

Green, R. F. 1987. Stochastic models of optimal foraging. In: Foraging Behavior. (A. C. Kamil, J. R. Krebs, and H. R. Pulliam, eds.) Plenum Press, New York, pp. 273-302.

Abstract. This paper treats the problem faced by an animal that forages for prey that are distributed in discrete patches. The number of prey may vary from patch to patch and the forager must use its experience in a patch in order to decide when to leave one patch and go on to another. Several different models of prey distribution and foraging pattern are considered and the optimal foraging strategy is found for each. The optimal strategies have different forms for different prey distributions, but in each case the optimal strategy has a simple form. Optimal strategies are also robust in the sense that the optimal strategy in one case will also be quite good in different cases of the same type.

INTRODUCTION

For many years, one of the central concerns of natural history was the food habits of animals; that is, what food animals eat, and how they go about obtaining it. The question was, How do animals forage? More recently, optimal foraging theory has posed the question, How should animals forage? This question may be asked in many forms, and the answers have cast light on the old subject of animals' food habits. Optimal foraging theory has helped make the study of foraging more interesting, which may account for the theory's remarkable popular success noted by Krebs et al. (1983).

In their review of optimal foraging theory, Pyke et al. (1977) listed four problems a forager might face. (1) What prey to take. (2) What patch types to search. (3) When to leave a patch. (4) How to move between patches. While much of the empirical work on foraging involves the first two problems, I will concentrate on the third problem, which is the most interesting theoretically. Much less quantitative work has been done on the fourth problem.

There is not universal agreement on what a forager should optimize (see Breck 1978). In this paper optimal foraging is taken to mean maximizing the long-term average rate of finding prey. This (energy maximization) is the criterion most often used, and it would be the correct one if fitness were a linear function of energy intake over any foraging period. As Schoener (1971) pointed out, when food (or time available for foraging) is short, it is plausible to assume that animals should maximize their average rate of energy intake.

This paper is concerned with stochastic models, in which a forager must deal with unpredictable variability in some aspect of the environment. There are three levels at which this variability can occur. (1) Variability may occur within patches. For example, the time between captures within a patch may be random. Considering stochasticity at this level and ignoring patch structure of the environment (as Charnov 1976b and Stephens

and Charnov 1982 do) can lead to the same result as using the analogous deterministic model in which the forager obtains prey continuously according to some fixed function. (2) Variability may occur between patches. For example, the number of prey per patch may vary randomly from patch to patch. Oaten (1977) pointed out that such variability cannot be ignored by substituting an analogous model, and he emphasized the importance of this level of variability. It is this level of variability that will be considered in this paper. (3) Variability may occur over time. For example, the overall density of prey in the environment may vary from time to time, perhaps from season to season, leading to changes in prey distribution among patches. This level of variability has not been modeled, but it is certainly of practical and theoretical importance.

The words "risk" and "information" have been used to distinguish two kinds of stochastic foraging models (Stephens and Charnov 1982). The words are catchy, but they do not represent very accurately models that differ in the level at which environmental variability occurs, in the choices that foragers can make, and in the time period over which optimization takes place. "Risk" refers to what happens to an animal while foraging, that is, to the short-term variability in foraging success. Risk models generally involve variability within patches and ask whether foragers should prefer patches (environments) with greater or less variability, but with the same average quality. Interesting experimental work on such models has been done by Caraco et al. (1980) and Caraco (1981). "Information" refers to what an animal learns about the quality of a patch (or environment) from its experience foraging there. Information models involve variability between patches as well as within patches and ask how foragers should maximize their long-term average rate of finding prey. This paper is concerned with information models.

This paper treats the problem of how an animal that feeds on prey distributed in superficially similar patches that vary in prey density should decide when to leave one patch and go on to another. Charnov's (1976a) marginal value theorem specifies the best leaving rule for a deterministic model in which the net energy gained after searching for time s in a patch is a fixed, continuous function, $G(s)$, with a negative second derivative. The marginal value theorem, illustrated in Fig. 1, says that the forager should leave a patch when the rate of finding prey in a patch (the "marginal rate," or derivative of $G(s)$) falls to the highest possible long-term average rate of energy intake.

(Put Fig. 1 here.)

One of the attractions of the marginal value theorem is that Fig. 1 simultaneously illustrates the relationships among several variables. (1) G itself is a function of s . (2) The rate of finding prey at any time, s is the derivative, $G'(s)$, which is shown as the slope of G at s . (3) The long-term rate of finding prey achieved by a forager leaving each patch at time S , namely, $R = G(S)/(S + \tau)$, is the slope of the secant line from $(-\tau, 0)$ to $(S, G(S))$, where τ is the travel time between patches.

For stochastic models a single illustration will be insufficient to show all these relationships. Fig. 2 shows a particular success trajectory, with the number of prey found by time s , $G(s)$, plotted against s . This figure shows why the marginal value theorem does not work for discrete prey. The derivative (the marginal value) in Fig. 2 is almost always zero, and the marginal value theorem does not specify a leaving rule. Pyke (1978) suggested a stochastic, discrete form of the marginal value theorem that does specify a leaving rule, but the rule specified is not always optimal (Oaten 1977; Green 1980; McNamara and Houston in press). Pyke's rule is not based on an animal's

present rate of finding prey, but on its expected rate in the immediate future, based on its past experience in the patch.

(Put Fig. 2 here.)

In this paper I consider several particular cases of Oaten's (1977) general stochastic foraging model. I follow the approach of Iwasa et al. (1981), who consider several different distributions of the number of prey per patch. For each prey distribution I treat the case of systematic as well as that of random search by the forager. For each case I find the optimal leaving rule, using a method similar to that used earlier for a discrete-time case (Green 1980). Treating special cases lacks generality, but it permits calculations that reveal principles that otherwise would remain hidden.

THE MODEL

A general stochastic model for the problem of when a forager should leave a patch was given by Oaten (1977). Other work on this problem has followed Oaten's lead (Breck 1978; Green 1980; Iwasa et al. 1981; McNair 1982; McNamara 1982). In this paper I will describe several particular cases of Oaten's model (some of which have been treated by Breck (1978) and by Iwasa et al. (1981)), and I will find the optimal leaving rule for each case.

Oaten's model contains the following assumptions:

1. Prey are found in patches, and the predator knows the distribution of the number of prey per patch.
2. The predator knows the joint distribution of the capture times, given the number of prey in a patch.
3. Prey are not replaced as they are captured.
4. The predator knows the time τ that it takes to go from patch to patch. If travel time is variable, τ is its mean.
5. The predator decides when to leave a patch based on its knowledge of 1, 2 and 4, and its experience in a patch.
6. Given a leaving rule, we can calculate

$$R = E(G) / [E(S) + \tau] ,$$

where

- $E(G)$ = the expected number of prey caught in a patch, using the leaving rule;
 $E(S)$ = the expected length of time searching for prey in a patch, using the leaving rule; and
 R = the long-term average rate of finding prey, using the leaving rule.

7. The predator uses the leaving rule that maximizes R .

I consider six cases, each characterized by (1) the distribution of the number of prey per patch, and (2) the pattern of search used by the forager. I assume that

- (1) The distribution of the number of prey per patch is
 - a) regular (all patches have the same number of prey), or
 - b) Poisson, or
 - c) negative binomial.

These distributions do not exhaust the possibilities, but they do represent a range of cases, corresponding to prey having uniform, random, or aggregated distributions.

- (2) Prey are distributed at random in a patch, and search is
 - a) systematic, or
 - b) random.

By systematic search I mean that a forager does not retrace its steps within a patch. It does not necessarily mean complete search. Mathematically, systematic search means that the rate of finding prey is constant throughout a patch. By random search, I

mean that a given spot within a patch is equally likely to be searched whether or not it has been searched before. Mathematically, random search means that the rate of finding prey within a patch decreases exponentially with time in the patch. The assumption of random search is a mathematical one, under which the proportion of a patch which has been searched after a particular amount of time in the patch is not random, but is a deterministic function of the amount of time. Random search is mathematically equivalent to systematic search at an exponentially decreasing rate. There is a logical difficulty with this kind of random search; it is impossible if prey are sought at discrete locations within a patch (as in Green 1980). There is a good deal of evidence (Krebs 1973) that foragers' movements between patches are non-random. Lima (1984) observed systematic foraging within patches in his field study of downy woodpeckers, and Baum (this volume) observed systematic foraging within patches in his laboratory study of pigeons. Random search is a mathematically convenient, but biologically unrealistic, assumption. It is widely used in optimal foraging theory, perhaps because it produces the depletion required for Charnov's marginal value theorem (1976a) to hold. I have included random search, along with the more realistic systematic search, so that my results will be comparable with those of others (especially Iwasa et al. 1981), but I think that the almost universal assumption of random search and the corresponding experimental emphasis on depleting schedules is misguided. This emphasis may be an example of the heavy weight of a young tradition, of models already referred to as "classical" (Krebs et al. 1983), long before their implications have been thought out properly. This emphasis on patch depletion distracts attention from path variability, which may be more important, and is certainly more interesting, than patch depletion.

I assume that all patches are superficially similar and of the same size. The theory may be extended to the case of several distinct patch types (each with a different prey distribution), as Charnov (1976) did, but such an extension will provide no insight into the foraging problem which Charnov has not already provided. Without loss of generality it may be assumed that all patches are of unit size (that is, it takes unit time to search a patch systematically, and the proportion of a patch left unsearched after searching randomly for unit time is $\exp(-1)$).

For each of the cases I consider I find the best leaving rule, directly or indirectly, by considering the average rate of finding prey at a given time. Charnov's (1976) marginal value theorem, illustrated in Fig. 1, is based on the instantaneous rate (the "marginal value") of gaining energy at a given time. When individual prey are caught at discrete times, as I assume happens, Charnov's idea does not apply. An alternative is to calculate the average rate of finding prey at a given time, at which a given number of prey have been found. The marginal value theorem may be applied to this average instantaneous rate, which may be thought of as the average rate of finding prey during a short time interval after a given search time, with the rate found by averaging over many patches in which the given number of prey have been found by the given search time. This idea, suggested by Pyke (1978) and used by Iwasa et al. (1981), sometimes does provide the best leaving rule, although it does not do so in general.

(Put Fig. 3 here.)

Fig. 3 illustrates the average instantaneous rate of finding prey for animals experiencing the particular success trajectory illustrated in Fig. 2. The average instantaneous rate is plotted against time spent by a forager in a patch, for each of the six

cases considered in this paper. For three cases (regular distribution and Poisson distribution with random search, and Poisson distribution with systematic search) the plot of the instantaneous rate against time is monotone non-increasing. (This is true for any success trajectory.) In these cases the best leaving rules are determined by the average instantaneous rate.

For a Poisson distribution of prey in patches the average instantaneous rate depends on the time in a patch, but not on the number of prey found. In such a case the forager should remain for a fixed time in each patch (until the patch has been searched completely for systematic search). For a regular distribution of prey and random search the average instantaneous rate depends on the number of prey found, but not on the time spent searching. In such a case the forager should remain in a patch until a fixed number of prey have been found. This is Gibb's (1958; Krebs et al. 1974) "hunting by expectation." It is interesting to notice that there is a case in which "hunting by expectation" is optimal.

It is more difficult to find the best leaving rules in the other three cases, in which the average instantaneous rate sometimes rises and sometimes falls. In each case, the best rule may be found by dynamic programming, using a method similar to that used by Green (1980). I give examples of the best leaving rule in all three cases, but I concentrate of the case of negative binomial prey distribution and systematic search, which is most realistic biologically. (The less realistic case of negative binomial prey distribution and random search was considered by Breck (1978) and Iwasa et al. (1981), but they did not find the best leaving rule.)

Several ideas are used to find the best strategy. First, there is a highest rate of finding prey that can be achieved--call it C^* . There is a rule (not necessarily unique) that achieves this rate. At any time an animal foraging in a patch could leave that patch, travel to another patch, continue foraging, and achieve the rate C^* . Therefore, an animal foraging in a patch should remain in that patch if there exists a strategy such that, if that strategy is used, the ratio of expected gain to expected time during the remainder of the patch visit exceeds C^* . For systematic search and random prey distribution within a patch all the information about patch quality obtained up to a given time is contained in the number of prey captured, and the actual capture times are irrelevant. This provides an important simplification of Oaten's (1977) more general model and makes my calculations relatively easy. Unfortunately, the value of C^* is not known. In order to find C^* , one guesses a value of C , uses dynamic programming to find a strategy that "tries" to achieve that rate, then calculates the rate, $R(C)$, actually achieved by the strategy. If $R(C) = C$, then the best rate has been found. Otherwise another value of C must be chosen. In practice it is simplest to choose an original $C = \alpha / (\beta + \mu)$, find $R(C)$, then use the value of $R(C)$ as the new choice of C , and continue until $R(C) = C$. Very few iterations are necessary because the function $R(C)$ is very flat in the vicinity of C^* . More details are given in the Appendix.

The negative binomial distribution is often used to represent the number of animals (in particular, insects) in a given location (Southwood 1966; Taylor 1984). In this paper I follow Pielou (1977; but I use different notation; see DeGroot (1970) for more details) and treat the negative binomial distribution as a mixture of Poisson distributions. I assume that the (random) number of prey, X , in a given patch has a Poisson distribution with parameter λ ,

$$P(X = x) = e^{-\lambda} \lambda^x / x! ,$$

while the parameter λ itself is random, having a gamma distribution with parameters α and β . The parameter λ has the probability density

$$f(\lambda) = \beta^\alpha \lambda^{\alpha-1} e^{-\beta\lambda} / \Gamma(\alpha).$$

The number of prey per patch will then have a negative binomial distribution with

$$P(X = x) = \binom{\alpha+x-1}{x} \left[\beta/(1+\beta) \right]^\alpha \left[1/(1+\beta) \right]^x,$$

where $\binom{\alpha+x-1}{x} = \alpha(\alpha+1)(\alpha+2) \dots (\alpha+x-1)/x!$, the combinatorial coefficient.

RESULTS

Fig. 4 gives examples of optimal leaving rules for each of the six cases I have considered in this paper. Notice that the rules are very different in the different cases. That is, fundamentally different ecological problems require fundamentally different forms of optimal behavior. This means that it is impossible to determine whether foraging behavior approximates that which is optimal without knowing how prey are distributed among patches.

(Put Fig. 4 here.)

The qualitative differences may be understood in terms of the information obtained when each prey is found, as illustrated in Fig. 3. A negative binomial distribution of the number of prey within a patch means that prey are clumped, and the more prey that have been found after searching part of a patch, the more prey the forager expects to find in the remainder of the patch. In this case a forager should leave a patch if relatively few prey have been found. For a regular distribution of prey, the more prey that have been found after searching part of a patch, the fewer prey the forager expects to find in the remainder of the patch. In this case a forager should leave a patch if relatively many prey have been found. For a Poisson distribution of prey per patch, the number of prey that have been found after searching part of a patch is independent of the number of prey in the remainder of the patch. In this case a forager should ignore the number of prey caught and stay for some fixed time in each patch.

The efficiency of different rules

An advantage of treating particular cases is that it permits a quantitative comparison of various foraging strategies. For each particular model we can imagine several possible leaving rules. For a Poisson distribution of prey per patch the best leaving rule, which might be called the "fixed-time rule," ignores the forager's experience in each patch. For a negative binomial prey distribution, and for a regular prey distribution, the best rules do depend on the forager's experience in each patch. The advantage of using this experience may be measured by comparing the rates of finding prey achieved using rules that do and do not use experience in a patch.

Table 1 gives the long-term average rates of finding prey for a forager using the best leaving rule for an example of each of the six models considered in this paper. In each case, the average number of prey per patch is five, and the average travel time between patches is 0.1. The fixed-time rules, which are the best of all rules for the Poisson cases, will also be the best

rules--within the set of rules that ignore experience in a patch--for the cases of regular and negative binomial prey distribution.

(Put Table 1 here.)

For each pattern of search, the rate of finding prey achieved by the best fixed-time rule will be the same for all three prey distributions. Thus, for systematic search, the best fixed-time rule (stay in each patch until it has been searched completely) will achieve a rate of 4.55 prey per unit time for each of the three prey distributions. For random search the best fixed-time rule (stay 0.416 units of time in each patch) will achieve a rate of 3.30 prey per unit time for each prey distribution.

For negative binomial and regular prey distributions we can compare the rate of finding prey achieved by the best rule and the best fixed-time rule. In each case listed in Table 1, except the Poisson cases, the best rule does substantially better than the fixed-time rule. The largest difference shown in Table 1 occurs in the case of a negative binomial prey distribution and systematic search, where the best rule achieves a rate of 8.07 prey per unit time, which is 78% more than the rate achieved by the fixed-time rule. This percentage might be considered as the advantage of taking account of experience in a patch to decide when to leave.

It is important to realize that these values depend on the parameter values I have chosen for illustration, as well as the assumption that the animals might actually use the best rule. For a negative binomial prey distribution, the advantage of the best rule over a fixed-time rule will be greater for more variable patches (smaller σ) and shorter travel time (smaller τ), and will be less for less variable patches and longer travel time.

Another leaving rule which might be considered is the "giving-up time" (GUT) rule, in which a forager would leave a patch as soon as it goes a fixed time (the GUT) without finding a prey. Iwasa et al. (1981) show how to calculate the rate of finding prey achieved by the GUT rule when search is random. Table 1 gives the rates achieved by the best GUT rules for each prey distribution. (I was unable to find the best value of the GUT for systematic search.) Krebs et al. (1974) equated optimal foraging and the marginal value theorem, and they suggested that the GUT rule be used to satisfy the marginal value theorem. In none of the three cases for which it is shown in Table 1 is the GUT rule best. Furthermore, the best value for the GUT in each case is not the one suggested to satisfy the marginal value theorem, that is, the reciprocal of the rate achieved. In each case the value of the GUT that achieves the highest rate is less than the reciprocal of the rate achieved (best GUT values = 0.233, 0.236, and 0.190 for regular, Poisson and negative binomial distributions, respectively). It is difficult to think of an example for which the GUT rule is best, but McNamara (1982) did provide an artificial example. The GUT rule is a simple rule that is based on a forager's experience in a patch, but in two of the three cases for which it is shown in Table 1 the GUT rule performs worse than the fixed-time rule, which ignores the forager's experience in a patch. In the third case (negative binomial prey distribution), the GUT rule performs better than the fixed-time rule, but its advantage over the fixed-time rule is only about two-thirds that of the best rule, which is itself as simple as the GUT rule.

The simplicity of the best rules

The calculations needed to find the best leaving rules may be difficult, but the rules themselves are simple. This is important because originally there was the fear that the theory was so difficult that animals would be unable to approximate optimality. Oaten (1977) suggested that animals might be expected to use simple, suboptimal foraging strategies. Others, describing a similar idea, followed economists if referring to "rules of thumb" (Cowie and Krebs 1979; Janetos and Cole 1981). "Rules of thumb" is a catchy phrase which has been loaded with a baggage of meaning too heavy for it to bear. Krebs et al. (1983) have claimed that rule of thumb models refer to proximal mechanisms, while optimal foraging models refer to ultimate mechanisms, and that rule of thumb models and optimal foraging models can be viewed as representing different points on a continuum of strategy sets. Both these claims are incorrect: the rules of thumb usually considered, such as the GUT rule or the fixed-time rule, are simply descriptions of possible behavior and do not refer to proximal mechanisms, and such rules may actually be optimal in some cases. Theorists have concentrated on rules of thumb not so much because they are biologically realistic as because they are possibilities that spring readily to mind (Krebs et al. 1974), and because they are mathematically convenient (Breck 1978). The important point is not whether the optimal strategy is difficult for the biologist to determine, but whether the optimal strategy is simple enough for an animal to use.

An example of a simple optimal foraging strategy is that illustrated in Fig. 4 for the negative binomial prey distribution and systematic search. Such a rule could be produced by a behavioral mechanism similar to that suggested by Waage (1979). Imaging the forager as a wind-up toy. On entering a patch the toy is wound up a certain amount. The toy unwinds at a constant rate, except when it encounters a prey, in which case it is wound up a fixed amount more. The forager leaves the patch when the toy winds down or the end of the patch is reached. Equally simple mechanisms could produce the best rules for the other cases I have considered.

Robustness of the best rule

Different types of prey distributions require different forms of optimal strategies. For negative binomial prey distributions a forager should leave a patch when few prey are found, while for a regular distribution a forager should leave a patch when many prey are found. Prey distributions of the same type require strategies of the same general form, but the optimal rules can be quite different for prey distributions with different parameters. However, the rates achieved are not very sensitive to which rule is used as long as the rule has the correct general form.

When finding the best rule for a particular set of parameters a guess is made of the rate a rule might achieve, and a rule is found that "tries" to achieve that rate. The rule found may not be the best rule for the given parameters, but, unless the guess is impossibly high, the rule will be the best possible for the given parameters of prey distribution and some value of travel time. In Fig. 5 the rate achieved by various rules of the best type for a particular negative binomial distribution are plotted against C , the rate that the forager "tries" to achieve. Leaving rules corresponding to three of the values of C are given as insets. These rules are quantitatively very different from each other, but they achieve rates that are quite similar. (The middle rule is the best, while the other two rules achieve rates of about 90% of the best.)

(Put Fig. 5 here.)

Such robustness, which was seen in a similar model (Green 1984), applies to other changes in the parameters as well. A biologically important conclusion is that a forager using a strategy that is optimal for a particular situation may also achieve a nearly optimal rate for a different situation for which the optimal strategy is quite different. A theoretically important point is that this robustness would not have been noticed without a quantitative treatment of particular cases.

The idea of judging behavior by the payoff it achieves may cast light on other foraging problems. For example, one model of prey choice shows that a forager searching for two different kinds of prey of different quality should show exclusive preference for the better prey if its density is high enough, but if the density of better prey is below a particular threshold, both types of prey should be taken indiscriminately. Looking only at the form of the optimal strategy, we see a sudden change from indiscriminate choice to exclusive preference at one particular density of the better prey. But, looking at the relative payoffs from different strategies shows a gradual change as the density of the better prey changes. For densities of the better prey around the threshold, the differences among the payoffs for different rules are negligible. The fact that animals do not shift suddenly from no discrimination to exclusive preference (references tabled in Krebs et al. (1983)) is not surprising if one considers quantitative theory. A number of possibilities have been suggested to explain away the "partial preferences" observed in prey choice experiments (Krebs and McCleery 1984), but perhaps if the rates achieved by different strategies are looked at, then no further explanation is needed.

DISCUSSION

This paper discusses several particular quantitative models for one problem that a forager might face: how to decide when to leave one patch and go on to another. Studying particular versions of this problem sheds light on the general problem. First, considering a particular problem makes it possible to find the optimal strategy for that problem. Second, because one particular case is not enough, it is natural to consider other particular cases and find the optimal strategy for each. When this is done it is seen that different cases require very different optimal strategies. Third, in all the cases considered, the optimal strategies were found to be simple enough that animals might be expected to use them. This is important because people have suggested that optimal foraging rules might be too difficult for animals to use (Janetos and Cole 1981; Krebs et al. 1983). This suggestion should not be used as an excuse by the theorist to avoid searching for the optimal strategy. Fourth, foraging strategies tend to be robust in the sense that a wide range of behavior results in similar payoffs. This means that foraging behavior cannot be judged solely by how close it is to optimal behavior.

It has been suggested that there is too much theory for the number of observations in behavioral ecology. I say there is too little theory, at least too little good theory. In the area of foraging behavior the number of empirical studies are out of proportion to the number of theoretical ideas they test. There have been hundreds of studies of foraging behavior in the last decade, and most have tested a few qualitative predictions: animals should prefer the best prey, they should be more selective when prey are more abundant, and animals foraging in patches should remain longer in patches which contain more prey. Not

surprisingly, these predictions are usually confirmed.

A research program

Here I outline a research program based on quantitative foraging models designed to study particular problems. This program is for the study of foraging behavior, and the work should be informed by a knowledge of the ecological problem the behavior is to solve, a mathematical analysis of possible foraging strategies to solve the problem, and a method to interpret the results statistically. There are four different aspects of the larger foraging problem.

1. The ecological problem. First, it is necessary to determine what ecological problem a forager has to solve. We need to know how prey are distributed among patches and what foraging pattern an animal may use. (For example, can the animal forage systematically?). This information is difficult to obtain, but there is little point in talking about optimal foraging unless we have some idea of the problem that an animal has to solve. This aspect of the foraging problem is often ignored. One of the values of optimal foraging theory (especially Iwasa et al. (1981)) is that it demonstrates the importance of specifying the ecological problem.

2. The mathematical problem. Next, it is necessary to evaluate the performance of various possible strategies for particular ecological problems. For a given problem we find the optimal strategy and compare the rate of obtaining prey that it would achieve with the rates achieved by other strategies. By considering particular ecological problems we see that optimal foraging behavior is different in different cases. This cannot be seen with a general approach which does not distinguish particular cases. A good start toward understanding the general problem treated in this paper has been made by Oaten (1977), Breck (1978), Green (1980, 1984), Iwasa et al. (1981), McNair (1982) and McNamara (1982).

3. The behavioral problem. After the ecological problem has been specified we must observe what animals actually do. Detailed observations of foraging behavior in the field are difficult to make, so laboratory studies may be more fruitful. I think that well-designed laboratory experiments that simulate important features of the natural environment provide the best hope of advancing foraging theory.

4. The statistical problem. Finally, it is necessary to analyze the results of the experiment. To do this we must know what the possible strategies are and what behavior we would observe if an animal were to use a particular strategy. The statistical problem is a difficult one, since even describing the data is complicated. For a given patch visit, the time at which each prey is found and the leaving time must be recorded. It is also difficult to test hypotheses about foraging strategies, since it is unreasonable to expect animals to use a mathematically precise strategy. It is important to remember that a strategy should not be judged by how close it is to the optimal strategy, but rather by how close its payoff is to that of the optimal strategy. The study of robustness shows that a wide variety of behavior can produce very similar results.

Two experimental studies

Two important experimental studies can be used to illustrate these four aspects of the foraging problem. Krebs et al. (1974) studied the foraging behavior of black-capped chickadees by allowing them to forage for five minute periods in aviaries containing five artificial "trees," each containing three groups of four artificial cones, each with six holes in which

pieces of meal-worm larvae could be put. The ostensible purpose of the experiment was to test Gibb's (1958) suggestion that birds might "hunt by expectation," that is, they might leave a patch after a particular number of prey have been found there. To test this prediction the birds were trained for four days with one prey in each cone, then their behavior was recorded over two more days with the same prey distribution. Then the birds were tested for two days with prey distributed as before (one larva per cone) in two "trees," but with three larvae per cone in three "high-intensity" trees. The results reported showed that the birds took more prey in the "high-intensity" trees during the last part of the experiment than they did from "low-intensity" trees when there were no "high-intensity" trees. Krebs et al. (1974) therefore concluded that the birds were not hunting by expectation.

The bulk of Gibb's (1962) observations provided evidence that prey had an aggregated distribution, and that foragers took a larger proportion of prey from patches with more prey, as would be expected if foraging were optimal. The evidence on which hunting by expectation was suggested was very slight. Krebs et al. (1974) performed an experiment, in part of which foragers were presented with a regular distribution of prey, in which case "hunting by expectation" might have been expected, but the observations were not analyzed to see whether hunting by expectation actually occurred. The conclusion that birds do not hunt by expectation was reached by using data from an experiment in which prey had an aggregated distribution and hunting by expectation would not be expected.

Two alternative strategies were considered, a fixed-time rule and a giving-up time (GUT) rule. Another experiment gave evidence that the birds were not using the fixed-time rule, and Krebs et al. (1974) concluded that the birds were using the GUT rule. However, evidence whether the GUT rule was actually being used, while readily available, was not considered. An animal using the GUT rule would leave each patch after going a fixed time without finding a prey. Instead of asking whether the giving-up times were actually constant, Krebs et al. (1974) treated the average of the giving-up times, which cannot be used to discriminate the GUT rule from other rules. In order to investigate whether animals use a GUT rule it is necessary to look at the distribution of the giving-up times and not just their average.

The paper of Krebs et al. (1974) is extremely important in the history of optimal foraging theory because it suggested that foraging behavior could be studied in the laboratory. However, the experiments performed did not mimic the environmental problem usually faced by the animals tested, the mathematical theory was not properly developed to determine the optimal foraging strategy, and the statistical analysis of the data was not sufficiently detailed to determine what strategy the experimental animals were actually using.

In a more recent study, on downy woodpeckers, Lima (1984) tested a theory that was properly worked out and analyzed his data carefully, although it is not clear whether the model used mimics nature. Lima (1984) performed a field experiment in which woodpeckers were offered 60 artificial "trees," in each of which 24 holes had been drilled. After a training period in which each hole contained a "prey" (a piece of sunflower seed), birds were tested by offering a mixture of trees that were either "empty" (with no prey in any of the holes), or "full" (with prey in some or all of the holes). In the first experiment, "full" trees had prey in all 24 holes. In the second experiment, "full" trees had prey in exactly 12 holes, chosen at random, and in the third experiment, "full" trees had prey in exactly six holes, chosen at random. If we ignore the complication that an

optimal forager might count the number of prey found in a tree (and leave a tree after finding exactly 12 prey in the second experiment, say), then the optimal strategy could be described as: leave a tree if no prey are found in the first n holes searched; if a prey is found in any of the first n holes, search every hole in the tree.

The birds in Lima's experiment did, for the most part, search every hole in a tree once they had found a prey, and their behavior was in good qualitative agreement with the computed optimal strategy. There were two deviations from optimality, one relatively large and the other small, but both instructive. First, the number of holes searched unsuccessfully before leaving a tree was not constant for a given experiment, but, with one exception, the most frequent number of holes searched unsuccessfully was equal to the optimal number. For example, in the second experiment, the optimal number of empty holes to search before leaving a tree is exactly three, but sometimes birds leave a tree after searching only two empty holes, sometimes they search three empty holes, or four, or even more. We might view this variability in behavior as being relatively large, but the important thing is to look at its consequences. While searching exactly three holes unsuccessfully before leaving a tree may be optimal, alternative behavior, such as leaving after searching two or four empty holes will achieve approximately the same average rate of finding prey. Thus, relatively large deviations from optimal behavior may have relatively small consequences.

The second deviation from optimal behavior is that, in one case, the most frequent number of holes searched unsuccessfully before leaving a tree was one more than the optimal number. While this difference may be of little importance, it is interesting that it occurred in the case (the second experiment) in which a bird might be most likely to stay in an empty tree longer than the calculated optimum. The reason is that the optimal behavior is calculated to maximize the rate of finding prey during the experiment. But Lima's experiments were performed on free-ranging wild birds. It is unlikely that such birds would use the artificial trees unless they were at least as good as the natural environment. If the artificial environment is better than the natural environment, then the overall best strategy to use in the artificial environment may not be that calculated for the artificial environment alone. Of the two cases in which the forager had a difficult choice (experiments two and three), the difference between the natural and artificial environment would be greatest in the second experiment, and it was in this case that the modal number of empty holes searches was one more than the calculated optimum. While this discussion may be an over-analysis of a single observation, it does point out the important fact that in field experiments the experimental world we create is only part of the world in which an animal lives.

The use of optimal foraging theory

The recent review by Pyke (1984) leaves the impression that most work on foraging theory is concerned with testing whether animals forage optimally. Most of the tests are of a very few simple, qualitative models. One of the points of this paper is to recommend that people develop and test a variety of quantitative foraging models, but the purpose of theory is not simply to provide something to test. One may test whether animals forage optimally in some sense, but even if they do the questions remain: 1) How do they learn to forage optimally? and 2) What are the ecological consequences of optimal foraging?

Some early work (MacArthur (1972); Hassell and May (1974); Murdoch and Oaten (1975)) incorporated foraging theory into a larger ecological context, but the development of a self-

conscious optimal foraging theory (especially after Charnov (1976a,b)) has tended to limit ecology to providing a context for foraging behavior. There has been some consideration of, or mention of, the influence of foraging behavior on the functional response of a forager to prey density (Comins and Hassell (1979); Abrams (1982); Krebs et al. (1983)), but this work has not treated foraging theory very seriously. It may soon be time to try seriously to put optimal foraging theory back into ecology to help us understand such things as territoriality, resource partitioning and functional response.

With the increasing involvement of psychologists in foraging theory we might expect an increasing interest in whether animals learn to forage optimally, and if so, how they learn. The foraging problems that animals face in nature are more difficult than those usually presented in the psychologist's laboratory, but approximately optimal solutions may be simple enough for animals to learn. In this paper I have considered variability in the times at which prey are found within a patch and variability in the number of prey in different patches, but I have assumed that the forager "knows" the distribution of prey numbers and capture times with patches in a given environment. I have not treated the problem of how an animal "learns" these things or, equivalently, how an animal adjusts its behavior to changes in the environment in time or space. I have found the optimal foraging strategy in a fixed, model environment, but if we think that an animal is not just foraging, but also simultaneously monitoring the environment as a whole for changes, then we might expect to see what appears to be suboptimal behavior. One interesting finding of my study is that, for a given type of prey distribution, the negative binomial, say, the exact form of the stopping rule is quite sensitive to the environmental parameters, but the rate of finding prey is not particularly sensitive to the precise form of the rule used. This means that foragers that use a rule of the correct general form may lose little while they learn to forage optimally.

ACKNOWLEDGMENTS

I would like to thank Steve Lima, Adeline Nunez, Ron Pulliam, Allan Stewart-Oaten and an anonymous referee for their comments on this paper.

LITERATURE CITED

- Abrams, P. A. 1982. Functional responses of optimal foragers. *American Naturalist* 120:382-390.
- Baum, W. M. 1986. Random and systematic foraging, experimental studies of depletion, and schedules of reinforcement. This Volume.
- Breck, J. E. 1978. Suboptimal foraging strategies for a patchy environment. Dissertation. Michigan State University, East Lansing, Michigan, USA.
- Caraco, T. 1981. Energy budgets, risk and foraging preferences in dark-eyed juncos (*Junco hyemalis*). *Behavioral Ecology and Sociobiology* 8:213-217.
- Caraco, T., S. Martindale, and T. S. Whittam. 1980. An empirical demonstration of risk-sensitive foraging preferences. *Animal Behaviour* 28:820-830.
- Charnov, E. L. 1976a. Optimal foraging: the marginal value theorem. *Theoretical Population Biology* 9:129-136.

- Charnov, E. L. 1976b. Optimal foraging: attack strategy of a mantid. *American Naturalist* 110:141-151.
- Comins, H. N., and M. P. Hassell. 1979. The dynamics of optimally foraging predators and parasitoids. *Journal of Animal Ecology* 48:335-351.
- Cowie, R. J., and J. R. Krebs. 1979. Optimal foraging in patchy environments. Pages 183-205 in R. M. Anderson, B. D. Turner, and L. R. Taylor, editors. *The British Ecological Society Symposium, Volume 20, Population Dynamics*. Blackwell Scientific Publications, Oxford, UK.
- DeGroot, M. H. 1970. *Optimal Statistical Decisions*. McGraw-Hill, New York, USA.
- Gibb, J. A. 1958. Predation by tits and squirrels on the eucosmid *Ernarmonia conicolana* (Heyl.), *Journal of Animal Ecology* 27:375-396.
- Gibb, J. A. 1962. L. Tinbergen's hypothesis of specific search images. *Ibis* 104:106-111.
- 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.
- Hassell, M. P., and R. M. May. 1974. Aggregation of predators and insect parasites and its effects on stability. *Journal of Animal Ecology* 43:567-594.
- Iwasa, Y., M. Higashi, and N. Yamamura. 1981. Prey distribution as a factor determining the choice of optimal foraging strategy. *American Naturalist* 117:710-723.
- Janetos, A. C., and B. J. Cole. 1981. Imperfectly optimal animals. *Behavioral Ecology and Sociobiology* 9:203-210.
- Krebs, J. R. 1973. Behavioral aspects of predation. Chapter 3 in P. P. G. Bateson and P. H. Klopfer, eds., *Perspectives in Ethology*, Plenum Press.
- Krebs, J. R., J. C. Ryan, and E. L. Charnov. 1974. Hunting by expectation or optimal foraging? A study of patch use by chickadees. *Animal Behaviour* 22:953-964.
- Krebs, J. R., D. W. Stephens, and W. J. Sutherland. 1983. Perspectives in optimal foraging. Pages 165-216 in A. H. Brush and G. A. Clark, Jr. editors, *Perspectives in Ornithology*. Cambridge University Press, Cambridge.
- Lima, S. L. 1984. Downy woodpecker foraging behavior: efficient sampling in simple stochastic environments. *Ecology* 65:166-174.
- MacArthur, J. H. 1972. *Geographical Ecology*. Harper and Row, New York.
- 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.
- McNamara, J. M., and A. I. Houston. 1985. Foraging in patches: there's more to life than the marginal value theorem. In press.
- Murdoch, W. W., and A. Oaten. 1975. Predation and population stability. *Advances in Ecological Research* 9:2-131.
- Oaten, A. 1977. Optimal foraging in patches: a case for stochasticity. *Theoretical Population Biology* 12:263-285.
- Pielou, E. C. 1977. *Mathematical Ecology*. Wiley-Interscience, New York, USA.
- Pyke, G. H. 1978. Optimal foraging in hummingbirds: testing the marginal value theorem. *American Zoologist* 18:739-752.
- Pyke, G. H. 1984. Optimal foraging theory: a critical review. *Annual Review of Ecology and Systematics* 15:523-575.
- Pyke, G. H., H. R. Pulliam, and E. L. Charnov. 1977. Optimal foraging: a selective review of theory and tests. *Quarterly Review of Biology* 52:137-154.
- Schoener, T. W. 1971. Theory of feeding strategies. *Annual Review of Ecology and Systematics* 2:369-404.
- Southwood, T. R. E. 1966. *Ecological Methods*. Chapman and Hall, London.
- Stephens, D. W., and E. L. Charnov. 1982. Optimal foraging: some simple stochastic models. *Behavioral Ecology and Sociobiology* 10:251-263.
- Taylor, L. R. 1984. Assessing and interpreting the spatial distribution of insect populations. *Annual Review of Entomology* 29:321-357.
- Waage, J. K. 1979. Foraging for patchily-distributed hosts by the parasitoid, *Nemeritis canescens*. *Journal of Animal Ecology* 48:353-371.

APPENDIX

The optimal strategy for a systematic forager when the number of prey per patch has a negative binomial distribution

Consider the problem faced by a forager searching for prey distributed in superficially similar patches of the same size, where size is measured in terms of the time required to search each patch. Without loss of generality we may assume that size = 1. Each patch encountered is searched systematically until it has been searched completely ($t = 1$), or until the forager decides to leave and go on to another patch. Within a given patch the prey are distributed at random, and their number has a Poisson distribution, with parameter λ , which itself is a random variable, having a gamma distribution with parameters α and β . Thus the distribution of the number of prey in a patch will have a negative binomial distribution. The average travel time between patches is τ .

This model has three parameters: α , β and τ . The problem is to find the optimal strategy and to determine the long-term average rate of finding prey that it achieves. This problem is simplified by the fact that prey are assumed to be distributed at random within a patch. This means that the number of prey discovered up to a certain time in a patch is a sufficient statistic to estimate λ in that patch. That is, the times at which the prey are found are irrelevant to the question of whether a patch is good or not, and the optimal strategy depends only on the number of prey discovered up to a given time. The optimal foraging strategy is a stopping rule of the form: leave a patch and go on to another patch if no prey have been found by time $t(0)$, or if only one prey has been found by time $t(1)$, and so on, up to n prey found by time $t(n)$. If $n+1$ or more prey are found the forager should remain in the patch to the end. The problem is to find n and the stopping times: $t(0)$, $t(1)$, $t(2)$, ..., $t(n)$.

A similar problem has been solved for a case in which patches consist of discrete bits in which prey may be found (Green 1980). The fact that search time is continuous adds slightly to the difficulty of the calculations in the present case, and the optimal rule is found using numerical methods.

The optimal strategy is one that maximizes the long-term average rate of finding prey, given by

$$R = EG / (ET + \tau), \quad (A1)$$

where EG = the expected number of prey found in each patch visited, using a given strategy, and ET = the expected time spent searching each patch visited. The best stopping rule is found by dynamic programming. First, a guess, C , is made at the maximum value of R , then a stopping rule is found that would achieve the rate C , if, in fact, C were the highest possible rate. Then $R(C)$, the rate actually achieved by the rule found by guessing C , is found. If $R(C) = C$, the highest possible rate has been found, as well as a rule that achieves it. Otherwise, another value of C is chosen, and a new $R(C)$ if found. In practice, the old value of $R(C)$ is a good choice for the new value of C . The process is continued until a solution to the equation $R(C) = C$ has been found.

The basic work of finding the best strategy consists of two parts. The first, and more difficult, is to find the rule that "tries to achieve" the guessed rate, C . The second part is to calculate the rate actually achieved by the rule found in the first part.

Part 1.

Once C has been chosen ($C = \alpha / (\beta + \beta T)$ is a good first choice), the first problem is to find n , the largest number of prey such that the forager would leave the patch before the end even if as many as n prey had been found. (If enough prey are found-- $n+1$ or more--the evidence is that the patch is good enough that it should be searched completely.)

The value of n . If x prey have been found at time t , the expected instantaneous rate of finding prey is

$$r(x,t) = (\alpha + x) / (\beta + t). \quad (A2)$$

Just before the patch has been exhausted, $t = 1$, and the expected rate would be almost $r(x,1) = (\alpha + x) / (\beta + 1)$. If this value exceeds C the forager should stay in the patch for the last instant. The number n will be the largest integer such that

$$r(n,1) \leq C. \quad (A3)$$

The value of $t(n)$. To find the stopping time, $t(n)$ (the idea will be similar for stopping times $t(k)$, for $k = 0, 1, 2, \dots, n-1$, as well) we imagine an animal that has found $k = n$ prey at time t and that decides to remain in the patch for a small additional time, t . If no prey are found during that time, the animal will leave; if a prey is found, the animal will stay until it hits one of the later stopping points (in the case $k = n$, the forager will stay until the patch has been searched completely). The ratio of expected additional prey found (EG') to expected additional time (ET') in the patch can be calculated. When this ratio equals C the stopping time, $t(n)$, has been found. The equality

$$EG' / ET' = C \quad (A4)$$

may be solved numerically, using the bifurcation method.

For $k = n$ and time = t we have

$$EG' = \Delta t r(n,t) [1 + (1-t) r(n+1,t)], \quad (A5)$$

$$ET' = \Delta t [1 + r(n,t) (1-t)]. \quad (A6)$$

The value of $t(n-1)$. As above, we numerically solve the equation (A4) for t . Here we have

$$EG' = \Delta t r(n-1,t) [1 + (1-t) r(n,t) - w(n,t,0) r(n,t(n)) (1-t(n))], \quad (A7)$$

$$ET' = \Delta t [1 + r(n-1,t) [(1-t) - w(n,t,0) (1-t(n))]], \quad (A8)$$

where

$$w(n,t,0) = \{[\beta + t] / [\beta + t(n)]\}^{\alpha + n}, \quad (A9)$$

the probability that a forager, having found n prey at time t , will find no more additional prey if it remains in the patch until time $t(n)$.

The value of $t(k)$, for $k = 0, 1, 2, \dots, n-2$. Having found $t(n)$ and $t(n-1)$, above, we find $t(n-2)$, $t(n-3)$, \dots , $t(0)$, successively, by solving (A4). Each stopping time, $t(k)$, is found by using the values found for $t(k+1)$, $t(k+2)$, \dots , $t(n)$. We have

$$EG' = \Delta t r(k, n) (1 + r(k+1, t) (1-t) - \sum_{i=0}^{n-k-1} w(k+1, t, i) r[k+i+1, t(k+i+1)] [1-t(k+i+1)]), \quad (A10)$$

$$ET' = \Delta t (1 + r(k, t) [(1-t) - \sum_{i=0}^{n-k-1} w(k+1, t, i) (1 - t(k+i+1))]), \quad (A11)$$

where $w(k+1, t, i)$ is the probability that a forager that has found $k+1$ prey at time t will stop at time $t(k+i+1)$, having found i additional prey.

To find the values of $w(k+1, t, i)$ we define, for $0 \leq i < j \leq n$,

$$p(i, j-i) = \{ [\beta + t(i)] / [\beta + t(j)] \}^{\alpha+i} \{ (\alpha+j-i-1)! / [(\alpha-1)!(j-i)!] \} \{ [t(j) - t(i)] / [\beta + t(j)] \}^{j-i}, \quad (A12)$$

the probability that there are $j-i$ prey in the interval $[t(i), t(j)]$, given i prey in the interval $[0, t(i)]$, and for $0 \leq i \leq j \leq n$,

$$v(i, t, j-i) = \{ [\beta + t] / [\beta + t(j)] \}^{\alpha+i} \{ (\alpha+j-i-1)! / [(\alpha-1)!(j-i)!] \} \{ [t(j) - t] / [\beta + t(j)] \}^{j-i}, \quad (A13)$$

the probability that there are $j-i$ prey in the interval $[t, t(j)]$, given i prey in the interval $(0, t)$.

Then, for $0 \leq k \leq n-2$, we have

$$w(k, t, 0) = v(k, t, 0), \quad (A14)$$

$$w(k, t, j-i) = v(k, t, j-i) - \sum_{l=i}^{j-1} w(k, t, l-i) p(l, j). \quad (A15)$$

Part 2.

Once n and the stopping times $t(0), t(1), \dots, t(n)$ have been found for a given choice of C , the rate achieved by the rule is calculated using ideas similar to those used in Part 1. Imagine a forager arriving at a new patch, and define

$$v(0, 0, j) = \{ [\beta] / [\beta + t(j)] \}^{\alpha} \{ (\alpha+j-1)! / [(\alpha-1)!j!] \} \{ [t(j)] / [\beta + t(j)] \}^j, \quad (A16)$$

the probability that there are j prey in the interval $[0, t(j)]$. Then, using the definitions of $w(k, t, i)$ and $p(i, j-i)$ given in (A12), (A14) and (A15), we have

$$EG = r(0, 0) - \sum_{i=0}^n w(0, 0, i) r[i, t(i)] [1 - t(i)], \quad (A17)$$

$$ET = 1 - \sum_{i=0}^n w(0, 0, i) [1 - t(i)]. \quad (A18)$$

The long-term average rate of finding prey achieved by the stopping rule determined by choosing the value C is given by (A1), using the values calculated from (A17) and (A18). The best rate will be achieved when $R(C) = C$. In practice, such a C can be found by first choosing $C = \alpha / (\beta + \beta t)$, finding $R(C)$, using this value as the next choice of C , and proceeding iteratively until $R(C) = C$. The number of such steps required is small.

green

TABLE 1. Long-term average rate of finding prey, calculated for two different patterns of search and three different distributions of prey in patches. In each case the average number of prey per patch is five, the search is at the rate of one patch per unit time, and average travel time between patches is 0.1.

Systematic search

Prey distribution	Rule used	
	Fixed time	Best
Regular	4.55	5.68
Poisson	4.55	4.55
Negative binomial*	4.55	8.07

Random search

Prey distribution	Rule used		
	Fixed time	Best	GUT
Regular	3.30	3.64	2.99
Poisson	3.30	3.30	3.19
Negative binomial*	3.30	5.61+	4.81

*The negative binomial distribution used here has (gamma) parameters $\alpha = 0.5$ and $\beta = 0.1$, as described in the text. The rate achieved by the best rule for the negative binomial distribution and random search is an approximation.

green

FIGURE LEGENDS

FIG. 1. Energy gained (G) per patch visit, plotted against time stayed in a patch (S). The slope of the G curve at any time is the rate of gaining energy at that time (the "marginal value"). The highest possible long-term average rate of gaining energy, given by $R = G/(S + \bar{T})$, is the slope of the tangent line from $(-\bar{T}, 0)$ to the G curve. The travel time between patches is \bar{T} .

FIG. 2. A success trajectory. The number of prey found by a certain time in a patch is plotted against time in a patch. In the example illustrated, prey are found at times 0.05, 0.2, 0.4, 0.7 and 0.9.

FIG. 3. Expected instantaneous rate of finding prey, for three different prey distributions and two different modes of search. The dotted lines show the expected instantaneous rates plotted against time in a patch, for various numbers of prey found, while the solid line shows the expected instantaneous rate for the particular example of a forager that finds prey, as illustrated in Fig. 2, at times 0.05, 0.2, 0.4, 0.7 and 0.9. Notice that finding a prey decreases the expected rate for a regular distribution, increases it for a negative binomial distribution, and has no effect for a Poisson distribution.

FIG. 4. Optimal stopping rules, for three different prey distributions and two different modes of search. An optimal forager should leave a patch when its success trajectory hits one of the points (or lines) shown in the figure. The average number of prey per patch is five in each case, and the average travel time between patches is $\bar{T} = 0.1$. The shape parameter for the negative binomial distribution is $\alpha = 0.5$.

FIG. 5. Robustness of stopping rules of the best type, for negative binomial prey distribution with $\alpha = 0.5$, $\beta = 0.1$ and $\bar{T} = 0.1$. The rates achieved by various rules are plotted against C , the rates that the foragers "try" to achieve. The three rules illustrated in the insets, for foragers which "try" to achieve rates 4.5, 8.07 (the optimum), or 11, show how different rules can be that achieve rates within 90% of the highest possible.

Fig 1




Fig 2






Fig 3



SYSTEMATIC SEARCH

RANDOM SEARCH

REGULAR



POISSON



NEGATIVE
BINOMIAL







Fig 4


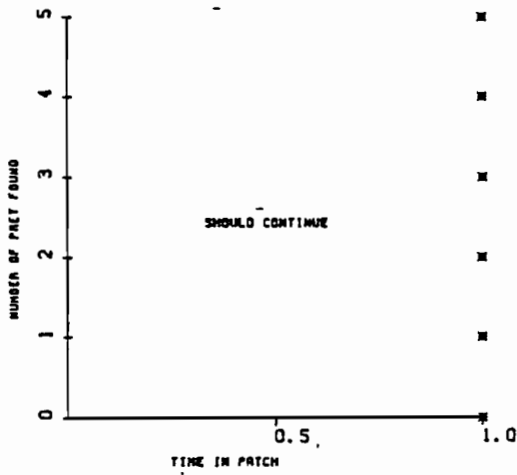
SYSTEMATIC SEARCH

RANDOM SEARCH

REGULAR



POISSON



NEGATIVE
BINOMIAL

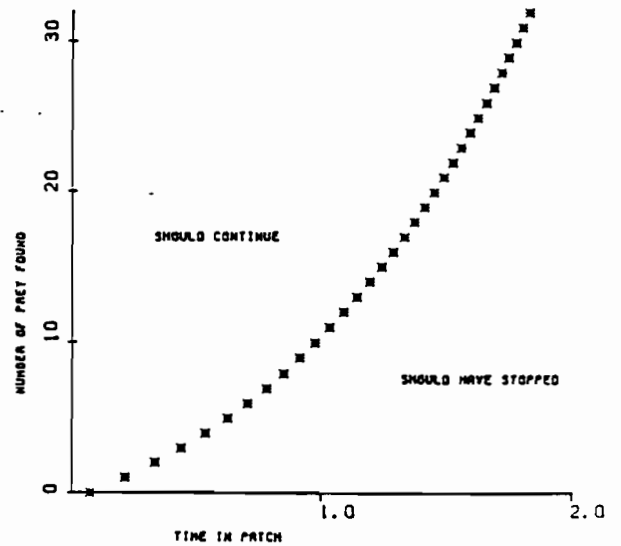
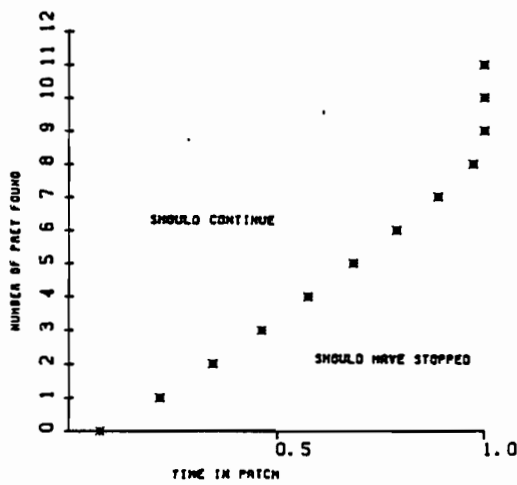


Fig 5

