake is maximized. The results of this analysis may be summarized as follows: (i) μ_τ and σ_τ are both unimodal functions of patch residence time, and (ii) the maximum standard deviation will always occur at a smaller residence time than the maximum mean and (iii) both μ_τ and σ_τ are zero when $t = 0$. These points are enough to determine the shape of the feasible set in μ_τ - σ_τ space and use the graphical solution of the previous section to make qualitative deductions about the patch-exploitation behaviour of a starvation minimizer.

Figure 4.5a shows the general shape of the feasible set. Figure 4.5b illustrates the qualitative deductions which can be made considering this shape, and the graphical method of solution presented in Section 4.4. The optimal residence time for a forager minimizing its probability of starvation will have the following properties: (1) It will usually be near the residence time predicted by mean-maximization (the marginal value theorem) because of the curve's oblong shape. (2) There will be a critical level of requirement, R^*_{net} , at which the optimal residence time is exactly the same as in the mean-maximization case. (3) If the actual requirement is less than R^*_{net} , then the optimal behaviour is to stay longer than predicted by mean-maximization. (4) If the actual requirement is greater than R^*_{net} , the optimal residence time is shorter than predicted by

Figure 4.5




Figure 4.5 Patch residence time and the risk of starvation. The curve shown in A shows the shape of the feasible set for an animal foraging in discrete patches, and choosing an optimal patch residence time as discussed in the text. The patch residence time (t_i 's) marked on the curve indicate how increasing residence time affects the mean and standard deviation, that is $t_4 > t^* > t_3 > t_2 > t_1$. B application of the graphical model of the previous section to finding optimal patch residence times. R_{net}^* is the net requirement at which mean maximization is preserved. But, as with R and as with $R_{\text{net}4}$ if the requirement is less than R_{net}^* it is best to stay longer (i.e., $t_4 > t^*$); and as with $R_{\text{net}3}$, if the requirement is greater than R_{net}^* it is best to leave sooner (i.e. $t^* > t_3$).

mean-maximization. (5) As the requirement becomes larger the optimal residence time becomes shorter and approaches the time at which the maximum standard deviation is attained, but will never become smaller than this.

As a general rule we may say that if the requirement is low compared to expectations, then the forager should stay longer than expected by mean maximization; if the requirement is high the forager should stay for a shorter period than expected by mean-maximization, but never shorter than the time producing maximum variance. The extreme-variance rule does not hold in this example, because neither the highest nor the lowest possible variance should be chosen, except in the limiting case. The extreme-variance rule does give some insight into why this result occurs. The forager chooses increasing higher variance as requirements become larger relative to expected gains. However, since the feasible set is not linear, the forager should never choose extremes of variance because it would have to accept a disproportionate loss in mean to do so.

Several authors have presented modifications of the marginal value theorem. Oaten (1977) and McNamara (1982) consider the a case in which the forager also gains information while foraging. McNair (in press) considers a modification where the forager is allowed to encounter new patches while searching a patch. In both of these cases the optimal residence time is longer than predicted by the marginal value theorem, (although the comparison with Oaten's model requires a redefinition of the Charnov's marginal capture rate, cf. McNamara 1982). McNair (in press) also analyses an optimal residence time model which minimizes the probability of starvation over an infinite time period. His model differs from mine in that he envisions no critical reckoning time

(nightfall in my example). McNair's 'starvation' is simply running out of fuel during the foraging process. McNair finds that the probability of starvation is minimized by staying longer than the marginal value theorem would predict. McNair's infinite time model is similar to mine when net requirements are very small relative to expected gains. My model suggests that over-staying may be the most common case; because requirements will normally be much less than expected gains. However, in the case of severe energetic stress relatively short stays may be optimal.

4.5 Discussion and Summary

In this chapter I present a theoretical analysis of risk-sensitive foraging behaviour based on the hypothesis of minimization of the probability of death due to starvation. I show that this hypothesis can account for both risk-prone and risk-averse foraging preferences. In particular, my derivation shows that it is not generally true that "a risk-prone forager accepts a greater probability of starvation in attempting to acquire a greater net energetic intake"(Caraco 1980). My model shows that there are conditions i.e. (4.7b) , where the forager may decrease its probability of starvation by risk-prone behaviour.

Perhaps most importantly I show that these Z-score models account for the empirically demonstrated expected energy budget rule of Caraco et al. (1980) and Caraco (1981).

In section 4.3 I argue that a useful theory of risk-sensitivity will have to go beyond the strict definitions of risk-proneness and risk-aversion. Such a theory will need to consider optimal choice from complex sets of mean-variance pairs; because the mean and variance that an animal may choose via its foraging behaviour will often be related. In section 4.3, I develop such a model based on the Z-score for the

risk of starvation hypothesis. The model shows that the form of the feasible set of mean-variance pairs is critical to a forager's preferences over variance.

In section 4.4 I present a worked example of the application of these results to a conventional foraging problem: choice of an optimal patch residence. The surprising shape of feasible sets I find for this problem suggests that it may be misleading to try to find general mean-variance relationships; because no straightforward functional relationship would have predicted this feasible set, (e.g. $\mu = k\sigma^2$ might seem a reasonable general relationship, but it is very far from what I find). This result suggests that the relationship between mean and variance will depend critically on the nature of the foraging problem at hand. For example, the feasible set may have been very different if the forager had been assumed to adjust the way it searched within a patch rather than simply the time spent. The shape of the feasible set found in this example means that the solutions will usually be near the solution for mean-maximization. It is tempting to conclude that this might be a general phenomenon but I think this conclusion should await further theoretical work.

Some Limitations and Difficulties. This model is no closer to an explanation of the troublesome risk-proneness found by operant psychologists (reviewed in chapter three) than the expected energy budget rule is in itself; since this is principally a mathematical justification of the expected-energy budget rule. Moreover, these experimental findings are doubly troublesome because there are many alternative explanations for risk-aversion (Real 1980a,b, Caraco 1980) but none for risk-proneness. As a first step towards solving this conundrum, I examine the dubious assumptions I have made here in the next chapter.

It should be pointed out that these models are not intended to replace mean-maximization, but to be a complementary part of a cohesive theory of optimal foraging. For example, the assumptions I have made may make sense for a small bird foraging in winter, but they will not make sense for the same bird in the breeding season. That is, the conditions in which the probability of survival may be expected to be the most important attribute of fitness are limited.

CHAPTER FIVE

Objections to the Theory

5.1 Introduction.

In this chapter I consider three objections to the Z-score models of risk-sensitivity discussed in chapter four. In other words, I relax three simplifying assumptions and consider their relative importance. Firstly, although I have considered models which purport to minimize the probability of death due to starvation, I have only "allowed" one kind of starvation, that is overnight starvation in the small bird example of chapter four. It is also possible that starvation could occur while foraging; the bird might simply run out of fuel in mid-day. In mathematical terms this may affect my estimate of the probability of starvation, because it restricts the number of admissible paths to overnight survival (i.e. the state where $S_0 > R$); because there is a lower absorbing boundary (death is a boundary which cannot be re-crossed). I call this the "lazy-L" problem because I am imagining two absorbing boundaries which define death due to starvation, the first is a horizontal line (mid-day death) followed by a vertical line of height R (the daily energy requirement); the death boundary is therefore the shape of a reclining L, see figure 5.1.

Secondly, I have assumed in chapter four, that an animal makes only one decision about whether to choose high or low variance. However, I have shown that an animal's food supply should affect this decision, since the food supply is a random variable it is possible that it might change in way which would make it suboptimal to continue the present policy. Suppose that $\mu < R - S_n$ at time n . Then it is optimal to

choose the highest variance, $\sum_{i=1}^n \sigma^2_i$, possible. However, what if the

Figure 5.1 The Lazy-L.




Figure 5.1 The figure shows the absorbing boundaries which represent death in the Lazy-L model. Time left in Day is on the abscissa and energy reserves are on the ordinate. The changing "state" of the forager's energy reserves with time would be a patch in this space, if it touches any of the hatched boundaries death occurs.

gain, X_n , at time n , is so large that $(n-1)\mu > R - (S_n + X_n)$? By the same logic which lead to choosing high variance at time n , the animal should choose low variance at time $n-1$. Since at time n , the highest possible variance for all remaining decisions was chosen, the optimal behaviour at time n contradicts optimal behaviour at time $n-1$, in this example. McNamara (in press) and Houston and McNamara (in press) have used this example to point out that animals ought to be permitted, in the most general case, to change their minds about their risk-sensitivity; in mathematical terms a full analysis must be sequential.

Finally, I have assumed that all amounts of energy gained over the requirement are equivalent; that is, it is no better to survive by having ten times as much as you need than to survive by having barely more than you need. There may be some cases in which this assumption is reasonable, but there must be many more in which it is not. I call this the problem of carry-over, because we need to know how energy "carried-over" to the next day is valued to find out how significantly this affects our results.

Throughout this chapter I am principally concerned with the effect of these changes on the extreme variance rule: Choose $\sigma_{\max}$ if requirements exceed expected gains, choose $\sigma_{\min}$ if expected gains exceed requirements. It is important to bear two features of this rule in mind. Firstly, it deals only with the case where the forager is supposed to choose optimal variance at a fixed mean; and this is the case treated in this chapter. Secondly, recall (as pointed out in chapter two) that to choose extreme variance is not the same as to be risk-prone or risk-averse; since these terms assume that the animal is offered a distribution vs. its mean with certainty. The extreme

variance rule deals with the more biologically realistic choice from a range of non-zero variances.

5.2 The Lazy-L: Death in the afternoon.

Starvation while foraging, as opposed to starvation by falling short, is similar to a class of mathematical and economic problems usually called ruin problems. A firm is ruined when it has zero capital, in the same sense that a forager must die if it has no energy reserves. Of interest here, as a motivating example, is a problem called the "insurance company's ruin," the principle result of which is given by Cox and Miller (1965, pp. 60-61) and which I cite without proof. This result gives the probability of eventual death due to starvation as a function of initial energy supply, (S_∞ , where ∞ denotes that nightfall is an infinite time away), the mean energy change per interval, $\mu_{\text{net}} = \mu - r$ (assume r is constant as chapter four), and σ^2 is the variance in energy gain per interval. Then

$$P(\text{Eventual Starvation}) = \begin{cases} 1, & \mu \leq r \\ \exp[-2S_\infty(\mu-r)/\sigma^2], & \mu > r \end{cases} \quad (5.1)$$

It is straightforward to show (by differentiation) that an animal minimizing its probability of starvation should choose low variance if it expects its energy reserves to increase ($\mu > r$) and indifferent to variance if it expects its supply to decrease. These results hold regardless of the initial energy supply (S_∞). The result, in the increasing reserves case, is exactly the same as the infinite time condition for minimizing the probability of a short-fall, (eqn 4.7). At first, it seems surprising that when reserves are declining that indifference to variance should occur. However, this is because of the infinite time assumption; risk-prone behaviour might prolong

survival, but starvation is nonetheless inevitable when reserves are expected to decrease ($\mu < r$). Biologically, it would be surprising if animals did not choose high variance when $\mu < r$. Therefore, it appears that the eventual starvation case is of fairly limited biological interest, although McNair (in press) has used it in a model of optimal patch residence time (see chapter four). We will need to consider the finite time case to make more appealing biological models of the lazy-L problem. In summary, we would not expect the lazy-L to make much difference to the extreme variance rule when expected gains exceed requirements, since low variance minimizes the probabilities of both being ruined and of falling short.

A Discrete Analogue. To investigate the lazy-L in more detail I consider a discrete analogue of the continuous case. Imagine that random energy gains are obtained at discrete intervals, as in the standard results of the chapter three. Say that there are n decision intervals left in the day, and that X_i represents the energy gain in interval i . In this discrete case I make a special assumption about the probability distribution of the X_i 's. Assume that $P(X_i = \Delta) = p$ and $P(X_i = -\Delta) = 1-p$, where $\Delta > 0$. For obvious reasons I call Δ the step-size. Note that while a given realization of X_i is simply "a step," $E(X_i)$ is the mean step-size.

The unit step case: The mathematical results I use to investigate the lazy-L are from the theory of random walks (see Feller 1950, Cox and Miller 1965). These results are most accessible when the case where there is step size of one ($\Delta=1$) and an arbitrary probability of a positive step, p , is considered (probability of a negative step is $1-p$). By appropriate re-scaling the results of the unit step case can be extended to any step-size. I modify the notation of chapter four

but try to maintain some similarity to enhance the comparison.

Let: s = the energy store at time n ($\approx S_n$).
 n = the number of intervals left in the day ($=n$).
 ρ = the minimum number of units of energy required
 to survive the night, ($\approx R$).

All of these variables must be integers in the discrete case.

The probability of survival: To find the probability of survival I begin by considering possible paths between the point (n,s) and an arbitrary point $(0,x)$ as shown in figure 5.2. The figure shows two paths. A forager taking path A would survive, while a forager taking path B would not. Both paths have 6 upward moves and 4 downward moves, as would any path between these two points, [i.e. between $(10,3)$ and $(0,5)$]. Given that only six upward moves occur any of the 210 (the binomial coefficient, 10 pick 6) paths between $(10,3)$ and $(0,5)$ are equally likely. This is the most important assertion of our model; that is, in a unrestricted unit-step random walk any of the possible paths between the points (n,s) and $(0,x)$ are equally likely. This means that our problem has become one of specifying the proportion of these paths which do not touch or fall below the time axis (i.e. that do not result in death in the afternoon).

There is a simple rule which can be used to find the number of paths which do touch or cross the time axis (i.e. that result in death). I simply state it here. The reflection principle: the number of paths from (n,s) to $(0,x)$ which do not touch or cross the time axis is exactly equal to the number of unrestricted paths from $(n,-s)$ to $(0,x)$. $(n,-s)$ is the reflection of (n,s) about the time axis. (See Feller 1950 for discussion and proof.) As stated earlier any path between the same two points has the same number of upwards and

Figure 5.2 Two Paths originating and ending at the same points.




Figure 5.2 The figure shows two paths from the point (10,3) [i.e. 10 time units left and 3 energy units in reserve]. to the point (0,5) [i.e. 5 units at "nightfall"]. A forager taking path A would survive until nightfall, but a forager taking path B would have died 7 units before nightfall. Notice also that both paths require the same number of upward and downward steps.

downwards moves. Let

$$u = \frac{x - s + n}{2} . \quad (5.2)$$

If u is not an integer then there is no possible path between (n,s) and $(0,x)$, but if u is an integer u = the number of upwards moves required for a path between (n,s) and $(0,x)$. The number of paths between (n,s) and $(0,x)$ is $\binom{n}{u}$ if u is an integer and 0 otherwise. By the same argument the number of upward moves required to get from the reflection point $(n,-s)$ to $(0,x)$ is

$$u' = \frac{x + s + n}{2} \quad (5.3)$$

or, $u' = s + u,$

and the number of paths from (n,s) to $(0,x)$ which result in death is given by $\binom{n}{s+u}$, if $s+u$ is an integer less than n , and 0 otherwise.

The probability of surviving the trip from (n,s) to $(0,x)$ is

$$P(\text{Surviving the Trip} | s, x) = \frac{\binom{n}{u} - \binom{n}{s+u}}{\binom{n}{u}} \quad (5.4)$$

where the binomial coefficients are interpreted as zero if u or $s+u$ are not integers. It is a simple matter to specify the total probability of survival.

$$P(\text{Survival}) = \sum_{u_{\min}}^n \frac{\binom{n}{u} - \binom{n}{s+u}}{\binom{n}{u}} \binom{n}{u} p^u (1-p)^{n-u} \quad (5.5a)$$

$$= \sum_{u_{\min}}^n \binom{n}{u} p^u (1-p)^{n-u} - \sum_{u_{\min}}^n \binom{n}{s+u} p^u (1-p)^{n-u} \quad (5.5b)$$

where

$$u_{\min} = \begin{cases} \frac{\rho - s + n}{2}, & \text{if } \frac{\rho - s + n}{2} \text{ is an integer} \\ \frac{\rho - s + n + 1}{2}, & \text{otherwise} \end{cases} \quad (5.6)$$

$u_{\min}$ is the minimum number of upward moves required to survive.

Now consider the right hand term in (5.5b). By changing the indices of summation I find that

$$\sum_{u=u_{\min}}^n \binom{n}{s+u} p^u (1-p)^{n-u} = \frac{(1-p)^s}{p^s} \sum_{i=s+u_{\min}}^n \binom{n}{i} p^i (1-p)^{n-i} \quad (5.7)$$

This gives the probability of starvation in the equivalent form:

$$\sum_{u=u_{\min}}^n \binom{n}{u} p^u (1-p)^{n-u} - \frac{(1-p)^s}{p^s} \sum_{i=s+u_{\min}}^n \binom{n}{i} p^i (1-p)^{n-i} \quad (5.8)$$

The advantage of this form is that it gives the probability of survival in terms of the cumulative distribution function (cdf) of the binomial. For n large, (numerical studies indicate that in this case n should be larger than about 30) the binomial cdf is well approximated by the normal cdf, i.e.

$$P(\text{Survival}) \approx 1 - \Phi\left(\frac{u_{\min} - np}{\sqrt{np(1-p)}}\right) - \frac{(1-p)^s}{p^s} \left\{1 - \Phi\left(\frac{s + u_{\min} - np}{\sqrt{np(1-p)}}\right)\right\} \quad (5.9)$$

These equations {(5.5), (5.8), (5.9)} give the probability of survival in the unit step case.

I use this result in two ways. Elsewhere (Stephens 1981) I have suggested that if the energy reserves are sufficiently large then the lazy-L will make little difference. I use equations (5.8) and (5.9) to make this claim more explicit. That is, I specify exactly how large reserves must be for the lazy-L to have no effect. Secondly, I analyse the effect of the lazy-L on the expected energy budget rule.

Conditions for the Lazy-L to have no effect. Notice that equation (5.8) is in two parts. The first part is simply the complement of the cumulative distribution function of the binomial. This term represents the total probability that an unrestricted path will lead to overnight survival (i.e. to a state where $S_0 \geq \rho$). The second term is the

probability that the forager arrives at the state $S_0 \geq \rho$ via a path which dips below or touches the time axis. In more biological terms, it is the likelihood that a forager arrives at a level of energy reserves which enables it to survive the night but does so by a route which requires that it dies in mid-day. If this second term is zero, then the Z-score (normal) approximation methods proposed in chapter four are unaffected by the lazy-L; because the first term is readily approximated by the normal.

This leads to the conclusion, for the unit step case, that the lazy-L has no effect as long as

$$s + u_{\min} > n. \quad (5.10)$$

That is, if the lower index of summation in the second term of (5.8) is greater than the upper index it must equal zero. By substituting the definition of $u_{\min}$ (5.6) into this expression (5.10), the condition for the lazy-L to have no effect becomes

$$s + \rho > n. \quad (5.11)$$

In words, if the sum of initial energy supply and the minimum energy requirement is greater than the number of trials (*time) remaining then the lazy-L has no effect, in the unit step-size case. To illustrate this three cases at different requirement levels are shown in figure 5.3. The figure illustrates that as the requirement becomes large the region of initial conditions over which the lazy-L makes no difference becomes larger. As general results, I may state that the lazy-L makes no difference when (a) initial energy reserves are sufficiently large, (b) requirements are sufficiently large, and (c) when the time remaining in the day is sufficiently small.

That large energy reserves should be important is intuitively clear, but the importance of large requirements and a small number

Figure 5.3 Regions where the Lazy-L has no effect. The Effect of ρ .




Figure 5.3 The shaded areas of these graphs show conditions (combination of time left on the abscissa, and energy reserve on the ordinate) when the Lazy-L has no effect. ρ is the discrete version of daily energy requirement, A,B,C is a series of increasing ρ . as ρ increases the size of the "no effect" region increases.

of trials remaining may be more surprising. For example, in figure 5.3c an initial state of $n=3$ and $s=3$ is in the region of conditions unaffected by the lazy-L. But, this state is unaffected not because the forager is unlikely to die before the end of the day, but because no path which might lead to overnight survival (in this case to $S_0 \geq 6$) also leads to death before nightfall, (i.e. dips below the time axis). It is worthwhile to state this condition more precisely: if no path which leads to overnight survival ($S_0 \geq \rho$) dips below, or touches, the time axis, the lazy-L makes no difference. This simply says that the lazy-L can only make a difference if overnight survival and mid-day starvation are jointly possible. This logic allows me to specify a more general condition.

The first bounce rule: a more general case. The unit step-size case, is obviously too simple to be of much general interest. However, the condition (5.11) does suggest a simple extension to a more reasonable continuous case. Imagine that at each interval the step X_i has the same distribution and that this distribution is bounded, specifically, I assume that $-\Delta_{\min} \leq X_i \leq \Delta_{\max}$. Assume that there are n steps left in day, the forager starts with energy supply S_n and requires R units of energy to survive the night. (This is the notation of chapter four.) The idea gained from the last section was that if no path from (n, S_n) to $(0, S_0)$ where $S_0 \geq R$ also dips below or touches the time axis then the lazy-L has no effect. It is easy to find, knowing S_n and $-\Delta_{\min}$, the first possible dip below the time axis. It is just the first interval after the axis is crossed if all the steps up to "now" have been $-\Delta_{\min}$. Let the first interval after the first crossing be k^* intervals away from the end of the day. In general at k^* , reserves (due to the previous $n-k^*$ steps of $-\Delta_{\min}$)

will be on or below the time axis at $S_{k^*} = S_n - (n-k^*)\Delta_{\min}$. Since this position (k^*, S_{k^*}) represents the earliest possible dip below the time axis, it follows that if the remaining k^* intervals all increase the reserves by $\Delta_{\max}$ the resulting S_0^* (reserves at nightfall) will be the largest possible dusk reserve level which also dips below the time axis. Consequently, if this S_0^* is less than the requirement, it follows that the lazy-L has no effect, since no path which leads to overnight survival also leads to mid-day starvation. The logic just described suggests the name the "first bounce rule;" because we consider the earliest and most extreme possible deflection off the time axis.

In figure 5.4 I present an analysis comparable to that in figure 5.3 for this more general case, that is I illustrate the effect of differing requirements. To do this I apply the logic of the first bounce rule "backwards." Starting at the overnight requirement I draw a line of slope $\Delta_{\max}$ until it crosses the time axis, and from this point of first crossing I extend a line of slope $-\Delta_{\min}$. This delineates the region of initial conditions for which the lazy-L has no effect. Note that all the qualitative effects of the unit step analysis are preserved: increasing S_n , increasing R and decreasing n all make the lazy-L less important. More interestingly, figure 5.5 suggests that decreases in the absolute sizes of both the minimum and maximum step-sizes ($\Delta_{\min}$ and $\Delta_{\max}$) will decrease the importance of the lazy-L. This leads to the conclusion that as the range of X_i increases that the lazy-L will become more important. Since the range is a measure of variability this suggests the tantalizing, and appealing idea that as variability in the distribution of food reward increases the lazy-L is likely to become more important.

Figure 5.4 The first bounce rule: Generality of the Effects of Increasing Requirement.




Figure 5.4 The figure illustrates the use of the first bounce rule to find conditions where the Lazy-L makes no difference . The case illustrated is when the minimum step-size $\Delta_{\min} = -1/2$ and the maximum step-size $\Delta_{\max} = 2$. The first bounce rule considers the first and most extreme possible deflection off the lower absorbing boundary. In both cases (a and b) a line of slope $\Delta_{\max}$ is constructed (in time unit steps) until the lower boundary is touched or crossed. A line of slope $\Delta_{\min}$ from; this point delineates the "no effect" region. Notice that increasing the net requirement increases the size of the "no effect" region, just as in the unit step case.

Figure 5.5 The First Bounce Rule: the Effect of Varying Extreme Step Sizes.




Figure 5.5 The figure shows the "no effect" region for three conditions of extreme step sizes (A) $\Delta_{\min} = -1, \Delta_{\max} = 1$; (B) $\Delta_{\min} = -1/2, \Delta_{\max} = 1$; (C) $\Delta_{\min} = -1, \Delta_{\max} = 1/2$. The conclusion is that decreases of the absolute value the extremes increase the size of the "no effect" region.

Although I have shown that there are conditions of large requirement, large reserves, and short time when the lazy-L doesn't have any effect there are two general cases where it will always have an effect which are worth mentioning. Firstly, if the time remaining (n) is large (recall that I assumed that n was large in the results of chapter four), then the lazy-L will normally make at least some difference. Secondly, if the gain at each trial is drawn from an unbounded distribution (i.e. if $-\Delta_{\min} = -\infty$ and $\Delta_{\max} = +\infty$) then the lazy-L always has some effect. This is because even a path extending far below the starvation boundary (the time axis) has a remote possibility of ending up with $S_0 \geq R$ (the condition for overnight survival), because an infinite positive step is possible.

Up to now, my analysis has been fairly simple. That is, I have only considered the all or none issue, will the lazy-L reduce the the probability of survival at all? A more complete analysis will ask how the lazy-L affects the forager's risk-sensitivity.

The lazy-L and the extreme-variance rule. The results from the first bounce rule suggest that as the range ($\sim$ variance) of stepsizes increases so does the importance of the lazy-L. Secondly, notice that from equations (5.5 to 5.9) above that the lazy-L, when it has an effect, will decrease the probability of survival. This, in turn suggests that as variance increases the probability of survival will decrease. It would seem reasonable to expect the lazy-L to have most effect on the condition of the extreme-variance rule under which preferences for high variance are predicted (as the ruin problems discussed above suggest). That is, one would expect that if $\mu_{\tau} > R_{\text{net}}$, low variance should be chosen, as in the unmodified results of chapter four; but when $\mu_{\tau} < R_{\text{net}}$ one would expect that

although higher variance should be chosen that there might be some limit to how high a variance is best in terms of the probability of survival set by the increasing probability of death in mid-day when variance increases. In other words, higher but not necessarily the highest variance should be chosen.

To illustrate this I analyse the initial conditions illustrated in table 5.1a. I consider 4 possible step-sizes (Δ 's) $\Delta = 1, 2, 5, 10$ at the same mean step-size. In my discussion of the unit step-size case, I said that any condition of s , ρ , n and Δ could be considered in terms of an equivalent unit step model. In table 5.1a I show the equivalent unit step-size case for each of the four stepsizes listed above. Recall that the distribution of step-sizes is given by $P(X_i = \Delta) = p$ and $P(X_i = -\Delta) = 1-p$, if we choose $p = (\mu + \Delta)/2\Delta$ then, at a fixed mean (μ), Δ is the only parameter of the distribution of the X_i 's, and $E(X_i) = \mu$ for all Δ , but $V(X_i) = \Delta^2 - \mu^2$. That is, variance increases monotonically with Δ . This distribution is convenient because we can change the variance without changing the mean by altering the step-size (Δ), and we can restrict our attention to the binomially distributed number of positive, or negative steps, made in the n remaining trials, in other words we may use the analysis of the unit step-size case.

In table 5.1b, I show the exact probabilities of survival for all four step-size (variance) conditions at ten different mean step-sizes. I also show the probabilities of survival calculated by Z-score approximation. The variance choice according to the exact, (lazy-L) calculation at each mean treatment is marked by an L, and the optimal choice according to the Z-approximation is marked by a Z. The starred (*) mean value ($\mu = 0.387$) indicates the threshold mean as

TABLE 5.1a
Equivalent unit step parameters for initial conditions:
R=21.6 $S_n=10$ n=30

	$\Delta = 1$	$\Delta = 2$	$\Delta = 5$	$\Delta = 10$
s =	10	5	2	1
$\rho =$	22	11	5	3
n =	30	30	30	30

TABLE 5.1b
Probabilities of survival at differing means and ranges

		Increasing Variance			
		$\Delta=1$	$\Delta=2$	$\Delta=5$	$\Delta=10$
$\mu = 1$	Exact=	1.00L	.976	.697	.234
	Z-Approx=	1.00Z	.974	.753	.632
$\mu = .75$	Exact=	.997L	.885	.611	.207
	Z-Approx=	.999Z	.858	.656	.579
$\mu = .50$	Exact=	.803L	.679	.517	.181
	Z-Approx=	.763Z	.626	.550	.525
$\mu = .387^*$	Exact=	.558L	.558	.558	.558
	Z-Approx=	.501Z	.500	.500	.500
$\mu = .25$	Exact=	.258	.407	.421L	.158
	Z-Approx=	.220	.353	.440	.470Z
$\mu = 0.00$	Exact=	.021	.178	.328L	.135
	Z-Approx=	.017	.145	.336	.416Z
$\mu = -.25$	Exact=	≈ 0	.053	.242L	.115
	Z-Approx=	≈ 0	.039	.242	.364Z
$\mu = -.50$	Exact=	≈ 0	.010	.170L	.096
	Z-Approx=	≈ 0	.006	.165	.226Z
$\mu = -.75$	Exact=	≈ 0	.001	.112L	.080
	Z-Approx=	≈ 0	≈ 0	.104	.266Z
$\mu = -1.0$	Exact=	0.0	≈ 0	.069L	.065
	Z-Approx=	≈ 0	≈ 0	.061	.223Z

* indicates the threshold mean; L following a number shows the optimal choice for the lazy-L; Z following a number marks the optimum for the Z-approximation.

predicted by the Z-score model, i.e. the mean above which risk-aversion is predicted and below which risk-proneness is predicted.

Table 5.1b shows that my expectations are closely met. Firstly, notice that the Z-score approximation is very good, except in the highest variance case ($\Delta = 10$). Secondly, at mean values less than the threshold requirement, $\mu < 0.387$, the optimal choice is to choose the second highest variance ($\Delta = 5$) and not the highest as predicted by the Z-score model.

Conclusions for the Lazy-L. In this section I find conditions in which the problem of dying in mid-day has no effect on the predictions of risk-sensitivity of the Z-score models discussed earlier. Specifically, I show, as one would expect, that when reserves are high the lazy-L has little effect. More surprisingly, I show that when requirements are large and when time remaining is short that the lazy-L has less effect. If the time remaining is short and the requirements are great one might say that the forager is in a situation of potential energy stress. Looked at in this way, the conditions which make the lazy-L less important coincide with our earlier assertions about when Z-score models may be appropriate; i.e. in conditions of energy stress when there is a premium on just surviving. This result also suggested that as range (and consequently variance) of the gain (X_i) distribution increases that the probability of survival will be decreased by the lazy-L effect. When the lazy-L does have an effect, the extreme variance rule will have to be modified to include the choice of moderately high, but not the highest, variance when expected requirements are greater than expected gains.

5.3 The value of changing your mind.

As I pointed out in the introduction of this chapter, the simple Z-score model of risk-sensitivity is a one decision model. That is,

given conditions of time remaining (n), current energy reserves (S_n), and the distribution of energy gain per interval, the animal is supposed to choose all remaining variances for the rest of the day. McNamara (in press and MS) and Houston and McNamara (in press) have shown that when the forager is allowed to change its mind, certain aspects of my model must be modified. I wish to emphasize that this section is, with one exception, simply a recitation of the important results of J. M. McNamara and A. I. Houston. However, no discussion of the objections to Z-score models would be complete without mentioning and explaining this work.

A simple model. To see the importance of changing your mind consider the following simple choice problem. At any interval a forager may choose between:

C. A certain choice where $P(X_i = \mu) = 1$, and

L. A lottery choice where $P(X_i = 0) = P(X_i = 2\mu) = 1/2$

These two alternatives have the same expected value $\{E(X_i) = \mu\}$ and different variances $\{V(X_i|C) = 0, V(X_i|L) = \mu^2\}$. I consider the optimal choices in two cases, when only one decision remains to be made in the day, and when two decisions remain, (the $n=1$ and $n=2$ cases). Optimal behaviour when a single decision remains is simple:

$$\begin{array}{ll} \text{Choose C,} & \text{if } R_1 \leq \mu \\ \text{Choose L,} & \text{if } R_1 > \mu \end{array} \quad (5.12)$$

Where R_i denotes the net requirement when i decisions remain in the day. The rule given above specifies a cost function over R_1 , (cost = the probability of death given R_1 and optimal behaviour at choice $n=1$). The cost function is

$$P_1(\text{Death}|R_1) = \begin{cases} 0, & R_1 \leq \mu \\ 1/2, & \mu < R_1 < 2\mu \\ 1, & 2\mu \leq R_1 \end{cases} \quad (5.13)$$

where P_1 indicates that this is the $n=1$ case. Now I use the idea of backward induction (see DeGroot 1970) to find the optimal behaviour when two decisions remain. That is, since the best procedure at $n=1$ is known, I know how the gains resulting from any decision taken at $n=2$ may best be used, consequently we know the value (in terms of the probability of death) of a decision at $n=2$.

To find the optimal behaviour at $n=2$ I consider the value in terms of ultimate survival (i.e. in terms of [5.13]) of choosing L or C. $R_1 = R_2 - X_2$ where X_2 is determined by the choice of C or L at $n=2$. First consider the value of choosing L. There are two cases $X_2 = 0$ and $X_2 = 2\mu$. Figure 5.6 illustrates the averaging of these two cases to obtain a cost function for the choice of L at $n=2$. This function is:

$$P_2(\text{Death}|R_2, L) = \begin{cases} 0, & R_2 < \mu \\ 1/4, & \mu < R_2 < 2\mu \\ 1/2, & 2\mu < R_2 < 3\mu \\ 3/4, & 3\mu < R_2 < 4\mu \\ 1, & 4\mu < R_2. \end{cases} \quad (5.14)$$

as shown in figure 5.6c. In the case where C is chosen I simply substitute $R_1 = R_2 - \mu$ into (5.13) to find

$$P_2(\text{Death}|R_2, C) = \begin{cases} 0, & R_2 < 2\mu \\ 1/2, & 2\mu < R_2 < 3\mu \\ 1, & 3\mu < R_2. \end{cases} \quad (5.15)$$

As one would expect, by the expected energy budget rule, the optimal behaviour when two decisions remain is to choose C if $R_2 \leq 2\mu$ and choose L if $R_2 > 2\mu$. It is interesting to notice that when $2\mu \leq R_2 \leq 3\mu$ that L and C provide the same probability of death. This is because under these conditions, if the forager chooses C now, it must choose L later. However, it seems reasonable that when faced with a choice of taking a risk now or later that the risk should be taken now. This slightly arbitrary assertion allowed, the first result

Figure 5.6 Calculating the Pay-off due to a risky decision when mind-changing is allowed.




Figure 5.6 The figure illustrates the pay-off averaging required to find the probability of death due to a risky decision (L_2) when there are two decisions left to make. (A) shows the probability of death as a function of current requirements if choice L_2 yields nothing (which it does with probability one-half). (B) shows the probability of death if L_2 yields a gain of 2μ . (C) shows the average of (A) and (B).

of this analysis is that the rule for the sequential ("mind changing") case is the same as for the one decision, non-sequential, model. This was also found in a much more general way by McNamara (MS and in press) and Houston and McNamara (in press), for the case in which the high and low variance options have the same mean.

Notice that in both cases ($n=1$ and $n=2$) that the optimal rule divides the requirement set (i.e. the set of all possible R_i 's) into risk-prone and risk-averse regions. In other words, there is a critical R_i^* value such that if $R_i > R_i^*$ then risk-aversion is optimal and if $R_i < R_i^*$ then risk-proneness is optimal. We can compare the pay-offs [$P(\text{Death})$] for the risk-prone regions in the two ($n=1$ and $n=2$) cases by considering a normalized risk-prone region $R_N = R_i - R_i^* = R_i - i\mu$. In figure 5.7a I plot the probabilities of survival due to risk-prone decisions as a function of R_N for both the $n=1$ and $n=2$ cases. Notice that the probability of survival due^a to risk-prone decision when $n=2$ is always greater than or equal to the probability of survival due to a risk-prone decision for $n=1$. Roughly speaking this says that a risk-prone decision taken far away from "dusk" is safer than a risk-prone decision taken close to "dusk."

Secondly, in the one decision case the extreme variance rule would yield:

$$P(\text{Death}|\text{Fixed choices}) = \begin{cases} 0, & R_2 \leq 2\mu \\ 3/4, & 2\mu < R_2 < 4\mu \\ 1, & 4\mu < R_2. \end{cases} \quad (5.16)$$

Compared with the sequential decision case:

$$P(\text{Death}|\text{Mind Changing}) = \begin{cases} 0, & R_2 < \mu \\ 1/4, & \mu < R_2 < 2\mu \\ 1/2, & 2\mu < R_2 < 3\mu \\ 3/4, & 3\mu < R_2 < 4\mu \\ 1, & 4\mu < R_2 \end{cases} \quad (5.17)$$

Figure 5.7b plots these functions simultaneously to show that

Figures 5.7a and 5.7b. The Probability of Survival due to Lottery Choice.




Figure 5.7 (A) The figure shows the probability of survival as a function of normalized requirement due to a Lottery. It shows that a lottery choice where there is another decision to make ($n=2$) has the same or higher probability of survival than a lottery choice when no decisions are left to make ($n=1$). (B) The figure shows a comparison of a fixed choice strategy when two intervals remain with a mind-changing strategy. The mind-changing strategy has the same or a lower probability of survival.

$P(\text{Death}|R_2)$ in the sequential case is always less than or equal to the $P(\text{Death}|R_2)$ in the non-sequential, one decision, case. This has also been shown in a more general case by McNamara (MS and in press) and Houston and McNamara (in press).

Conclusions for the fickle forager. In summary this says that although the rules for the sequential and non-sequential cases are very similar the pay-offs ~~may~~^{be} different. Specifically the probability of survival due to taking a "high variance" decision may be much higher than predicted by Z-score models. This is because Z-score models ignore the possibility that a high variance choice may produce a gain large enough to put the forager in a position where it can switch to low variance.

McNamara (MS) has shown that in the pure risk-prone vs. risk-averse case (i.e. the animal chooses either $\sigma=0$ or $\sigma=\sigma_1$, $\sigma_1 > 0$) that the probability of survival due to single risk-prone decision is twice the value that would be estimated from the Z-score method. This suggests that the increase in the probability of survival due to a "high variance" decision in the sequential case will be most important near the threshold requirement level (R_{j^*}) and less important as the actual requirement becomes less than R_{j^*} . Table 5.2 shows some typical probabilities of survival calculated using McNamara's (MS) sequential equations for the time (∞) large case and the Z-score approximation. Note that the difference is greatest near the threshold requirement.

One difficulty with using the probability of starvation to explain the risk-sensitive preferences shown by foraging animals (see Chapter two), is that by applying the logic of the Z-score models one would expect at least one-half of all animals choosing risk-prone decisions

Table 5.2 Probabilities of Survival when High Variances should be Chosen: The Value of Changing Your Mind

Distance below Requirement	One Decision Case $\Phi(Z)$	Many Decision Cases	
	Choices: Only σ is chosen	0 or σ	$\sigma/2$ or σ
0	.5	1.0	0.667
1σ	.1587	0.3174	0.2116
2σ	.0228	0.0456	0.0304
3σ	.0013	0.0026	0.0017

to die of starvation. However, McNamara's and Houston's work (McNamara MS and in press and Houston and McNamara in press) shows that the probability of survival may actually be much higher than one half, especially if the requirement is not much higher than the threshold requirement.

5.4 Carry-over.

In the Z-score models of chapter four I assumed that a forager should maximize its probability of acquiring a fixed amount of energy by a given time (e.g. nightfall). Putting this assumption in the general framework of utility theory, this is the same as saying that the forager has the simple "step" utility function illustrated in figure 5.8a. In this case all amounts of energy gained above the requirement (R) are of equal value, (or utility). That is, I assume there is no value to large amount of energy "carried-over" to the next day. This may be reasonable in the case of a small over-wintering bird. However, consider the value of differing amounts of food for a female insect larva about to pupate. Imagine an insect which must pupate at some fixed time of year. If the insect requires a minimum amount to survive pupation, then it has a similar problem to a small bird trying to survive the night. However, larger reserves at pupation may allow earlier emergence, or larger egg production; carry-over will be especially important in those insects in which adults feed very little, or not at all (e.g. may-flies and ant-lions).

In a simple-minded, but fairly general way, the step utility function of the Z-score models can be converted to consider the problem of carry-over. I illustrate several cases of such modified step functions in figure 5.8b. These modified step functions are made up of three pieces; an unchanged portion below the requirement (R) where

Figure 5.8 Three types of Utility Functions.




Figure 5.8 The figure shows three types of utility functions. (A) A simple step-function with no carry-over [R =requirement], V_s = value of survival]. (A) is the case of my standard result, (B) illustrates a number of possible modifications with positive (increasing) carry-over. (C) illustrates the step-linear utility when food in excess of the requirement increases utility linearly ; this is the case studied in section 5.6.

utility is zero, a vertical line of height V_S , representing the value of just surviving, at $S_0 = R$, and an increasing function of S_0 extending upwards from the point (R, V_S) . This function $U_C(S_0 - R)$ represents how energy carried-over is valued in terms of fitness (or utility). I call this the carry-over function; and I define it so that $U_C(0) = 0$, that is, $U_C(S_0 - R) + V_S$ is the utility of survival at supply level S_0 , given that $S_0 \geq R$. Let $f(s_0)$ be the distribution of the random variable S_0 , energy supply at pupation (or nightfall), and let $F(s_0) = P(S_0 < s_0)$ be the cumulative distribution function associated with $f(s_0)$. Then expected utility can be given quite generally as

$$\begin{aligned} \text{or, } E\{U(S_0)\} &= 0 \cdot F(R) + \int_R^{\infty} \{U_C(s_0 - R) + V_S\} f(s_0) ds_0 \\ &= \int_R^{\infty} U_C(s_0 - R) f(s_0) ds_0 + \{1 - F(R)\} V_S \end{aligned} \quad (5.18)$$

This equation could be used to solve any particular carry-over problem; however, here I wish to use it to consider how carry-over affects the extreme-variance rule. To do this I consider a simple example, making specific assumptions about the forms of $U_C(S_0 - R)$ and $f(s_0)$, with surprising results.

The achieved mean effect. I consider a case where the mean of S_0 (μ) is held fixed, and ask whether the extreme-variance rule of chapter four (Choose $\sigma_{\max}$ if $\mu < R$, Choose $\sigma_{\min}$ if $\mu > R$) is the optimal behaviour when carry-over is important. The answer is surprising. Firstly, as one would expect in all cases when expected gains are less than requirements ($\mu < R$) $\sigma_{\max}$ should be chosen. However, in general, when expected gains are greater than requirements

it is not true that $\sigma_{\min}$ should always be chosen; in fact I show that there are conditions when the highest possible variance should be chosen regardless of expected requirements.

To show this I make specific assumptions about the carry-over function $U_c(S_0 - R)$ and the probability distribution function of total energy gain $f(s_0)$. Assume that the carry-over function is linear, i.e.

$$U_c(S_0 - R) = \gamma(S_0 - R), \quad (5.19)$$

where γ is a constant. The form of the utility function resulting from this assumption is shown in figure 5.8c. Also assume that the probability distribution of food gain is given by the double exponential (or Laplace) distribution:

$$f(s_0) = \frac{1}{2\beta} \exp\left\{-\frac{|s_0 - \mu|}{\beta}\right\} \quad (5.20)$$

where $E(S_0) = \mu$, $V(S_0) = 2\beta^2$, i.e. $\sigma = (\sqrt{2})\beta$. Given these assumptions it is straightforward to find the terms required to specify the expected utility:

$$1 - F(R) = \begin{cases} \frac{1}{2} \exp\left\{\frac{\mu - R}{\beta}\right\}, & \text{I. } R > \mu \\ 1 - \frac{1}{2} \exp\left\{\frac{R - \mu}{\beta}\right\} & \text{II. } \mu \geq R \end{cases} \quad (5.21)$$

$$\int_R^{\infty} s_0 f(s_0) ds_0 = \begin{cases} \frac{R + \beta}{2} \exp\left\{\frac{R - \mu}{\beta}\right\} & \text{I. } R > \mu \\ \frac{R - \beta}{2} \exp\left\{\frac{\mu - R}{\beta}\right\} & \text{II. } \mu \geq R \end{cases} \quad (5.22)$$

when U_c is linear eqn (5.18) becomes:

$$E\{U(S_0)\} = \gamma \int_R^{\infty} s_0 f(s_0) ds_0 + (V_S - \gamma R)\{1 - F(R)\} \quad (5.23)$$

I choose the double exponential distribution instead of the more reasonable normal because, unlike the normal, it is readily integrated

and hence used in eqn (5.23). It is also convenient because it allows me to manipulate the mean and variance independently; since one of its parameters (β) is monotonically related to variance.

Case 1. $\mu < R$. This example breaks conveniently into two cases of interest. Firstly, consider the case where expected reserves are less than requirements, ($\mu < R$). The expected utility is found by substitution and re-arrangement of eqns (5.21), (5.22) and (5.23):

$$E\{U(S_0)\} = \frac{1}{2} \exp\left\{\frac{\mu - R}{\beta}\right\} \{V_S + \gamma\beta\}. \quad (5.24)$$

To see how variability changes expected utility I differentiate with respect to β :

$$\frac{\partial E[U(S_0)]}{\partial \beta} = \frac{1}{2} \exp\left\{\frac{\mu - R}{\beta}\right\} \left\{\gamma + \frac{R - \mu}{\beta^2} (V_S + \gamma\beta)\right\} \quad (5.25)$$

since $R > \mu$ this expression is always positive. Therefore, when requirements are greater than expected gains the highest possible variance should be chosen, as predicted by the extreme-variance rule.

Case 2. $\mu \geq R$. Secondly, I consider the case in which expected gains are greater than or equal to requirements. The expected utility is found by substitution and some algebra to be:

$$E[U(S_0)] = \gamma(\mu - R) + V_S - \frac{1}{2} \exp\left\{\frac{R - \mu}{\beta}\right\} \{V_S - \gamma\beta\} \quad (5.26)$$

and

$$\frac{\partial E[U(S_0)]}{\partial \beta} = \frac{1}{2} \exp\left\{\frac{R - \mu}{\beta}\right\} \left\{\gamma - \frac{(\mu - R)(V_S - \gamma\beta)}{\beta^2}\right\} \quad (5.27)$$

In figure 5.9 I graph a two examples of $E[U(S_0)]$ as a function of β . The figure shows that expected utility is a concave-up function of β . As β gets large $\exp\{(R - \mu)/\beta\}$ approaches 1, so that expected utility

Figure 5.9 Two plots of expected utility as a function of variability β , the $\mu > R$ case.




Figure 5.9 The figure shows two plots of expected utility as a function of variability β , at two different carry over coefficients $\gamma = 0$ and $\gamma = 10$. The $\gamma = 0$ case shows that the extreme variance rule is a special case of the carry-over model (i.e. no carry-over); because in the $\mu > R$ case expected utility always decreases with variability. The dashed lines are linear asymptotes (eqn. 5.28). β^* is the β at which expected utility is minimized.

asymptotically approaches the line given by:

$$\gamma(\mu - R) + \frac{V_S}{2} + \gamma\beta \quad (5.28)$$

From the concave-up shape of the curve we conclude that when offered a range of β 's at a fixed mean, the forager should choose either $\beta_{\min}$ or $\beta_{\max}$ depending on exactly what range is offered. Surprisingly if $\beta_{\max}$ is sufficiently large it will always be the optimal choice, if $\gamma > 0$, since (5.26) increases to infinity with β .

From the derivative (5.27) I can show that expected utility increases with β if,

$$\gamma \beta^2 + \gamma(\mu - R)\beta - (\mu - R)V_S \quad (5.29)$$

is positive, and decreases with β if (5.29) is negative. Several interesting special cases follow from this condition.

Recall that in the extreme-variance rule the forager should be indifferent to variance if expected gains are exactly equal to requirements. However, in all cases of carry-over (with symmetrical gain distributions) it is optimal to choose the highest possible variance when expected gains equal requirements. This can be seen from condition (5.29) by setting $\mu - R = 0$, in which case (5.29) is always positive (since $\gamma > 0$ when there is carry-over). This is the most extreme example of what I have called the achieved mean effect. The achieved mean effect is responsible for all of the deviations from the extreme variance rule due to carry-over. It is easily seen in the case where expectations equal requirements ($\mu = R$) as shown in figure 5.10. When $\mu = R$ the probability of survival is one half for any symmetric probability distribution, and changing the variance, therefore, cannot change the risk of starvation. However, increased variance will increase the mean of the distribution truncated below R , as shown in

Figure 5.10 The Achieved Mean Effect




Figure 5.10 The figure shows the achieved mean effect in a simple case. Where $\mu = R$, changing variability (β) cannot change the probability of starvation, but can, as the figure shows, can increase the mean for which the forager is "paid". Compare case A (low variability) to case B (high variability). AM indicates "achieved mean".

figure 5.10. In the double exponential case when $\mu = R$ this "achieved" mean increases linearly with β (in fact it is $R + \beta$). I call this the achieved mean effect because it is the mean for which the forager is "paid."

This example suggests a rather counter-intuitive result. If expected gains are larger than, but close to, requirements then high variance is more likely to be the optimal choice than when expected gains are much larger than requirements. This is surprising because intuition would suggest that if μ is only barely larger than R that it is more important to reduce variance than if μ is much greater than R . The answer to this conundrum is that in the carry-over case intuition is only half right. Consider the graph of expected utility as a function of β (variability) shown in figure 5.9. The curve has a minimum at β^* , and β^* is the positive root of the quadratic equation given by (5.29). Notice that the range of β 's offered determines whether $\beta_{\max}$ or $\beta_{\min}$ is the optimal choice. If $\beta_{\max} < \beta^*$ then low variance is always the best choice. However, β^* is an increasing function of $\mu - R$, so that the region in which low variance is optimal actually decreases as μ approaches R . Similarly, if $\beta_{\min} > \beta^*$ high variance ($\beta_{\max}$) is always preferred. This happens because if μ is close to R only very low variances have much effect in lowering the probability of starvation, so that if all the feasible variances are fairly large then the achieved mean effect arbitrates for high variances. To make this clear figure 5.11 shows this effect with uniform distributions.

There are two other interesting effects in this model. Firstly consider the slope of the carry-over function, γ . Notice that when $\gamma = 0$ the utility function is exactly the same as the step utility

Figure 5.11 The effect of distance between μ and R .

Figure 5.11a




Figure 5.11b




Figure 5.11. One counter-intuitive prediction of the carry-over model is that high variance is more likely to be preferred close to the threshold requirement. This figure illustrates this, for uniform distributions with mean μ and maximum $\mu+l$, minimum $\mu-l$, l is a measure of variance. (a) if all feasible l 's are small enough so that $\mu-l > R$, then low variances are preferred. (b) However if only l 's are so large that $\mu-l < R$ are offered then the achieved mean effect makes high variance preferable. The range of l 's for which low-variance is preferred becomes smaller as μ approaches R .

function illustrated in figure 5.8c, and the condition (5.29) is always negative; predicting that low variance should be preferred in all cases where $\mu > R$. That is, this model has the results of the extreme variance rule as a special case. It seems reasonable that increasing γ would enhance the value of the achieved mean, and consequently increase the value of high variances. This is borne out by the graphs shown in figure 5.9, where I show what is otherwise the same function at two values of γ ($=0$, $=10$). The effect is most clearly seen by noticing that γ is the slope of the asymptote given by (5.28). We can conclude that the smaller γ is, the more closely the carry-over case approximates the extreme variance rule.

Secondly, consider the minimum value of survival, V_S . As above it is reasonable to expect that as V_S increases, the achieved mean will increase and consequently, that higher V_S 's will increase the value of high variance. This is not the case; because increasing V_S increases the relative costs of death due to starvation (i.e. it puts more weight on the probability of survival). So that if the minimum value of survival, V_S , is large, low variance becomes more likely. This can be seen from the condition (5.29), when $V_S = 0$ (5.29) is always positive, suggesting that high variance should always be chosen when V_S is sufficiently small.

Summarizing the results for linear carry-over functions and a double exponential distribution of energy gain, I find that the achieved mean is the most important principle in understanding the carry-over problem. When expected gains are less than requirements high variance should always be chosen. But, when gains exceed requirements it is not true that low variance should always be chosen; in this case, when high variance is preferred it is due to the achieved

mean effect. The preference for high variance is increased: when expected gains are close to the requirement, when the rate of increase in the value of energy carried-over (γ) is high, and when the value of surviving (V_S) is low.

Is this true? These results are sufficiently surprising for it to be important to consider whether they are artefacts of the simplifying assumptions I made to obtain them. There are two tenuous assumptions; that gain has a double exponential distribution, and that the carry-over function is linear. It is possible that the achieved mean effect will be enhanced by using the double exponential distribution; because like the simpler negative exponential, its mean has special properties under truncation. Also the assumption of linear carry-over may have implications for risk-sensitivity; since we know from utility theory that functions which are not from linear (i.e. having differing concavities) act differently in conditions of risk.

To analyse the importance of these two assumptions, I have evaluated the expected utility [expression (5.18)] using the trapezoidal rule for numerical integration, assuming that the distribution of gain follows the more general normal distribution. The basic scheme of these calculations is constant; I assume that 25 units of energy are required, $R = 25$. At each of several mean values I calculated the expected utility for 50 different variances $\sigma^2 \in \{2, 4, \dots, 100\}$. In the resulting tables I present the variance at which the largest expected utility was found. I do this for several treatments. Three carry-over functions: (i) linear $[\gamma(S_0 - R)]$, (ii) Concave down $[\gamma\sqrt{S_0 - R}]$ and (iii) concave-up $[\gamma(S_0 - R)^2]$. I also vary γ the carry-over co-efficient and, V_S , the minimum value of survival.

The linear case. In table 5.3 I present the numerical results for linear carry-over functions. In this case I have changed only the probability distribution function of gain from that considered in the analytical results above. I find that when carry-over is linear the results are qualitatively unchanged from those above. Firstly, notice that either maximum or minimum variance is always preferred; intermediate variances are never preferred, in the linear case or in subsequent analyses. Notice in table 5.3a that the rate of change of the value of energy carried over, γ , has the predicted effect: low variance is preferred when γ is small, high variance when γ is large. Table 5.3b shows that the minimum value of survival (V_S) also has the predicted effect: at low V_S , high variance is preferred; and at high V_S , low variance is preferred. I conclude that when carry-over is linear that the double exponential distribution is an excellent model of the more general normal distribution.

The importance of concavity. In tables 5.4 and 5.5 I present calculations for concave-down and concave-up carry-over functions. From the basic results of utility theory we would expect concave-down functions to mediate for low variance and concave-up functions to mediate for high variance. This is exactly what the tables show, so much so that we see no effect of either γ or V_S . However, this may be an effect of the range of parameters tested; this range was chosen mainly to consider the linear case. In any case concavities (up or down) probably reduce the importance of these parameters.

Finally notice that the tables do not show the effect of increased preference for high variance when expected gains are close to (but greater than) requirements. Closer examination of the numerical results shows, as in fact one would expect, that this is an artefact of

Table 5.3 Linear Carry-over: $U_C(S_0 - R) = \gamma(S_0 - R)$

OPTIMAL VARIANCES

from the range 2 to 100

Table 5.3a Effect of γ

Mean	$\gamma = 0$	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	indiff	100	100	100
$\mu = 30$	2	2	100	100
$\mu = 35$	2	2	100	100

Initial Conditions:

$V_S = 10, R = 25$

"indiff" indicates that all variances provide equal expected utilities

Table 5.3b Effect of V_S

	$V_S = 1/2$	$V_S = 1$	$V_S = 10$	$V_S = 50$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	100	100	100	100
$\mu = 30$	100	100	100	2
$\mu = 35$	100	100	100	2

Initial Conditions:

$\gamma = 5, R = 25$

Table 5.3c Effect of $\mu - R$ (Variance range 50 to 100)

Mean:	25(R)	30	35	40
Optimal Variance:	100	100	50	50

Initial Conditions:

$\gamma = 5, R = 25, V_S = 50$

Table 5.4 Concave-down Carry-over: $U_C(S_0 - R) = \gamma \sqrt{(S_0 - R)}$

OPTIMAL VARIANCES

from the range 2 to 100

Table 5.4a Effect of γ

Mean	$\gamma = 0$	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	indiff	100	100	100
$\mu = 30$	2	2	2	2
$\mu = 35$	2	2	2	2

Initial Conditions:

$V_S = 10, R = 25$

"indiff" indicates
that all variances
provide equal
expected utilities

Table 5.4b Effect of V_S

	$V_S = 1/2$	$V_S = 1$	$V_S = 10$	$V_S = 50$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	100	100	100	100
$\mu = 30$	2	2	2	2
$\mu = 35$	2	2	2	2

Initial Conditions:

$\gamma = 5, R = 25$

Table 5.5 Concave-up Carry-over: $U_C(S_0 - R) = \gamma(S_0 - R)^2$

OPTIMAL VARIANCES

from the range 2 to 100

Table 5.5a Effect of γ

Mean	$\gamma = 0$	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	indiff	100	100	100
$\mu = 30$	2	100	100	100
$\mu = 35$	2	100	100	100

Initial Conditions:

$V_S = 10, R = 25$

"indiff" indicates
that all variances
provide equal
expected utilities

Table 5.5b Effect of V_S

	$V_S = 1/2$	$V_S = 1$	$V_S = 10$	$V_S = 50$
$\mu = 15$	100	100	100	100
$\mu = 20$	100	100	100	100
$\mu = 25(R)$	100	100	100	100
$\mu = 30$	100	100	100	100
$\mu = 35$	100	100	100	100

Initial Conditions:

$\gamma = 5, R = 25$

the range of variances offered. Table 5.3c shows that if the range of variances is reduced to 50 to 100, this effect of proximity is easily seen. More importantly, notice that when expected gains equal requirements that high variance should be chosen in all cases of positive carry-over ($\gamma > 0$) regardless of the shape of the carry-over function. This is a product of the achieved mean effect which I expect to be quite general.

Discussion and Conclusions--Carry-over. To what extent does the achieved mean effect mean that we will have to modify the extreme-variance rule? The numerical results presented above suggest that this will depend on how much the carry-over function departs from linearity. The most reasonable case is where the value of energy carried-over increases at a decreasing rate, or is concave down. When this is true the achieved mean effect may change things rather little. However, the idea of linear carry-over functions should not be thrown away too quickly. All mean-maximizing models of optimal foraging tacitly assume linear utility functions, and these are not without successes. Secondly, it is the case which provides the most insight into the conditions which enhance or diminish the achieved mean effect.

The analytical model presented has three important parameters, (γ , V_S , and $\mu - R$). All of which might be found in a real animal example by considering aspects of the forager's natural history and environment. Here I discuss the implications of natural history for γ , the coefficient of the carry-over function; in a sense γ is a direct measure of the importance of carry-over. Many things may increase the importance of carry-over. For example, the proximity (in time) of the breeding season, and short life expectancies may increase γ . We might expect that animals of differing life-histories, or the same animal at

different times of the year, or stages of life may have different γ 's. To choose a specific example, the larval insect mentioned earlier is likely to be more of a risk-taker (i.e. have high γ) than the small over-wintering bird of chapter four. The idea that such (high γ) risk-takers may exist presents a significant challenge to the ideas of Real (1980a, 1980b) that animals ought always to avoid risk.

Throughout my discussion of risk (Part I) I have presented the 'probability of starvation' models simply as counter-examples to the larger literature of mean-maximization. This leaves the problem of finding a more general currency of maximization, which takes into account both the importance of the risk of death as well as the more general problems of gaining energy quickly and efficiently as dealt with by mean-maximization models. The results of this section present a means of doing this (i.e. eqn 5.18); which can be recommended as fairly natural formulation. McNair (in press) has suggested that risk of starvation and mean maximizing currencies might be combined by maximizing the product of the probability of survival and the mean energy gain. However, McNair's currency is an ad hoc one, which does not subsume the empirically supported energy budget rule in the way that this currency does. Although it would need further development and generalization, the technique presented here is promising step towards solving the combination problem set above.

In this section I have considered fairly vague and general shapes for the carry-over function. In chapter two, I criticised this "arbitrary" utility function approach as used by Caraco (1980) and Real (1980a, 1980b), because the utilities they proposed have little biological meaning. However, I feel that this approach is valuable as a technique to examine the assumptions of the more biologically

meaningful models of chapter four. That is, I have not considered what energy carried-over is used for in this section; I have simply asserted that it is valuable.

McNamara and Houston (in press) have considered a particular use for energy carried-over from day-to-day. They consider a small bird feeding in winter, and assume that fitness is maximized by maximizing the probability of over-winter survival. Further, they assume as I do in chapter four, that a fixed amount of energy is required to survive each day. The amount in excess of this requirement can then be valued in terms of its contribution to the total probability of over-winter survival. Using simple and reasonable assumptions, they find numerically the probability of surviving over the next k days as a function of reserves at the end of the present day. That is, they find utility functions given that "winter" ends in the next k days. They find that after the first few days, these functions change very little with increasing k . This says that if a forager maximizes its probability of surviving over a long period, it may do just as well to only consider maximizing its probability of survival over a much shorter period. This result seems to lend support to the analysis of chapter four, which maximizes the probability of surviving over a given short critical period. Of course, this is a distinct problem from the carry-over problem discussed above, in which food has a value other than for survival, since McNamara and Houston assume that survival is the only attribute of fitness.

Perhaps the most exciting implication of carry-over is that it might be used to explain the apparently anomalous risk-proneness reported by operant psychologists (see chapter three). It is

especially suggestive that operant subjects are normally maintained at deprivation levels close to their minimum requirement; a condition which corresponds to the choice of high variance in the carry-over model discussed above. Nothing more conclusive can be said until more energy budgets are measured in operant studies, and until the value of energy carried-over is better understood; what is the advantage of being at 95% of ad lib. weight instead of 80%? The model suggests some experiments. If the risk-proneness observed is due to carry-over, it should be possible to change it to risk-aversion by presenting appropriate mean-variance combinations, as discussed above.

5.5 Summary Discussion--Objections to the Theory.

I have examined the importance of three changes in the simplifying assumptions of the Z-score models of risk-sensitivity presented in chapter four. In each case I find that some modifications must be made, but that there are fairly generous and reasonable conditions in which the results of chapter four are preserved. In the first section, I consider how the possibility of starvation while foraging (as opposed to starvation at the end of the day) affects the extreme variance rule; I call this problem the lazy-L. The main effect is to limit the size of variance that should be chosen when expected gains are less than requirements; an intermediate, rather than the highest, variance should be chosen. In the second section I consider the effect of allowing the forager to change its mind about risk-sensitivity, following the work of McNamara and Houston (McNamara MS and in press, and Houston and McNamara in press). It is found that the extreme variance rule is not changed, but that the probability of survival due to acting according to the extreme variance (applied sequentially) is increased. In the third section I assume that energy in excess of the amount required for

survival has some positive value. If this value increases linearly, there are startling changes in the extreme variance rule, due to what I have called the achieved mean effect. However, when the value of energy carried-over increases more slowly than linearly, (i.e. it is a concave-down function; as will often be reasonable) deviations in the extreme variance rule due to the achieved mean effect will be small. I conclude that the Z-score models are generally robust, and moreover, that those changes that have been found are themselves interesting phenomena.

Many problems are left unsolved. The lazy-L could be extended to the continuous step-size case, possibly resulting in a more appealing formulation. When decisions are made sequentially there are suggestions that more important changes will have to be made when means as well as variances are chosen by the forager, (McNamara pers com). These may be more that can be done analytically with the carry-over problem; it would be helpful to replace the numerical results of the carry-over section by analytical ones. Finally, I have not investigated one rather obvious problem. That is, the case where the requirement itself is a random variable. For example, overnight temperature, which is almost certainly unpredictable, may affect the amount of energy required to survive the night. There will almost certainly be conditions in which this makes little difference; however there may also be interesting conditions, say of skew in the distribution of energy gain, in which the forager should "hedge its bets" differently then when the requirement is fixed and known.

Implications. McNamara and Houston's work (McNamara MS and in press, and Houston and McNamara in press) on sequential risk-taking has important implications for the interpretation of risk-sensitivity

experiments. Specifically, one should not expect that greater than one-half of foragers choosing high variance should die, as one would predict applying the logic of the Z-score models. The model of carry-over presented here has potential implications as a general currency of maximization subsuming both risk of starvation and mean maximization models. Perhaps the most surprising prediction made here is that when the rate of increase in the value of energy carried-over (γ) is large one would expect foragers, in a fairly general way, to be risk-takers. I suggest that carry-over might explain the risk-proneness found by operant psychologists.

CHAPTER SIX

Risk: Conclusions and Discussion

List of Conclusions. In Part I, I have considered theoretical aspects of the risk-sensitive preferences of foraging animals. My argument consists of the following main points

(1) It is an empirically demonstrable fact that animals are not indifferent to variability in their food reward distributions. Animals show both risk-proneness and risk-aversion, and in some cases the forager's preferences can be changed by altering the forager's energy budget: i.e. Caraco et al.'s (1980) expected energy budget rule.

I have shown that a simple hypothesis based on minimizing the probability of death due to starvation may account for risk-sensitive foraging preferences. In particular both risk-aversion and risk-proneness can be predicted, and these Z-score models predict the expected energy budget rule.

(3) The Z-score models can be extended to consider the case of complex mean-variance sets. This allows me to consider how the risk of starvation hypothesis changes the predictions of a classical optimal foraging problem, the choice of an optimal patch residence time. This also illustrates the application of Z-score models to foraging problems beyond the simple dichotomous choice problems which inspired them.

(4) The Z-score models have well defined limitations, three of which I study in detail. Firstly, I consider the effect of allowing starvation during "the day" as well as "at nightfall," and I find that this limits the desirability of increased variance. Secondly, I consider (following McNamara in press) the importance of allowing the forager to change its variance choice in "mid-day," this is seen to

increase the pay-off (measured as the probability of survival) due to the extreme-variance rule. Thirdly, I consider the effect of allowing the energy acquired at the end of the day to be carried-over. This increases the value of high variance. This effect may account for the risk-proneness reported by operant psychologists.

Discussion: Implications and Limitations. I conclude that there are at least some cases in which mean-maximization must be abandoned in optimal foraging theory. I proposed that the minimization of the probability of starvation due to an energetic short-fall may be a reasonable alternative in cases where survival is paramount. Risk may, of course, be important in other situations, but the risk of starvation hypothesis is important as a simple and tractable alternative to mean-maximization. In short, the risk of starvation models developed here should not be applied to an animal if its natural history does not warrant it. I have illustrated my derivation with a small bird feeding in winter as an example of a case where survival is the chief attribute of fitness.

Further experimental work is clearly required. Firstly, we cannot determine whether the probability of starvation hypothesis may account for the known risk-sensitive foraging preferences until more energy budgets have been measured and manipulated. Secondly, the significance of carry-over in explaining the risk-proneness found by operant psychologists could be investigated by manipulating both mean and variance of the reward distributions offered (see chapter five). There is also room for theoretical refinement of the models discussed in Part I. Some problems are listed here: (1) What feasible sets of mean and variance are most biologically reasonable, and to what extent will they "preserve" mean maximization? (2) What are the effects of adding

continuous reward distributions and sequential decision-making to the Lazy L? (3) Can the carry-over model be extended to consider the more reasonable normal distribution of reward and the effects of correlated mean and variance?

I conclude along with Caraco (1980), McNamara & Houston (in press), Real (1980a,b), Stephens (1981) and Stephens and Charnov (in press) that there are compelling reasons for animals to be sensitive to risk in their foraging decisions. With hindsight this is not a surprising conclusion, but it is surprising that foraging theory has gone so far without reaching it.

PART II: INFORMATION