CENTRAL-PLACE FORAGING IN A STOCHASTIC ENVIRONMENT: A MULTIPLE-PREY LOADER

By RICHARD F. GREEN,* ADELINE TAYLOR NUÑEZ† AND ALEJANDRO KACELNIK‡

*Department of Mathematical Sciences, University of Minnesota, Duluth, MN 55812, U.S.A. †Forest Wildlife Project, Department of Fisheries and Wildlife, University of Minnesota, Cloquet Forestry Center, 175 University Road, Cloquet, MN 55720, U.S.A. and ‡Edward Grey Institute of Field Ornithology, Department of Zoology, South Parks Road, Oxford OX1 3PS.

SUMMARY

(1) This paper gives the theoretical development of a stochastic model for a central-place foraging experiment with starlings. Adult birds obtained food for their nestlings by visiting a feeder that rewarded pecks at a rate which was high on some visits and low on other visits (at random).

(2) We consider a number of different types of rules which the birds could use to decide when to leave the feeder and return to the nest. For each rule we calculate the long-term average rate of bringing food to the nest that could be achieved.

(3) For long enough travel time to and from the nest, the best rule is to remain at the feeder until a full load of prey is obtained, whether the reward rate is high or low. For shorter travel times, the best rule is, roughly, to remain on each visit until either a full load is obtained, or experience shows that the odds are low that the patch is good.

(4) While different rules achieve different rates of obtaining prey, the rates achieved by rules of a given type are insensitive to the specific rules used.

INTRODUCTION

Animals that forage on clumped resources may find that some clumps (patches) have more prey than others. This variability was not taken into account in most "classical" optimal foraging models. For example, the model for which Charnov (1976) stated the marginal value theorem (MVT) assumed that the amount of energy gained by a predator that has spent some time, t, foraging in a patch is determined uniquely by t. Empirical tests of this model either ignored patch variability when accounting for observed results (Cowie 1977, Krebs, Ryan & Charnov 1974, Carlson & Moreno 1982, Bryant & Turner 1982) or artificially eliminated between-patch variance by using deterministic experimental patches (Kacelnik, Houston & Krebs 1981, Kacelnik 1984, Ydenberg 1982).

There is, however, a whole family of theoretical models that include environmental stochasticity. First, Oaten (1977) showed that Charnov's MVT does not apply generally to foraging for discrete prey items in stochastic environments in which predators may gain information about patch quality; he proposed a general stochastic model and described how to find the optimal strategy. Pyke (1978) suggested a discrete stochastic version of Charnov's MVT. Breck (1978) considered the performance of suboptimal rules in environments with different distributions of numbers of prey per patch. Green (1980) found the optimal strategy for a special case of Oaten's stochastic model and compared the rates of finding prey achieved by the optimal strategy and other strategies. Iwasa et al. (1981) showed how the form of optimal (or nearly optimal) rules depends on the variability of patch quality. Krebs, Kacelnik & Taylor (1978) showed that, when animals exploit stochastic patches which can be visited repeatedly, the advantages resulting from learning the properties of each patch depend on the expected stability of the environment (the "time horizon"), and Houston, Kacelnik & McNamara (1982) explored the performance of a number of decision rules in the same kind of situation.

This paper describes a stochastic model of central-place foraging, in which the forager must return periodically to its nest with food (Orians & Pearson 1979). In general, the strategy of a central-place forager includes a choice of which prey to take, how many to take, and when to leave a patch and return to the nest.

We model an experiment performed with starlings that were provided food at a feeder that was located some distance from their nests to which they returned to feed their young (Kacelnik et al.   ). The feeder was an operant food dispenser with a random ratio schedule, with a prey item delivered at random with a fixed probability of reward on each peck during a visit. It is assumed that all prey were the same and that all foraging took place at the feeder. On each visit to the feeder the reward probability for the entire visit was one of two values, assigned randomly for successive visits. Using the terminology of other foraging studies, we refer to the feeder as the "patch."

The foraging strategies we consider consist of stopping rules that the birds could use to decide when to leave the patch and return to the nest with a load of food. We will compare several such stopping rules.

The model of the environment contains four parameters:

α : the proportion of patch (feeder) visits with good reward probability.

p_1: the probability of a reward on each peck for a good patch.

p_2: the probability of a reward on each peck for a poor patch.

τ: the return trip travel time between the central place and the patch, including time at the central place. Travel time and patch time are measured in peck equivalents. For example, if travel time is 30 seconds and the average interval between pecks is 1.5 seconds, then $\tau = 20$.

We assume that the bird stays in the patch until it has taken at least one prey, and that it does not eat any prey for itself. The maximum number of prey that can be transported, M, is a constraint for the model. Our calculations are done for a sample set of environmental parameters: $\alpha = 0.5$, $p_1 = 0.2$, $p_2 = 0.08$, $M = 6$, and several values of τ. These values of α, p_1 and p_2 are the ones used by Kacelnik et al. ( ).

We compare different stopping rules by calculating the long-term average rate of delivering prey to the nest that each would achieve. This rate is given by

$$R = E(G)/(\tau + E(T)) , \qquad (1)$$

where $E(G)$ = the average number of prey found per patch visit, and $E(T)$ = the average number of pecks per patch visit. Notice that R is the gross, rather than the net, rate of finding prey; that is, we do not include the energy costs in the equation. For a discussion of the importance of this simplification see Kacelnik (1984) and Kacelnik & Houston (1984).

We consider three types of rules as well as the best rule, that is, the rule, of any type, which achieves the highest possible rate R. There may be several different rules of a particular type. For example, there are many different Odds rules, each corresponding to a different choice of the odds. We can compare the rates achieved by different rules of a particular type and find the best rule of that type. Then we can compare the rates achieved by the best rule of each type. The types of rules we consider are:

1. Fulloader. The forager obtains M prey on all visits. Since we assume that M = 6, there is only one rule of this type.

2. mT. A rule of this type is specified by a pair of numbers, (m,T). The forager remains in each patch until a minimum number, m, prey have been found. If the mth prey is obtained after time T, the forager leaves the patch. Otherwise it remains until M prey have been obtained.

3. Odds. This rule is specified by a single number, O, the odds. If, at any time during a visit, after one prey has been found, the odds that the animal is foraging in a good patch (determined by the number of prey found and the number of pecks taken) fall below O, the forager leaves the patch. Otherwise it remains until M prey have been obtained.

The best rule is the one which achieves the highest possible rate subject to the constraints that at least one prey but no more than M are taken from each patch. This rule yields the maximum possible rate. The rule is specified by M - 1 pairs of numbers, (t_m,m), where m takes all the integer values from 1 to M - 1. Each (t_m,m) pair represents a stopping point; that is, the patch is left if time t_m is reached and only m prey have been found.

Rules of the Odds and mT types are fairly realistic and are proposed as working hypotheses for analyzing the performance of the starlings. We use the Fulloader rule and the best rule to explore the limits of performance in a given environment and to evaluate the rates achieved by the other stopping rules.

For each type of rule, we find the best rule of that type (for example, by specifying the values of m and T for the mT rule which produce the maximum possible rate of finding prey). We compare the rates achieved by the best rule of each type for different travel times, τ. The calculations that determine the best rule of each type are given in the Appendix. The performance of the rules is discussed in the next section.

ANALYSIS OF THE STOPPING RULES

The Fulloader Rule.

An animal using this rule would remain in a patch until a full load of M prey is found. The long term average rate of finding prey would be

$$R = \frac{M}{\tau + M\left[\alpha/p_1 + (1-\alpha)/p_2\right]} = \frac{1}{\tau/M + \alpha/p_1 + (1-\alpha)/p_2} . \qquad (2)$$

The Fulloader rule will be best if the travel time τ is long enough (that is, if τ exceeds some critical value, τ_{crit}), since for a long travel time the animal cannot achieve a high average rate of finding prey. For travel time exceeding τ_{crit}, the average rate would be lower than the rate achieved in a poor patch, so the animal should remain in a poor patch even if the patch is known to be poor. As shown in the Appendix,

$$\tau_{crit} = M\,(1/p_2 - 1/p_1)\,. \qquad (3)$$

In Table 1 we can see how the average rate achieved by the Fulloader rule declines as the travel time increases.

(Put Table 1 here.)

If τ is small, or M is large, the rate achieved by the Fulloader rule is approximately equal to the harmonic mean of p_1 and p_2, the rates of finding prey achieved in good and poor patches.

<u>The</u> <u>mT</u> <u>rule</u>.

Like the Fulloader rule, the mT rule is very simple. An animal using this type of rule would stay in a patch until some minimum number of prey, m, have been found. If finding the m prey takes T pecks or more, the animal leaves the patch and returns to the nest. If m prey are found in fewer than T pecks, the animal remains in the patch until a full load, M, has been found.

We have calculated the rates for travel times $\tau = 0$, 5, 10, 15 and 20. For each travel time, it is interesting to see which rule (which set of m and T) yields the highest rate and also to see how sensitive the rate is to changes in the values of m and T.

Table 2 gives the optimal values of m and T for several travel times.

(Put Table 2 here.)

The best values of m and T depend on τ. Compare the optimal values of m and T for $\tau = 0$, 5 and 10 in Table 2. The best value of m is 2 in each case, and the best T values are 12, 17 and 21, respectively. To see why these stopping times are different for these different travel times, we can consider the ratio of expected gain to expected loss for a forager that decides, when m prey have been found in T pecks, to remain in the patch until a full load has been obtained. An animal using the best mT rule will achieve some rate, call it R^*. An animal having found m prey in a good patch and having decided to remain until a full load is obtained would, on average, find

$$B = (p_1 - R^*)(M - m)/p_1$$

more prey during the remaining stay than would be found during an average interval of the same length (not necessarily spent in a good patch) by an animal using the best mT rule. Similarly, an animal having found m prey in a poor patch and deciding to remain until a full load is obtained would, on average, find

$$C = (R^* - p_2)(M - m)/p_2$$

fewer prey during the remaining stay than would be found during an average interval of the same length by an animal using the best mT rule. We may think of

$$B/C = (p_2/p_1)(p_1 - R^*)/(R^* - p_2)$$

as the benefit/cost ratio of continuing in a patch, where the benefit results from staying in good patches and the cost results from staying in poor patches.

Now the question is, given an animal's experience in a patch, what are the odds that the patch is good? For an animal that has found m prey in T pecks, the odds that the patch is a good one are

$$\begin{aligned} \text{Odds} &= (\alpha/(1-\alpha))(p_1/p_2)^m \left[(1-p_1)/(1-p_2)\right]^{T-m} \\ &= (\alpha/(1-\alpha)) \left[(p_1(1-p_2))/(p_2(1-p_1))\right]^m \left[(1-p_1)/(1-p_2)\right]^T \end{aligned} \qquad (4)$$

which is an exponentially decreasing function of T. A forager should remain in a patch as long as the product of the odds and the benefit/cost ratio is at least one, and leave when that product is less than one. The critical value of T is the smallest number of pecks such that the product of the odds and the benefit/cost ratio is less than one.

For the case $\alpha = .5$, $p_1 = .2$ and $p_2 = .08$ we have, from Table 1, $R^* = .1275$, .1105 and .0988 for $\tau = 0$, 5 and 10, respectively. For these cases we have that the benefit/cost ratios are .6105, 1.1738 and 2.1532. For the best values of m and T, the odds that the patch is good are 1.5449, .7681 and .4392.

Now consider, for example, the case $\tau = 10$. Here $R^* = .0988$, benefit/cost = 2.1532, and for $T = 21$, odds = .4392 (while, for $T = 20$, odds = .5050). The product of odds and benefit/cost is .9457 for $T = 21$ (and is 1.0874 for $T = 20$). Thus, for $\tau = 10$, a forager that takes 20 (or fewer) pecks to find $m = 2$ prey stands to gain more from good patches than it would lose from poor patches if it decides to remain in the patch. A forager that takes 21 (or more) pecks to find $m = 2$ prey stands to lose more from poor patches than it would gain from good patches if it decides to remain.

Why are the best values of T different for different travel times, τ? When travel time is longer, the highest possible rate of finding prey will be lower and will thus be closer to the rate of finding prey in poor patches. Thus, the ratio of potential gain from foraging in good patches to potential loss from foraging in poor patches (B/C) will be greater. This, in turn, means that when travel time is longer a forager should remain in a patch when the odds that the patch is good are lower.

Therefore, foragers should tend to stay longer in patches when travel time is longer.

For $\tau = 0$, 5 and 10 the best mT rules have m = 2, but for the longer travel times, $\tau = 15$ and 20, the best rules have m = 3. Having m = 3 rather than 2 means that the forager must wait longer (and find one more prey) before deciding whether to remain in the patch or leave. Staying longer provides more information and enables the animal to make a better decision. However, if the patch is poor a forager using a larger value of m will waste time which could be used feeding in good patches. It is best to delay the decision only when there is little to lose by staying in poor patches; this happens when travel time is long.

When travel time is long enough ($\tau \geq 22.5$ for $\alpha = .5$, $p_1 = .2$ and $p_2 = .08$), the Fulloader rule is best, and the forager should stay in all patches until a full load (M = 6) of prey has been found. When travel time is short, however, Table 1 shows that the best mT rule can achieve a substantially higher rate than can the Fulloader rule.

It is important to note the insensitivity of the rate achieved to the rule used. Consider, for example, the case $\tau = 10$. The highest rate for an mT rule will be achieved for m = 2 and T = 21. For m = 2, Figure 1 plots the rates achieved for various values of T. Within a wide range of values, the rate changes very little.

(Put Figure 1 here.)

The Odds rule.

An animal using the mT rule decides to remain in a patch until a full load has been obtained if the odds that the patch is good are high enough at the point when it finds the mth prey. An animal using the Odds rule will remain in a patch as long as the odds that the patch is a good one exceed some critical value, call it 0*. The animal will leave the patch after a full load has been obtained or if, at any time after having found at least one prey, the odds that the patch is good falls to (or below) the critical value 0*. Each choice of 0* determines the behavior of a forager that follows some Odds rule.

Figure 2 illustrates, for several different choices of 0*, the odds lines, that is, the sets of values of m and t such that the odds (given by (4)) equal a given 0*. An animal using an Odds rule will leave a patch if it reaches any point on or to the right of the chosen odds line. The actual stopping points for the best Odds rules are given in Table 3 for several values of $\mathcal{T}$.

The Odds and mT rules are proposed as realistic stopping rules that an animal can use. We saw that the long-term rate of finding prey achieved by an mT rule was quite insensitive to the particular value of T; how sensitive is the performance of an Odds rule to the choice of 0*? In Figure 3 the rates achieved for various values of O^* are plotted against $-\log(O^*)$. Note that the rate achieved by an Odds rule is not very sensitive to the value of the odds used. Figure 2 illustrates the differences in the Odds rules determined by different choices of O^*. Figure 3 shows that these very different rules yield quite similar rates of finding prey. As with the mT rule, if an animal uses an Odds rule, it does not matter very much exactly which rule is used.

(Put Figures 2 and 3 here.)

<u>The best rule</u>.

The best rule is very similar to the Odds rule, and, as Table 1 shows, is not much better, but it is much more difficult to find. The best rule is found by dynamic programming. Details are given in the Appendix.

The best rules for τ = 0, 5, 10, 15 and 20 are given in Table 3, where they are compared with the corresponding best Odds rules. While the rules are very similar, it can be seen that the set of stopping points for the best rule is slightly steeper than that for the best Odds rule. This is because, as McNamara (1982) has pointed out, an animal foraging in a variable environment gains information as it searches a patch. When the animal has spent a small time in a patch it still has much information to gain about patch quality, and the value of this information is added to the value of potential prey. Thus, because of the value of gaining information, an animal should be more reluctant to leave a patch shortly after beginning to forage there than if the decision were based solely on current opinion of patch quality. As more time is spent in a patch (and more prey are found) the value of additional information is less and an optimal forager becomes more willing to leave the patch.

(Put Table 3 here.)

Since the best Odds rule is so similar to the best rule, both in its form and in the rate it achieves, it is convenient to use the Odds rule to evaluate the starling experiments. It should be borne in mind, however, that the best Odds rule is not (quite) the best rule.

DISCUSSION

We have described four types of rules that could be used by animals foraging in a stochastic environment. Can we compare these rules with stopping rules suggested for other models?

Charnov (1976) stated his marginal value theorem for a deterministic model in which the amount of energy gained by a forager is a strictly determined function of the time the forager spends in a patch. Charnov assumed this function is continuous and has a positive first derivative and a negative second derivative. Under these conditions, the best foraging strategy is to leave a patch when the rate of gaining energy in that patch falls to a value equal to the highest long-term rate of obtaining energy possible in the environment.

One trouble with the marginal value theorem is that it does not provide a stopping rule that a forager could use in a stochastic environment. Krebs et al. (1974) suggested that animals might use the "giving-up time" rule, in which a patch would be abandoned if the animal had not found a prey for some fixed time (the "giving-up time," or GUT). Krebs et al. seem to have identified the GUT rule with the marginal value theorem and with optimal foraging, but, in fact, the GUT rule is not, in general, optimal, nor does it satisfy the marginal value theorem (see McNair 1982, Green 1984).

Pyke (1978) devised a "stochastic discrete" version of the marginal value theorem in which an animal foraging in a variable environment should leave a patch when its expected rate of finding prey in the patch reached the best long-term rate. Pyke's rule is described implicitly, but an explicit version can be found in many cases. In our model Pyke's rule would be an example of the Odds rule, but it would not be the best Odds rule, which is itself not the best rule.

However, there is a sense in which an idea much like the marginal value theorem can be used to find the optimal foraging strategy in a stochastic model. For any particular strategy it is possible, in principle, to calculate the expected energy gain (E(G)) and the expected length of time (E(T)) in a patch for each visit. If we consider a set of possible strategies, we can plot the points (E(G),E(T)) for each strategy. This is done in Figure 4 for a set of Odds rules. Each point in Figure 4 is determined by a particular Odds rule, which is itself characterized by a set of stopping points such as are illustrated in Figure 2. For a given travel time, τ , the (E(G),E(T)) point corresponding to the best strategy in the set will be such that the line from $(-\tau,0)$ through the point will pass through or above all the (E(G),E(T)) points corresponding to the other strategies.

(Put Figure 4 here.)

While our Figure 4 resembles the figure with which Charnov illustrated his marginal value theorem, and while the reason that each leads to the optimal strategy is the same, there is no marginal value theorem in our case. It is meaningless to talk about derivatives in our model, since the set of possible points (E(G),E(T)) is discrete. Further, while Figure 4 does help show how to find which Odds rule is best, it does not show anything about the form or nature of Odds rules.

In this paper we develop the theory for a particular model of central-place foraging, itself a special case of foraging. The strategy that we are interested in does not involve patch choice or prey choice, but simply the rules an animal could use to decide when to leave a patch. The advantage of a particular model is that it is possible to calculate the rates of finding prey for various strategies under various environmental conditions. For more general models, it is impossible to compare different types of strategies and find the best one. It is also impossible to see how sensitive the rates achieved are to differences in strategies.

For our model, we see that the Fulloader rule is best if travel time is long enough. (The best mT rule and the best Odds rule are equivalent to the Fulloader rule for long travel time.) For shorter travel time, the best mT rule is better than the Fulloader rule, while the best Odds rule is better still. In fact, the best Odds rule is almost as good as the best rule, which has a very similar form.

To decide what rule a forager actually uses, we must check the details of the leaving rules. Do animals leave a patch immediately after finding a prey or do they leave after going for some time without finding a prey? Animals using the Fulloader rule or an mT rule would leave immediately after finding a prey. Animals using an Odds rule or the best rule would often leave a patch after one or more unsuccessful pecks.

It is possible to calculate the proportion of patch visits that would end with a successful peck, for various stopping rules. For the environmental parameters we have used, an animal using the best Odds rule would leave a patch immediately after a successful peck with probabilities .577, .601, .631, .661 and .719 for τ = 0, 5, 10, 15 and 20, respectively. Thus, in only about two-thirds of patch visits would an animal using the best Odds rule or the best rule leave a patch immediately after finding a prey.

Our most important conclusion is that for an animal using a certain type or rule, the rate achieved is not very sensitive to the exact rule used. This insensitivity implies that, since it makes very little difference exactly what the animal does, we might expect to see highly variable behavior in the experiments. Such behavior would not mean that the animal is behaving suboptimally, but simply that the behavior is varying over the wide range of approximately optimal behavior. It is important to judge behavior, not by how similar it is to optimal behavior, but rather by how similar the results it achieves are to the results of optimal behavior.

REFERENCES

Breck, J. E. (1978). Suboptimal foraging strategies for a patchy environment. Ph. D. Diss. Michigan State University.

Bryant & Turner (1982)

Carlson & Moreno (1982)

Charnov, E. L. (1976). Optimal foraging: the marginal value theorem. Theoretical Population Biology, 9,129-136.

Cowie (1977)

Green, R. F. (1980). Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging. Theoretical Population Biology, 18,244-256.

Green, R. F. (1984). Stopping rules for optimal foragers. American Naturalist, 123,30-43.

Houston, Kacelnik & McNamara (1983)

Iwasa, Y., Higashi, M., & Yamamura, N. (1981). Prey distribution as a factor determining the choice of optimal foraging strategy. American Naturalist, 117, 710-723.

Kacelnik (1983)

Kacelnik, Gorter & van Beek (unpublished ms.)

Kacelnik & Houston (1983)

Kacelnik, Houston & Krebs (1981)

Krebs, J. R., Kacelnik, A. & Taylor, P. (1978). Test of optimal sampling by foraging great tits. Nature, 275,27-31.

Krebs, J. R., Ryan, J. C. & Charnov, E. L. (1974). Hunting by expectation or optimal foraging? A study of patch use by chickadees. Animal Behavior, 22,953-964.

McNair, J. N. (1982). Optimal giving-up times and the marginal value theorem. American Naturalist, 119,511-529.

McNamara, J. (1982). Optimal patch use in a stochastic environment. Theoretical Population Biology, 21,269-288.

Oaten, A. (1977). Optimal foraging in patches: a case for stochasticity. Theoretical Population Biology, 12,263-285.

Orians, G. H. & Pearson, N. E. (1979). On the theory of central place foraging. Analysis of Ecological Systems (Ed. by D. J. Horn, R. Mitchell & G. R. Stair), pp. 155-177. Ohio State University Press, Columbus, Ohio.

Pyke, G. H. (1978). Optimal foraging in hummingbirds: testing the marginal value theorem. American Zoologist, 18,739-752.

APPENDIX

Here we show how to calculate the long-term average rates of finding prey for animals using each of the types of rules we have discussed. In general,

$$k = EG/(\tau + ET) \ , \qquad \text{(A1)}$$

where k = long-term average rate of finding prey; EG = expected number of prey found per patch visited; ET = expected time (number of pecks) in each patch visited; τ = the average time (measured in peck equivalents) spent travelling on one return trip to the nest.

a) For the <u>Fulloader</u> rule we have

$$k = \frac{M}{\tau + M\left[\alpha/p_1 + (1-\alpha)/p_2\right]} \ , \qquad \text{(A2)}$$

where M = number of prey collected in each patch (maximum load); α = proportion of good patches (and $1-\alpha$ = proportion of poor patches); p_1 and p_2 = probabilities of reward per peck in good and poor patches, respectively. The average numbers of pecks needed to find one prey in good and in poor patches are $1/p_1$ and $1/p_2$, respectively.

A bird using the Fulloader rule remains in a patch until it obtains a maximum load of prey, M. If the travel time τ is long enough, this rule is the best; that is, it is optimal to stay in all patches until a maximum load is collected, even in poor patches. This will happen if the overall rate of finding prey is lower than the rate in poor patches,

$$R < p_2 \,,$$

or, equivalently, if τ is greater than a critical value, τ_{crit}, which satisfies the equation

$$p_2 = \frac{M}{\tau_{crit} + M\left[\alpha/p_1 + (1-\alpha)/p_2\right]}\,,$$

which is equivalent to

$$\tau_{crit} = \alpha M\left[1/p_2 - 1/p_1\right] \;.$$

If $\tau < \tau_{crit}$, then the Fulloader rule will not be best, and it will pay the forager to use a rule which assesses patch quality. Three such rules are considered.

b) For the <u>mT</u> rule the forager stays in each patch until some minimum number of prey, m, have been found. If at least some critical number of prey, T, are required to capture the m prey, then the forager leaves the patch immediately and returns to the nest. If m prey are found in fewer than T pecks, then the forager does not leave the patch until a full load M has been obtained.

For an animal using an mT rule, the question is whether or not it finds m prey in fewer than T pecks. In order to calculate the probability of various outcomes it is convenient to distinguish between whether a patch is good or poor, and whether the forager leaves the patch after finding m prey or continues until it obtains a full load. We calculate EG and ET using the probabilities that an animal will reach particular points at which it decides to remain in a patch or not.

An animal that decides to remain in the patch until capturing M prey will have found its m<u>th</u> prey at one and only one time, m, m+1, ..., T-1. The decision would be made when the animal reached the point (i,m), for i = m, m+1, ..., T-1. The probabilities of making the decision at these points would be negative binomial

$$n_1(i,m) = \binom{i-1}{m-1} p_1^{\,m} (1 - p_1)^{i-m} \qquad (A4)$$

for good patches, and

$$n_2(i,m) = \binom{i-m}{m-1} p_2^{\,m} (1 - p_2)^{i-m} \qquad (A5)$$

for poor patches.

An animal that finds its m<u>th</u> prey at time $t \geq T$ will decide to leave the patch. To avoid summing infinitely many terms, we note that an animal that finds its m<u>th</u> prey at time $t \geq T$ would have found fewer than m prey at time t = T-1. Thus, an animal that reaches one of the points (T-1,j), for j = 0, 1, ..., m-1, may be considered to have "decided" to leave the patch as soon as m prey have been found, even though it will still remain in the patch until m prey have been found.

The probabilities of reaching the points (T-1,j), for j = 0, 1, ..., m-1, are binomial

$$b_1(T-1,j) = \binom{T-1}{j} p_1^{\,j} (1 - p_1)^{T-1-j} \tag{A6}$$

for good patches, and

$$b_2(T-1,j) = \binom{T-1}{j} p_2^{\,j} (1 - p_2)^{T-1-j} \tag{A7}$$

for poor patches.

Using (A4) - (A7) we have

$$\begin{aligned} EG = M \sum_{i=m}^{T-1} \left[\alpha\, n_1(i,m) + (1-\alpha)\, n_2(i,m)\right] + \\ m \sum_{j=0}^{m-1} \left[\alpha\, b_1(T-1,j) + (1-\alpha)\, b_2(T-1,j)\right], \end{aligned} \tag{A8}$$

The summation terms in (A8) represent the probabilities of remaining in a patch until M prey are found, or leaving a patch when m prey are found. We also have

$$\begin{aligned} ET = \sum_{i=m}^{T-1} \Big\{ &\left[i + (M - m)/p_1\right] n_1(i,m) + \\ &(1-\alpha) \left[i + (M - m)/p_2\right] n_2(i,m) \Big\} + \\ \sum_{j=0}^{m-1} \Big\{ &\alpha \left[T-1 + (m - j)/p_1\right] b_1(T-1,j) + \\ &(1-\alpha) \left[T-1 + (m - j)/p_2\right] b_2(T-1,j) \Big\}, \end{aligned} \tag{A9}$$

where, for example, the expression $i + (M - m)/p_1$ represents the average number of pecks taken in a good patch by a bird which finds its <u>mth</u> prey on its <u>ith</u> peck and then stays until it finds $M - m$ more prey, each of which takes an average of $1/p_1$ more pecks to find.

c) An <u>Odds</u> rule is characterized by the choice of a value for the odds that a patch is good. For a given set of values of α, p_1 and p_2, a particular value of the odds, say O, determines a set of stopping times, $t_1, t_2, \ldots, t_{M-1}$. Each t_i is the smallest integer, t (number of pecks), such that the odds of a patch being good if the animal reaches the point (t_i, i) are less than O.

An animal using this rule returns to the nest immediately if

(1) a full load of M prey has been found, or

(2) i prey have been found in t_i pecks, or

(3) the first prey is found after more than t_1 pecks.

An animal that reaches one of the points in category (3), namely, $(t,1)$, for $t > t_1$, must have passed through the point $(t_1, 0)$, which we may think of as a "virtual" stopping point; that is, an animal that reaches the point $(t_1, 0)$ has "decided" to stop as soon as the first prey is found. Our calculations become simpler if we define $t_0 = t_1$, and consider $(t_0, 0)$ as a stopping point.

While each stopping point has two coordinates, we can index the points by the integers $k = 0, 1, 2, \ldots, t_{M-1}$;

where k = 0 indicates the "virtual" stopping point $(t_0, 0)$; k = 1, 2, ..., M-1 indicate the stopping points (t_k, k); and k = M, M+1, ..., t_{M-1} indicate the stopping points (k,M).

The probability of reaching a particular stopping point will be the product of the number of paths to the point, w_k, and the probability, q_k, of reaching the point along a single path. For q_k we have

$$q_k = \alpha p_1^{\,k} (1 - p_1)^{t_k - k} + (1-\alpha)\, p_2^{\,k} (1 - p_2)^{t_k - k} \qquad \text{(A10)}$$

for k = 0, 1, 2, ..., M-1, and

$$q_k = \alpha p_1^{\,M} (1 - p_1)^{k-M} + (1-\alpha)\, p_2^{\,M} (1 - p_2)^{k-M} \qquad \text{(A11)}$$

for k = M, M+1, ..., t_{M-1}. For w_k we have

$$w_0 = 1, \quad w_1 = \binom{t_1}{1} = t_1, \text{ and}$$

$$w_k = \binom{t_k}{k} - \sum_{i=0}^{k-1} w_i \binom{t_k - t_i}{k-i} \qquad \text{(A12)}$$

for k = 2, 3, ..., M-1, and

$$w_k = \binom{k-1}{M-1} - \sum_{\{i:\ k-t_i \quad M-i\}} w_i \binom{k-t_i-1}{k-i-1} \qquad \text{(A13)}$$

for k = M, M+1, ..., t_{M-1} .

Now, using (A10) - (A13), we may calculate EG and ET.

$$EG = w_0 q_0 + \sum_{k=1}^{M-1} k w_k q_k + M \sum_{k=M}^{t_{M-1}} w_k q_k , \qquad \text{(A14)}$$

and

$$ET = \alpha (t_1 + 1/p_1)(1 - p_1)^{t_1} + (1-\alpha)(t_1 + 1/p_2)(1 - p_2)^{t_1} + \sum_{k=1}^{M-1} t_k w_k q_k + \sum_{k=M}^{t_{M-1}} k w_k q_k \,. \qquad \text{(A15)}$$

d) The <u>best</u> rule (subject to the constraints that at least one, but no more than M prey, must be taken in each patch visited) is very similar to the Odds rule, and, once the best rule has been found, the rate that it achieves can be calculated in the same way as that for the Odds rule. The problem is to find the stopping times, t_1, t_2, ..., t_{M-1}, similar to those for the Odds rule.

The best rule is found by dynamic programming, using backward recurrence. The method is very similar to that used by Green (1980), but here the rule is found, not by starting at the end of a patch (here patches are assumed to be endless), but rather at the largest number of prey that can be collected, M-1, without leaving a patch. To find the best rule six steps are followed.

Step 1. First, guess a "candidate" value, C, for the highest long-term average rate of finding prey. Then find a rule that would be best if C were, in fact, the highest possible rate (namely, C*). This is done in Steps 2-5.

Step 2. Next, find the stopping time, t_{M-1}. Notice that if a given number of prey have been found, then the fewer pecks it has taken to find them, the larger is the chance that the patch is good. That is, if x prey have been found in t pecks, the probability that the patch is good is given by

$$p(t,x) = \frac{\alpha p_1^{x} (1 - p_1)^{t-x}}{\alpha p_1^{x} (1 - p_1)^{t-x} + (1-\alpha) p_2^{x} (1 - p_2)^{t-x}} , \quad \text{(A16)}$$

which is a decreasing function of t.

If an animal has reached a point (t,x), the probability of finding a prey with the next peck is given by

$$P(\text{Prey} \mid t,x) = p(t,x)\ p_1 + \left[1 - p(t,x)\right] p_2 , \quad \text{(A17)}$$

which is also a decreasing function of t for a fixed x. The stopping time, t_{M-1}, is the smallest value of t such that

$$P(\text{Prey} \mid t,x) < C . \quad \text{(A18)}$$

This inequality means that a bird taking a peck after the point (t,x) would have a lower expected rate of gain for that peck than the long-term rate, C.

Step 3. Since the points, (t,M-1), for t = M-1, M, M+1, ..., $t_{M-1}-1$, are "continuation" points, we may calculate the expected numbers of additional prey found and additional pecks made, EG(t,M-1) and ET(t,M-1), respectively, for a bird which reaches the point (t,M-1), and which will continue in the patch until it has made

t_{M-1} pecks or has found the Mth prey. We have

$$EG(t,M-1) = p(t,M-1)\left[1 - (1 - p_1)^{t_{M-1}-t}\right] + \left[1 - p(t,M-1)\right]\left[1 - (1 - p_2)^{t_{M-1}-t}\right], \tag{A19}$$

and

$$ET(t,M-1) = p(t,M-1)\left[1 - (1 - p_1)^{t_{M-1}-t}\right]/p_1 + \left[1 - p(t,M-1)\right]\left[1 - (1 - p_2)^{t_{m-1}-t}\right]/p_2 . \tag{A20}$$

Step 4. New find t_{M-2} (and, using the same method, t_{M-3}, t_{M-4}, ..., t_1), using the values found from (A19) and (A20). Notice that the point, $(t_{M-1}+1, M-2)$, will not be a continuation point, since a bird reaching that point cannot continue in the patch for more than one additional peck, and from (A16) and (A17),

$$P(\text{Prey} \mid t_{M-1}-1, M-2) < P(\text{Prey} \mid t_{M-1}, M-1) < C .$$

Now, imagine a bird that, at point (t,M-2), for $t = M-2, M-3, \ldots, t_{M-1}-2$, considers adopting the rule: Make one more peck. If a prey is not found, leave the patch immediately. Otherwise, continue in the patch, using the rule determined by the stopping time, t_{M-1}. The expected number of additional prey found and pecks taken using such a rule would be given by

$$EG'(t,M-2) = P(\text{Prey} \mid t,M-2)\left[1 + EG(t+1,M-1)\right], \tag{A21}$$

and

$$ET'(t,M-2) = 1 + P(\text{Prey} \mid t,M-2)\ ET(t+1,M-1), \qquad (A22)$$

respectively. The values of EG' and ET' are found successively, starting with $t = t_{M-1}-2$ and working backward as long as

$$EG'/ET' \angle C\ . \qquad (A23)$$

Then t_{M-2} will be the smallest value of t such that (A23) is true (or $t_{M-2} = t_{M-1}-1$ if the inequality does not hold for $t = t_{M-1}-2$).

The primes in EG' and ET' indicate that these are provisional expectations, assuming that the bird decides to take another peck and leave the patch if it finds no prey on that peck. Once the stopping time, t_{M-2}, has been determined, the actual expectations, EG(t,M-2) and ET(t,M-2), for t = M-2, M-1, M, ..., t_{M-2}, can be found recursively, starting at $t = t_{M-2}$ and working backward, assuming that the bird uses the rule that has been built up so far (that is, t_{M-2} and t_{M-1}). We have

$$EG(t_{M-2},M-2) = 0,\ ET(t_{M-2},M-2) = 0$$

and, for $t = t_{M-2}-1,\ t_{M-2},\ \ldots,\ M-2$,

$$EG(t,M-2) = P(\text{Prey} \mid t,M-2)\left[1 + EG(t+1,M-1)\right] + \left[1 - P(\text{Prey} \mid t,M-2)\right] EG(t+1,M-2), \qquad (A24)$$

and

$$ET(t,M-2) = 1 + P(Prey \mid t,M-2)\ ET(t+1,M-1) + \left[1 - P(Prey \mid t,M-2)\right] ET(t+1,M-2) \quad \text{(A25)}$$

Step 5. The stopping time, t_{M-3}, is found by repeating the process given in step 4, by defining EG'(t,M-3) and ET'(t,M-3) and comparing their ratio with C, as in (A23). When the value of t_{M-3} has been found, the values of EG(t,M-3) and ET(t,M-3) can be found by using formulas similar to (A24) and (A25). The process is repeated to find t_{M-4}, and so on down to t_1.

Step 6. Once the values of t_1, t_2, ..., t_{M-1} have been found, the rate achieved by the rule is determined by the same method as for the Odds rule. This rate, R(C), is a function of the guessed rate, C. If $R(C) = C$, the best rule has been found. If $R(C) \neq C$, another guess of C is made (the last R(C) is a good guess), and the process is repeated until the value of C is found such that $R(C) = C$. This value is the best possible rate, C*.

TABLE 1. Long-term average rates of finding prey achieved by a forager using the best rule of each of four different types, for different values of the travel time, τ. In all cases the environmental parameters are $\alpha = .5$, $p_1 = .2$, and $p_2 = .08$.

Travel time τ	Type of rule: Fulloader	mT	Odds	Best
0	.114286	.127541	.141650	.141714
5	.104348	.110514	.117109	.117156
10	.096000	.098757	.101985	.102013
15	.088889	.089948	.091247	.091261
20	.082759	.082948	.083199	.083203

TABLE 2. Best mT rules, for $\alpha = .5$, $p_1 = .2$, $p_2 = .08$. The best values of m and T are given for various values of travel time, τ.

τ	m	T
0	2	12
5	2	17
10	2	21
15	3	34
20	3	43
25	6	∞

TABLE 3. Comparison of stopping times for the best Odds rules and the best rules, for $\alpha = .5$, $p_1 = .2$, $p_2 = .08$ and for various values of travel time, τ. For a given type of rule and a given travel time, the table gives the best stopping time (the number of pecks at the feeder), if the specified number of prey have been found.

Travel time		Number of prey found				
		1	2	3	4	5
$\tau = 0$	Odds Rule	12	19	27	34	42
	Best Rule	12	19	26	32	38
$\tau = 5$	Odds Rule	18	25	33	40	48
	Best Rule	19	25	32	38	44
$\tau = 10$	Odds Rule	22	30	37	45	52
	Best Rule	24	30	37	43	49
$\tau = 15$	Odds Rule	28	35	43	50	58
	Best Rule	29	36	42	49	55
$\tau = 20$	Odds Rule	37	44	52	59	67
	Best Rule	39	45	52	58	64

Figure Legends

FIGURE 1. Average rates of finding prey achieved by mT rules with m = 2, plotted against T, for T = 3, 4, 5, ..., 50. Here $\alpha = .5$, $p_1 = .2$, $p_2 = .08$ and $\tau = 10$.

FIGURE 2. The odds lines for various values of odds (or, log(odds)), for $\alpha = .5$, $p_1 = .2$, $p_2 = .08$. Each of the diagonal lines represents points where the odds that a patch is good equal a constant. An animal using an Odds rule would remain in a patch until it crossed the appropriate odds line or hit one of the horizontal lines.

FIGURE 3. Rates achieved by various Odds rules, plotted against log(odds). Each point represents the rate achieved for a particular Odds rule, specified by the value of odds and the fact that $\alpha = .5$, $p_1 = .2$ and $p_2 = .08$. The five sets of points are (from top to bottom) for the travel times, $\tau = 0, 5, 10, 15$ and 20. The relative insensitivity of the rate achieved to the rule used may be seen by comparing the rates (illustrated in Figure 3) with the rules (illustrated in Figure 2).

FIGURE 4. Expected gain E(G) plotted against expected time E(T) per patch for various Odds rules, for $\alpha = .5$, $p_1 = .2$ and $p_2 = .08$. If travel time, τ, is measured to the left of the origin, the line from the point $(-\tau, 0)$ through the point (E(T),E(G)) corresponding to the best Odds rule will have the steepest slope of all such lines.

FIGURE 1

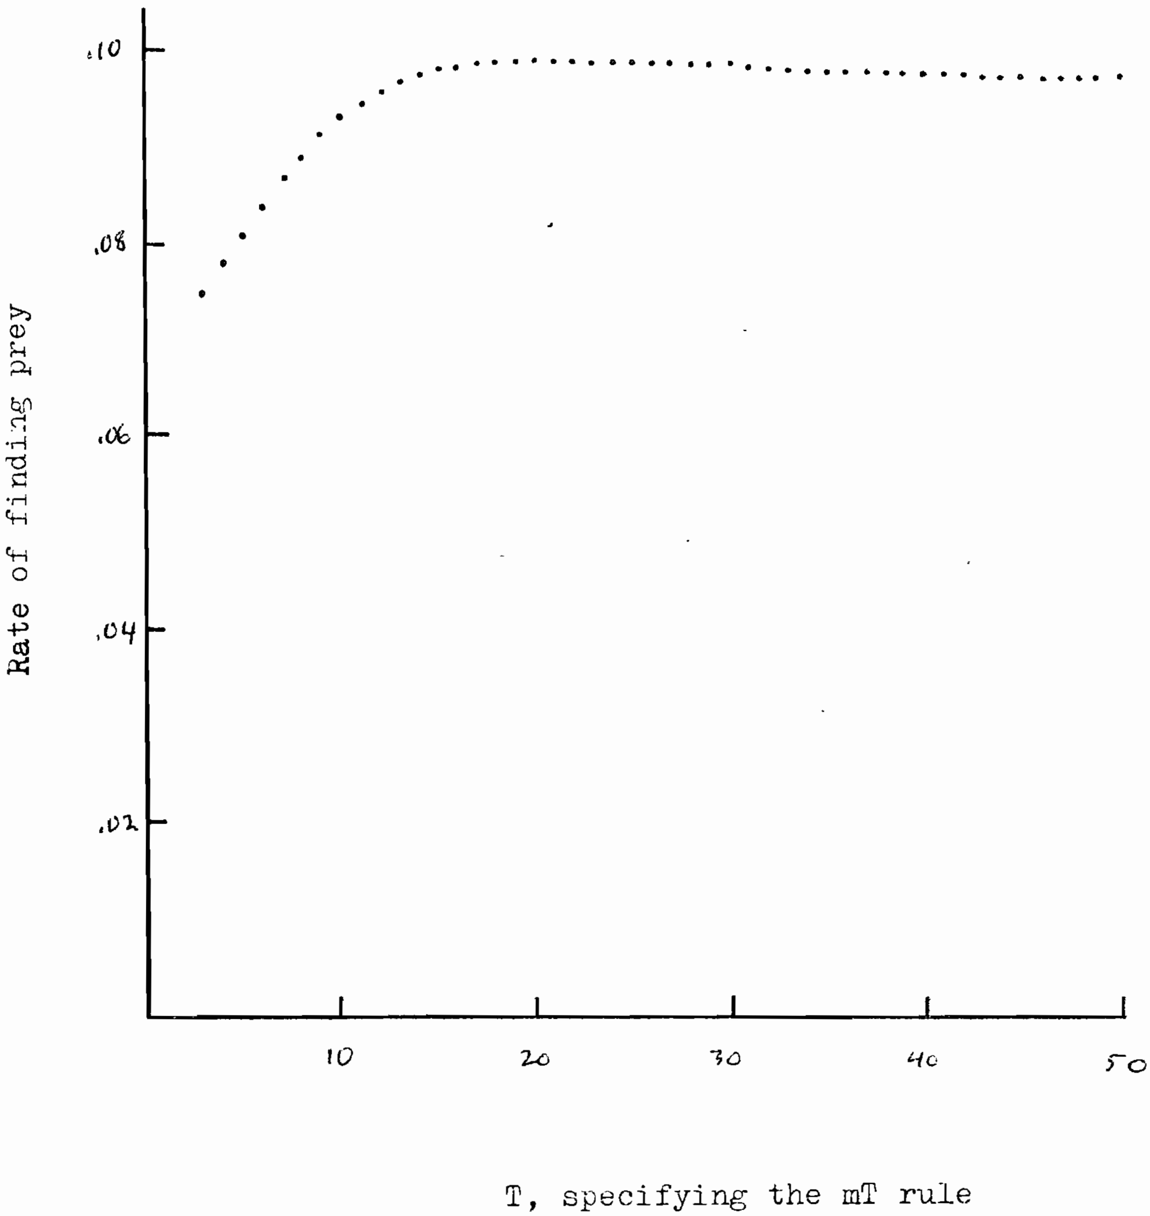
.10
.08
.06
.04
.02
10
20
30
40
50
Rate of finding prey
T, specifying the mT rule

FIGURE 2

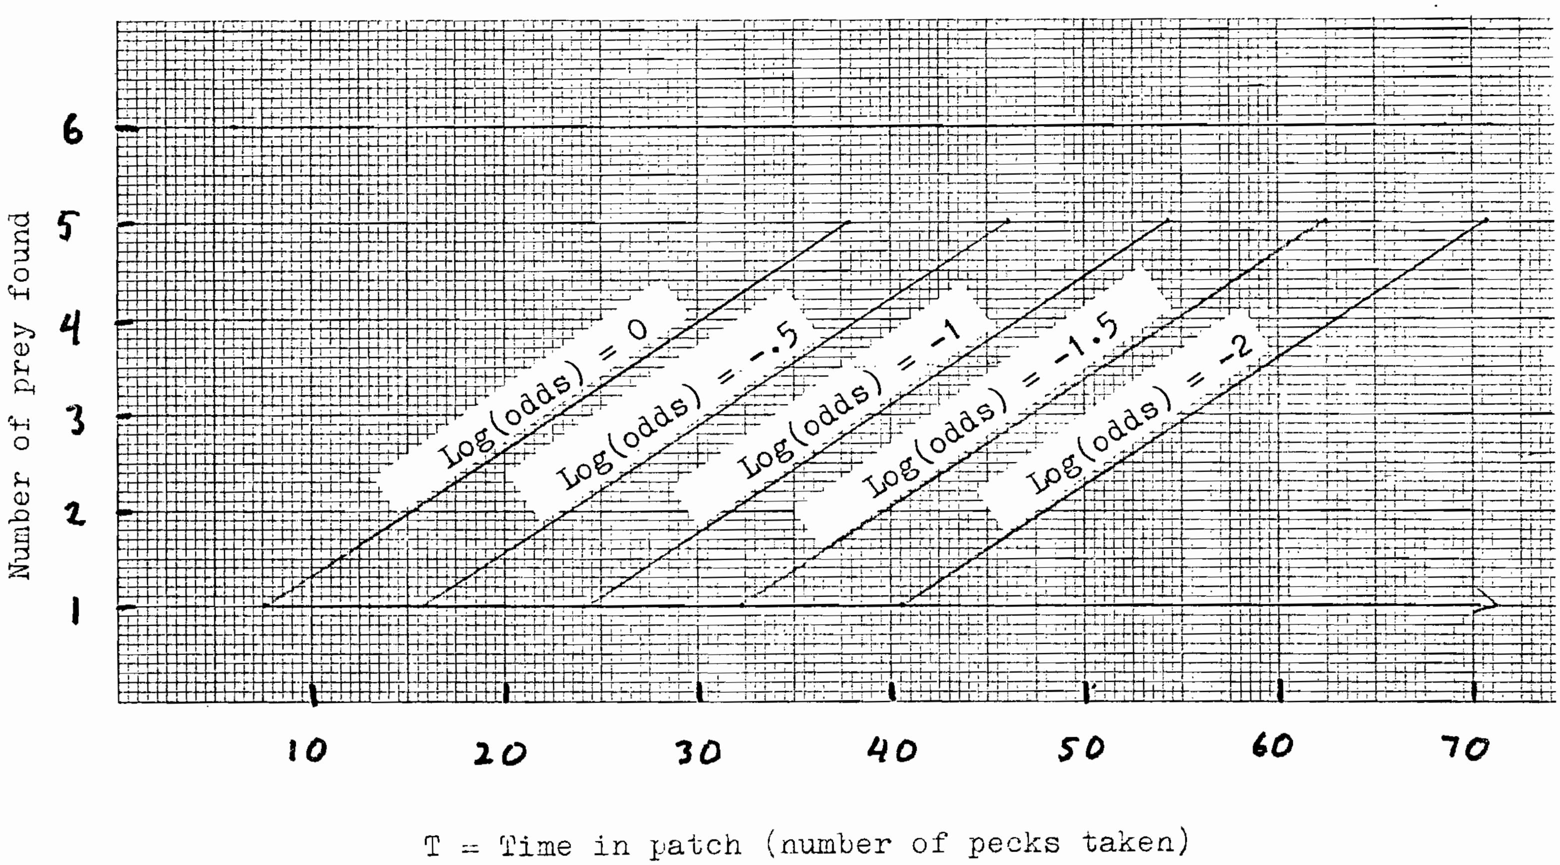

Number of prey found
6
5
4
3
2
1
Log(odds) = 0
Log(odds) = -.5
Log(odds) = -1
Log(odds) = -1.5
Log(odds) = -2
10
20
30
40
50
60
70
T = Time in patch (number of pecks taken)

FIGURE 3

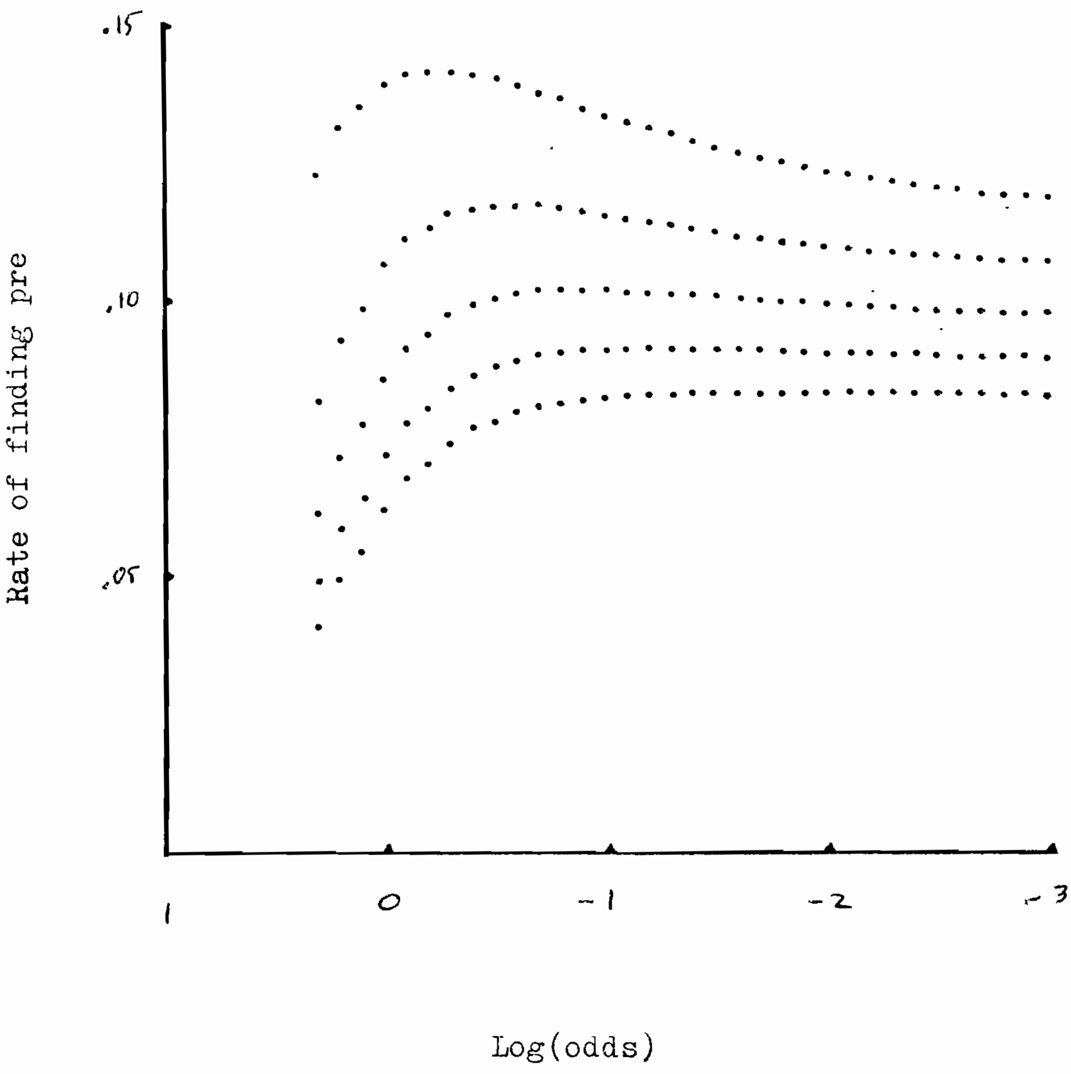

FIGURE 4