BAYESIAN BIRDS: A SIMPLE EXAMPLE OF OATEN'S
STOCHASTIC MODEL OF OPTIMAL FORAGING

BY RICHARD F. GREEN

Technical Report No. 50

Department of Statistics
University of California
Riverside, California 92521

February, 1979

BAYESIAN BIRDS: A SIMPLE EXAMPLE OF OATEN'S
STOCHASTIC MODEL OF OPTIMAL FORAGING

Richard F. Green
Department of Statistics
University of California
Riverside, California 92521

ABSTRACT

Allan Oaten (1977) has argued that stochastic models of optimal foraging may produce results qualitatively different from those of the analogous deterministic models. Oaten's model is very general and difficult to understand intuitively. In this paper a simple, tractable model is considered in which the predator searches each patch systematically (without going over the same area twice) until he exhausts the patch or decides the patch is not very good. It is assumed that each patch contains a fixed number of bits, each of which may contain a prey. The number of prey per patch is assumed to have a binomial distribution with n equal to the number of bits and p being a random variable having a Beta distribution. After searching each bit the predator decides whether to leave the patch or not according to how many prey it has found. In this paper the best strategy is determined and the long-term rate of feeding is compared with that of the naive animal that searches each patch completely. The advantage of being a Bayesian is determined for a variety of environmental conditions.

BAYESIAN BIRDS: A SIMPLE EXAMPLE OF OATEN'S
STOCHASTIC MODEL OF OPTIMAL FORAGING

INTRODUCTION

Animal behaviorists and ecologists have been interested in the problem of predation for a number of years. The question is, how do predators operate? What prey do they hunt, and how and where do they hunt it?

More recently biologists have begun to think about "optimal foraging." That is, how should predators operate if they are to be efficient? The question is one of optimization, but it is not immediately clear what is to be maximized or minimized. Biologists are quick to say that what is maximized is fitness. The currently fashionable slogan is that "natural selection tends to maximize inclusive fitness." This is no help because it is not known exactly how different patterns of foraging affect fitness.

The assumption usually made is that a predator should maximize its net rate of energy intake while foraging. This assumption and many aspects of foraging are discussed by Schoener (1971). Other more recent and more specialized reviews are by Krebs (1974) and Pyke, et al. (1977). The assumption that the rate of net energy intake should be maximized is reasonable if either (1) feeding time is strictly limited, as it is for songbirds that winter in the north where days are short, or if (2) there is more time available for feeding than is needed but the predator is itself in danger from other predators while it forages. This could be true for songbirds during the late summer when they are no longer feeding young.

In order to maximize energy intake the predator must choose the correct prey and must search in an efficient way. In this paper I will ignore prey choice and will consider only the way the predator searches for its prey. The problem I will consider involves a predator that searches for prey that is distributed in discrete patches. The problem that faces the predator is how to decide when to leave one patch and go on to another.

Eric Charnov (1973, 1976) has considered this problem and assumes that the predator depletes the patch as it forages. Charnov concludes that the optimal forager should leave each patch when the rate of energy gain from feeding in that patch falls to the highest rate that can be sustained for a long time. Charnov, with economics in mind, refers to this as his "marginal value theorem." Charnov's result is for a deterministic model but he asserts that "qualitatively similar results may be shown to follow from a corresponding stochastic model".

This is an important point. Most models used by biologists are deterministic. Since the phenomena being modelled are usually stochastic it is good to know whether introducing a stochastic element changes the conclusions.

Allan Oaten (1977) has shown that, in fact, a stochastic version of Charnov's model can produce qualitatively different results. Oaten constructs an example in which a predator that takes account of randomness by assessing the quality of each patch it enters can do arbitrarily much better than a predator that uses the strategy suggested by Charnov. Unfortunately, the particular example used by Oaten is an unlikely special case and his approach is mathematically too general to be tractable.

In this paper I will describe a tractable and biologically plausible case of Oaten's model and will mention a few conclusions. Throughout this paper the predator will be described as a well-informed, rational strategist capable of fairly sophisticated calculations. The result of the calculations will show an optimal strategy which is quite simple. The predator doesn't have to perform the calculations itself. All that it has to be able to do is to follow the simple strategy.

THE MODEL

In his model Oaten assumes that:

- 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.
- 5) The predator decides when to leave a patch based on its knowledge of 1), 2) and 4) and its experience in the patch.
- 6) Given a strategy we can calculate

$$R = E(G)/(E(S) + \tau) \quad (1)$$

where

$E(G)$ = the expected number of prey caught in a random patch,

$E(S)$ = the expected length of time in each patch, and

R = the long term rate of gain using the given strategy.

- 7) The predator uses the strategy that maximizes R .

I make the same assumptions but confine myself to the case that each patch contains exactly n bits -- places where prey may be found -- and the number of bits containing prey has a binomial distribution with parameters n and p . The predator takes unit time to search each bit. It is assumed that the quality of patches varies and that this is reflected in the assumption that p has the beta distribution with parameters α and β ($p \sim B(\alpha, \beta)$). This is all "known" to the predator.

I assume that the probability of finding prey in each bit searched does not change as a patch is searched. This is equivalent to saying that the patch is searched systematically in the sense that the predator does not go over the same part of a patch twice during the same visit. This certainly is plausible for certain predators like treecreepers and woodpeckers that follow a fairly regular search pattern.

The optimal foraging strategy for the Bayesian bird consists of using a stopping rule that tells the bird when to leave a patch. Any stopping rule will consist of a function t_k , $k = 0, 1, 2, \dots, n$ such that the bird will leave a patch if after examining t_k bits it has found only k prey.

FINDING THE BEST STRATEGY

Since there are only finitely many strategies there must exist one that produces the maximum rate $R = c^*$. Of course, c^* is not known. What we do to find the best strategy is to pick some value, c , and find the strategy that would be used if that c is the correct one. Then, with that strategy we find the actual value of R and compare it with c . If they are the same then we have found the correct value of c . If the value found for R is greater than the c chosen then we try a larger c . If the value found for R is smaller than c we try a smaller c .

The general idea of the stopping rule is this: if at any time the bird has found so few prey in a patch that it knows that it will probably not do as well staying longer in that patch as it knows it can do in the long run then it should leave that patch and get on with the search. We find the best rule by starting at the end of the patch and working toward the beginning. This method of solving a problem by backward recurrence is described in a number of books on dynamic programming, for example, one by Nemhauser (1966). Another recent application of dynamic programming in this field is the work of Krebs, Kacelnik and Taylor (1978) who consider a "two-armed bandit" model for the choice of the best feeding place by foraging birds.

Finding the Best Strategy, Given c

We start by picking some constant, c , as our guess of the maximum rate obtainable. The chosen c will determine a strategy which will satisfy a condition which must be satisfied by the best strategy if c is the maximum possible rate. The strategy determined by c need not be the best unless the value of c used is actually c^* , the maximum rate possible.

Use the notation:

$x(t)$ = the number of prey found in the first t bits searched in a patch,
and

$p(x,t)$ = the probability that the next bit searched will contain a prey
given that x prey have been found in the first t bits.

It is well known that $p = (\alpha + x)/(\alpha + \beta + t)$. (See, for example, DeGroot, 1970, p. 160.) Thus, when all but the last bit have been searched the probability that the n^{th} (and last) bit will contain a prey is

$(\alpha + x(n-1))/(\alpha + \beta + n-1)$. If this probability is greater than c then the bird should go on to the n^{th} bit, otherwise it should leave the patch and go on to the next patch. Thus, if the bird has found enough prey (if $x(n-1) > s(n-1)$ for some critical value $s(n-1)$) it should continue in the patch. If it has not found enough prey it should leave the patch. This tells the bird what to do if it reaches the last bit.

Now go back to two bits before the end of the patch. If the bird has found $x(n-2)$ prey in the first $n-2$ bits searched then the probability that it would find another prey if it searched the next bit is $(\alpha + x(n-2))/(\alpha + \beta + n-2)$. If the bird does search the $n-1^{\text{st}}$ bit then it already has a rule that tells it whether to search the last bit when it gets to it. The bird decides to go on and search the next bit if the ratio of the expected number of prey gained to the expected time spent until finishing or leaving the patch is greater than c . Using the rule that tries to achieve rate c building up from the end, we have the following notation:

$EG(x,t,c)$ = the expected number of prey found after time t if x prey have been found by time t .

$ES(x,t,c)$ = the expected time in the patch after time t if x prey have been found by time t .

$EG'(x,t,c)$ = the expected number of prey found after time t if x prey have been found by time t and the predator decides to go on (and search the $t+1^{\text{st}}$ bit).

$ES'(x,t,c)$ = the expected time in the patch after time t if x prey have been found by time t and the predator decides to go on.

We have then,

$$\begin{aligned} EG'(x, n-2, c) = & p(x, n-2)(1 + EG(x+1, n-1, c)) \\ & + (1 - p(x, n-2)) EG(x, n-1, c), \text{ and} \end{aligned} \quad (2)$$

$$\begin{aligned} ES'(x, n-2, c) = & 1 + p(x, n-2) ES(x+1, n-1, c) \\ & + (1 - p(x, n-2)) ES(x, n-1, c). \end{aligned} \quad (3)$$

The strategy will then be to continue in the patch if $EG'(x, n-2, c)/ES'(x, n-2, c) > c$ and to leave the patch otherwise.

This method is then used for when the bird has searched $n-3$ bits and so on. In general we have

$$\begin{aligned} EG'(x, n-i, c) = & p(x, n-i)(1 + EG(x+1, n-i+1, c)) \\ & + (1 - p(x, n-i)) EG(x, n-i+1, c), \text{ and} \end{aligned} \quad (4)$$

$$\begin{aligned} ES'(x, n-i, c) = & 1 + p(x, n-i) ES(x+1, n-i+1, c) \\ & + (1 - p(x, n-i)) ES(x, n-i+1, c) \end{aligned} \quad (5)$$

and the strategy must satisfy

Rule 1: After $n-i$ bits have been searched and x prey found

a) search the next bit if

$$EG'(x, n-i, c)/ES'(x, n-i, c) > c, \quad (6)$$

b) leave the patch if

$$EG'(x, n-i, c)/ES'(x, n-i, c) < c, \text{ and}$$

c) either search the next bit or leave the patch if

$$EG'(x, n-i, c)/ES'(x, n-i, c) = c.$$

I arbitrarily choose to have the predator leave the patch when $EG'(x, n-i, c)/ES'(x, n-i, c) = 0$.

That the best stopping rule must satisfy Rule 1 for all $i=1,2,\dots,n$ may be seen by considering any strategy which does not satisfy Rule 1 for some value of i and achieves some rate c . Then this strategy may be improved by changing it at bit $n-i$ to make it satisfy Rule 1, either by having the predator remain in the patch if $EG'(x,n-i,c)/ES'(x,n-i,c) > c$ or leaving it $EG'(x,n-i,c)/ES'(x,n-i,c) < c$.

If the chosen constant c is the highest rate of finding prey possible, then the rule found in this way will be the best. If the value c is lower than the best possible rate R , then the bird using a strategy based on the assumption c is the best rate will be too reluctant to leave a patch, while if c is higher than the best possible R , then the bird will be too anxious to leave a patch. The actual rate R can be calculated for each rule and compared with c . When the c chosen and the R actually achieved are the same then the best stopping rule will have been found.

Finding the Rate R , Given the Strategy

When the stopping rule has been determined for a given value c the actual rate R may be found if we know the probability of each possible outcome. We can characterize the possible outcomes by a pair (k, t_k) , for $k = 0, 1, \dots, n$, where k is the number of prey found and t_k is the number of bits searched in finding the k prey. The predator leaves the patch after t_k bits. If the predator exhausts the patch before leaving it then $t_k = n$, the number of bits in the patch.

Let $q(k, t_k)$ be the probability of achieving outcome (k, t_k) in one particular way (one particular order of finding prey or not in each bit) and let $w(k, t_k)$ be the number of distinct ways of achieving outcome

(k, t_k) while obeying the rule given by (6). This is the number of ways of arranging k prey in t_k bits minus the number of such ways for which the predator would have stopped before reaching point (k, t_k) . Since k determines t_k uniquely we can use the simpler notation q_k and w_k .

Then we have

$$q_k = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \int_0^1 x^{\alpha+k-1} (1-x)^{\beta+t_k-k-1} dx \quad (7)$$

$$= \frac{(\alpha+k-1)^{(k)} (\beta+t_k-k-1)^{(t_k-k)}}{(\alpha+\beta+t_k-1)^{(t_k)}}$$

where $n^{(r)} = n(n-1)(n-2)\dots(n-r+1)$ is the r^{th} factorial power of n .

Again, we have

$$w_k = \binom{t_k}{k} - \sum_{i=0}^{k-1} w_i \binom{t_k-t_i}{k-i}. \quad (8)$$

The probability of achieving outcome (k, t_k) will be $w_k q_k$. Finally, combining (7) and (8) we have

$$E(G) = \sum_{k=1}^n k w_k q_k, \quad \text{and} \quad (9)$$

$$E(S) = \sum_{k=0}^n t_k w_k q_k. \quad (10)$$

The values found in (9) and (10) may be substituted into (1) to find the actual rate R when the strategy determined by the chosen constant c is used. When $R = c$ the best strategy has been found and c is the rate of finding prey that it achieves.

Notice that this solution will be for particular values of the parameters α , β , n and τ . These mathematical parameters correspond roughly to ecological parameters. The value n represents the size of each patch, τ the time between patches (based on unit time between bits within a patch) and α and β represent the patch quality and variability.

OTHER STOPPING RULES

a) The naive strategy.

The naive predator does not assess patch quality but relies on his knowledge of the average probability of finding a prey in each bit. Since this value is the same for all bits the naive predator remains in each patch until he has searched each bit.

Since the purpose of this paper is to investigate the advantage of assessing patch quality the overall rate of finding prey using the "best" strategy is compared with that of the naive predator.

b) The omniscient strategy.

If the predator can tell the quality of a patch without searching it then the best strategy will be to search the best patches completely and avoid the poorer patches. The rule the omniscient predator follows is to search only those patches whose expected quality is greater than or equal to the highest long term rate of finding prey possible for the predator (Charnov, pers. comm.). The omniscient predator will use Rule 2: search a patch completely if the probability of finding a prey in each bit is greater than c and avoid all poorer patches. Here c is the solution of the integral equation

$$c = \int_c^1 x^\alpha (1-x)^{\beta-1} dx / (\tau/n + \int_c^1 x^{\alpha-1} (1+t)^{\beta-1} dx) .$$

An omniscient predator can do better than a predator that must search for prey in order to assess patch quality. I assume, however, that searching is necessary to assess a patch.

c) The instantaneous rate strategy.

This strategist assesses the quality of patches by searching but, unlike the best strategist, he does not take account of future possibilities. Using a rule reminiscent of Charnov's marginal value theorem, the instantaneous rate strategist leaves a patch when the probability of finding a prey in the next bit falls below some critical value. Such a predator will follow Rule 3: continue searching a patch after having x prey in t bits if and only if

$$(\alpha+x)/(\alpha+\beta+t) > c .$$

Unlike the case in Charnov's deterministic model, however, the critical value, c , that maximizes the long term rate of finding prey is not equal to that rate itself. (In the cases I have considered the critical value is lower than the long term rate of finding prey.)

RESULTS

The several strategies may be compared by choosing particular parameters α , β , τ and n for the model and finding the best long term rate of finding prey for each strategy. There are a number of ways to vary the parameters. The table below gives the best rate achievable for several cases all having the same patch size, $n = 20$, the same time between patches, $\tau = 1$, and the same average quality of patches,

$\alpha/(\alpha+\beta) = .25$, but different amounts of variability in the quality of patches. Lower values of α and β mean more variable patches.

Table 1: Long term rate of finding prey achievable using different strategies, for various amounts of variability in patch quality. In each case $n = 20$, $\tau = 1$, $\alpha/(\alpha+\beta) = .25$.

<u>Patch quality</u>		<u>Rate of finding prey for various strategies</u>			
α	β	Best Assessment Strategy	Instantaneous Rate Strategy	Naive Strategy	Omniscient Strategy
.25	.75	.51390	.51345	.238	.636
.50	1.50	.41968	.41814	.238	.542
1.00	3.00	.33887	.33778	.238	.452
2.00	6.00	.29053	.28862	.238	.379

As the values of α and β increase in these cases, the distribution of the number of prey approaches a binomial with $n = 20$ and $p = .25$. For this limiting case the number of prey per patch is still variable but assessment of patch quality by the predator does no good and the best strategy is the same as the naive strategy. This limiting case bears out Charnov's assertion that there are stochastic models which give results similar to his deterministic model. This is only a special case, however.

As the values of α and β approach zero, the patches tend to have either no prey at all in a patch or prey in every bit in a patch. In this limiting case it does a great deal of good to assess the quality of patches. This case is essentially the example given by Oaten to show that the predator that assesses the quality of patches as it searches them can do arbitrarily much better than the naive predator.

The omniscient strategy clearly is best if it can be used. I assume that this strategy is impossible because patches are superficially similar

and must be searched to be assessed. However, the advantage of the omniscient strategy shows that it would be good for a predator if he could evaluate patch quality without actually searching.

For predators that must assess patch quality by searching, the instantaneous rate strategy is almost as good as the best strategy. They are so similar that I have calculated the long term rates of finding prey for these two strategies to five significant figures to show the difference.

While there is very little difference between these two assessment strategies they are both clearly better than the naive strategy that does not use assessment. This advantage of assessment is greatest when patch quality is most variable (α and β are small). For less variable patches (α and β large) the advantage of assessment is less.

If the size and distribution of quality of patches are kept the same but the time between patches is allowed to vary then both the best rate and the naive rate will vary. This is shown in the following table.

Table 2: Long term rate of finding prey for best assessment strategy and naive strategy, for various travel times between patches, τ . Here $n = 10$, $\alpha = 1$, $\beta = 1$.

τ	Best rate	Naive rate
2.0	.476	.417
1.0	.557	.455
.5	.614	.476
.1	.678	.495

As the time between patches decreases, the rate of finding prey increases for the best strategy and the naive strategy as well. The advantage of the best strategy over the naive strategy increases as the time between patches decreases.

The advantage of using the best strategy (the "assessment" strategy) will be greatest when patches vary most in quality and the time between patches is shortest.

The Stopping Rule

The stopping rule itself will be quite simple. The rule may be represented by the sequence $\{t_k\}$, where t_k indicates the time (number of bits searched) at which the predator would leave the patch if it had observed only k prey at that time. The following table gives the stopping rules for the two extreme cases considered in Table 1.

Table 3: The stopping rule for the best assessment strategy for $n = 20$, $\tau = 1$, and two different pairs of values for α and β . Here t_k is the number of bits such that the optimal forager will leave the patch if it has found only k prey after searching t_k bits.

k	0	1	2	3	4	5	6	7	8	9	10	≥ 11
$\alpha=.25, t_k$ $\beta=.75$	1	3	5	7	9	11	13	15	16	18	19	20
$\alpha=2, t_k$ $\beta=6$	2	5	9	12	15	17	20	20	20	20	20	20

For a given average prey density (for given $\alpha/(\alpha+\beta)$), the best strategy will require the predator to leave a patch more readily when the environment is more variable (α and β are smaller).

DISCUSSION

The Model

The model considered in this paper is a simple, fairly tractable version of Oaten's stochastic model of optimal foraging. Using my model it is possible to find the best strategy for a predator and to find the

long-term rate of finding prey for a predator that uses this strategy. This strategy involves assessing the quality of patches as they are searched.

There are two advantages in using a model simple enough for calculations to be made. First, it is possible to see how much better the assessment strategy is than the naive strategy in particular cases. Second, it is possible to see how the rate of finding prey and the advantage of using the assessment strategy depend on the parameters of the model.

While the advantage of the assessment strategy is not as striking for reasonable cases as it was in one example used by Oaten, it seems to be large enough to be of practical importance. Even a 20% increase in feeding rate should be important to a predator if it is capable of using the strategy that would make that increase possible.

One thing that is seen by varying the parameters of the model is that predators using the assessment strategy can do better in a highly variable environment than they can in a more uniform environment that has the same overall prey density.

The model presented in this paper is reasonable in several ways. Some predators do feed in discrete patches which may vary in quality although they are superficially similar. Within a patch, prey may only be present in certain places which the predator searches one at a time. It may take time and energy to go from patch to patch. The family of beta mixtures of binomial distributions is quite a rich family and could well represent the actual distribution of the number of prey per patch.

Other assumptions are almost certainly not correct but some of them may not matter very much. For example, patches are probably not all the same size and they are probably not superficially indistinguishable. If the predator can tell how good a patch is without actually visiting it then there is no need to assess the quality of the patch by searching it.

My model assumes that the rate of finding prey during a visit does not change as prey are found. This is a reasonable assumption for predators that search a patch systematically (unless they start in the best part and search parts with progressively worse quality).

The idea that the rate of finding prey decreases as a patch is searched is critical to Charnov's model of optimal foraging. The purpose of my paper is not to dispute the importance of this idea but rather to show with some calculations that, as Oaten has said, stochasticity is important too. I suspect that some of the behavioral aspects of prey depression discussed by Charnov, Orians and Hyatt (1976), such as predators frightening away prey, are more important than depletion due to the predator having eaten some of the prey in a patch and therefore not having them to eat again. I have ignored depletion to simplify the model and to isolate the effect of stochasticity. In fact, I believe that both depletion and stochasticity matter.

The main problem with the model discussed here is the question of how the predator knows what to do. I have followed Oaten in describing the predator as well-informed and capable of a difficult calculation (difficult for a bird without a computer, at least). In fact, it is I (and a computer) who have done the calculation. What the calculation

shows is a very simple rule which it is reasonable to believe an animal could follow. (A rule would be roughly of the form: if after the first few bits searched you have found a prey in at least a third of the bits searched continue, otherwise stop.)

Even if the strategy to be used by the Bayesian bird is simple enough for a real bird to follow there is the question of how the strategy is acquired. Are birds innate Bayesians or do they learn to be Bayesians? Since the best strategy depends on environmental quality, which must change over time, it would be better if foraging strategy was learned and could be adjusted to changing conditions. One important result that can be seen in this model is that suboptimal assessment strategies are often better than the naive strategy.

The model shows that if patches vary in quality there are simple strategies of searching patches based on assessment of patch quality that enable a predator to do better than a predator that does not assess patch quality. The question remains of whether predators actually do use strategies that take advantage of variability in patch quality.

Recently Breck (1978) has considered several "suboptimal foraging strategies." He considers both the long term rate of finding prey and the chance of not finding any prey during some time interval. Among his conclusions is one that it becomes more important to use the information in the environment as the variability of the environment increases.

ACKNOWLEDGMENTS

Most of the work on this paper was done while I was on leave from the University of California, Riverside, at the Department of Zoology, Oxford. I would like to thank Richard Cowie for bringing this problem to my attention. I would also like to thank Eric Charnov, Adeline Nunez and Allan Oaten for their comments on this paper.

REFERENCES

- Breck, J. E. 1978. Suboptimal foraging strategies for a patchy environment.
Ph.D. Thesis, Michigan State University.
- Charnov, E. L. 1973. Optimal foraging: some theoretical explorations.
Ph.D. Thesis, University of Washington.
- Charnov, E. L. 1976. Optimal foraging: the marginal value theorem.
Theor. Pop. Biol., 9:129-136.
- Charnov, E. L., Orians, G. H., and Hyatt, K. 1976. Ecological implication
of resource depression. Amer. Natur., 110:247-259.
- DeGroot, M. H. 1970. "Optimal Statistical Decisions," McGraw-Hill,
New York.
- Krebs, J. R. 1973. Behavioral aspects of predation, in "Perspectives
in Ethology" (P. P. Bateson and P. H. Klopfer, Eds.), pp. 73-111,
Plenum Press, New York.
- Krebs, J. R., Kacelnik, A., and Taylor, P. 1978. Test of optimal sampling
by foraging great tits. Nature, 275:27-31.
- Nemhauser, G. L. 1966. "Introduction to Dynamic Programming," John Wiley
and Sons, New York.
- Oaten, A. 1977. Optimal foraging in patches: a case for stochasticity.
Theor. Pop. Biol., 12:263-285.
- Pyke, G. H., Pulliam, H. R., and Charnov, E. L. 1977. Optimal foraging:
a selective review of theory and tests. Quart. Rev. Biol., 52:137-154.
- Schoener, T. W. 1971. Theory of feeding strategies. Ann. Rev. Ecol. Syst.,
11:369-404.