Published in final edited form as:

J Exp Psychol Anim Behav Process. 1996 October ; 22(4): 480–496.

Bayesian Analysis of Foraging by Pigeons (*Columba livia*)

Peter R. Killeen, Gina-Marie Palombo, Lawrence R. Gottlob, and Jon Beam

Department of Psychology, Arizona State University.

Abstract

In this article, the authors combine models of timing and Bayesian revision of information concerning patch quality to predict foraging behavior. Pigeons earned food by pecking on 2 keys (patches) in an experimental chamber. Food was primed for only 1 of the patches on each trial. There was a constant probability of finding food in a primed patch, but it accumulated only while the animals searched there. The optimal strategy was to choose the better patch first and remain for a fixed duration, thereafter alternating evenly between the patches. Pigeons were nonoptimal in 3 ways: (a) they departed too early, (b) their departure times were variable, and (c) they were biased in their choices after initial departure. The authors review various explanations of these data.

In this article, we analyze foraging strategies in a simple experimental paradigm in terms of optimal tactics and constraints on their employment. Evolutionary processes drive organisms and their parts toward optimality by selecting individuals that are better able to exploit their environment to the benefit of their progeny. Whereas the ultimate criterion for selective advantage is measured by the number of viable offspring in the next generation, it is the proximate characteristics such as sensory acuity, plumage, and foraging strategies that are selected in the current generation. Individuals who survive are better in some of these key respects than those who do not, recognizing inevitable trade-offs among the aspects selected; ornate plumage may interfere with foraging, foraging with nest tending, and so on. When we observe a species-specific behavior, it is natural to presume that it is adaptive and to seek to understand the environmental pressures that make it so.

How, though, do we justify the jump from *adapted* to *optimal*? These observations set the stage. First, *better* and *best* must always be defined in terms of the alternate strategies that an organism might “choose” or that its competitors have chosen. As long as a structure or function is better than that of its competitors, the nature of the best (i.e., optimal) is irrelevant to any organisms other than ecologists; in the exponential mathematics of generations, better is all that matters. Second, these strategies are subject to side effects and structural-epigenetic constraints (e.g., bright plumage attracts predators as well as mates, the memorial requirements for optimal foraging compete with those for song, and so on). It is the system as a whole that must compete successfully; some behaviors may be inferior to those they replace but survive because they are part of a package, that is, on the whole, superior. Is there any sense then in speaking of optimal strategies when the constraints are on systems, not subsystems such as foraging, and when the ultimate criterion of relative genetic success is so intractable to experimental manipulation? The arguments on this point continue to compete and evolve: For reviews, see Krebs and Davies (1978, 1991), Lea (1981), and Shettleworth (1989). Stephens and Krebs's (1986) last chapters provide a thoughtful consideration of just what foraging models can do and some of the achievements and pitfalls of optimal foraging arguments.

What is good about optimal foraging theories is that they guide our understanding of the constraints under which an organism labors and thus the driving forces in its niche. They provide the antithesis of the null hypothesis, telling us not the lower bound (no advantage) but the upper bound (the best that could be expected). If we find an organism using a strategy that is obviously inferior to the best one that we can imagine, we are either imagining environmental pressures different from those under which the behavior evolved or not taking into account the epigenetic constraints that bridle the organism. The deviation between the ideal and real instructs us in these constraints and pressures.

Many of the experimental paradigms in which optimality analyses are invoked were designed for purposes other than to test models of optimal foraging. Consider, for instance, the traditional experimental paradigm in which reinforcement is delivered with a constant probability over time for one of two concurrently available responses. In such situations the proportion of time that animals spend responding to one of the two schedules approximately equals—or matches—the relative rate of reinforcement delivered by that schedule. Is this behavior optimal? It is almost (Staddon, Hinson, & Kram, 1981; Williams, 1988); although it is not a bad strategy, many other strategies would do about as well. Such schedules have a “flat optimum.”

Many experimental schedules pit long-term optima (e.g., maximizing overall rate of reinforcement) against short-term optima (e.g., a stimulus correlated with a higher local probability of reinforcement) and find that the immediate contingencies overpower the long-term ones (viz., the infamous weakness of deferred rewards in self-control) or act conjointly with them (Williams, 1991). However, these results do not provide arguments against optimality, so much as a clarification of the time scales over which it may be prudent for an organism to optimize. The future should be discounted because it is uncertain, but the calculation of just how many “birds in the bush” are worth one in the hand is, in general, a near-intractable problem of dynamic programming (e.g., Stephens & Krebs, 1986). Recourse to situations in which many of the variables are under strong experimental control (e.g., Commons, Kacelnik, & Shettleworth, 1987) weakens the subject's control and minimizes the variability characteristic of the field. This ancient abstract-tractable versus real-complex dilemma is resolvable only by cycling through Peirce's abductive-retroductive program: hypothesis generation in the field, model construction and testing in the laboratory, redeployment in the field (Cialdini, 1980, 1995; Killeen, 1995; Rescher, 1978). The laboratory phase of this cycle is engaged here: formalization and experimental testing of a quantitative optimal foraging model.

Optimal Search

Rather than apply optimality arguments to traditional scheduling arrangements, it is possible to design a scheduling arrangement in which the strategy that maximizes long-term rate of reinforcement also maximizes short-term probability of reinforcement and in which the optimal search behavior is well defined.

How long should a forager persevere in a patch? Intuition suggests it should stay as long as the probability of the next observation being successful is greater than the probability of the first observation in the alternate patch being successful, taking into account the time it takes to make those observations. In our experimental design, this is the foundation of the ideal strategy because it optimizes both immediate and long-term returns. It is an instance of Charnov's (1976) marginal value theorem, “probably the most thoroughly analysed model in behavioral ecology, both theoretically and empirically” (Stephens & Dunbar, 1993, p. 174). However, as Stephens and Dunbar continue, “although it is considered the basic model of patch-use in behavioral ecology, the marginal-value theorem does not provide a solution of the optimal (rate-maximizing) patch residence time; instead, it provides a condition that optimal

patch residence times must satisfy” (p. 174). Further specification of the foraging conditions is necessary, and those are provided here in the context of optimal search theory.

The theory of optimal search (Koopman, 1957) was developed for situations in which it is possible to (a) specify a priori the probability that a target would be found in one of several patches (the priors) and (b) specify the probability of discovering the target within a patch as a function of search time or effort (in foraging theory this is the gain function; in search theory it is the detection function; see, e.g., Koopman, 1980; Stone, 1975). This theory was designed for naturalistic situations in which pilots are searching for survivors at sea, for enemy submarines, and so on. It is applicable not only to those “foraging” situations but also to those in which the depletion and repletion of patches are at a steady state, to situations in which prey occurs and moves on at constant rate that is only minimally perturbed by a prior capture, and to the initial selection of patches after an absence (Bell, 1991).

The most common detection function assumes a constant probability of detecting the prey over time, which implies an exponential distribution of times to detection (Figure 1). How should an organism distribute its time in such patches to maximize the probability that the next observation will uncover the reward? Consider a situation in which on each trial the target, a prey item, is in one or the other of two patches, with the prior probability of it being in Patch i being $p(P_i)$ and where $p(P_1) + p(P_2) = 1.0$. It is obvious that the searcher should start by exploring the more probable patch first: Patch 1 if $p(P_1) > p(P_2)$, Patch 2 if $p(P_1) < p(P_2)$, and either if $p(P_1) = p(P_2)$.

There are two ways to derive the optimal giving-up time corresponding to the point of equality. The more general is the Bayesian analysis given in the Appendix. It yields the same prediction (Equation 2) as the following, intuitively simpler analysis. We assume that there is a constant probability of finding the prey item during each second of search: In the patch that contains the target on that trial, the probability of finding it in any epoch is λ , and in the other patch it is 0. Given the constant probability λ of finding the prey, the continuous curves in Figure 1 show that the probability that an organism will have found it in Patch i by time t is as follows:

$$p(F_{i,t}) = p(P_i)(1 - e^{-\lambda t}). \quad (1)$$

The slope of this exponential detection function is the marginal rate of return from the patch and is given by the time derivative of Equation 1:

$$\dot{p}(F_{i,t}) = p(P_i)\lambda e^{-\lambda t}. \quad (2)$$

Notice that as time in a patch increases, the marginal rate of return decreases exponentially. (This is called “patch depression,” but in the present model it results not from a depletion of the patch but rather from the logic of a constant-probability sampling process: The probability of long runs before a success decreases with the length of the runs.) The first time at which the marginal values for the two patches are equal is when the slope on the more probable side, $\dot{p}(F_{1,t})$, has fallen to the value of the slope on the inferior side when that is first sampled (i.e., at $t = 0$ for Patch 2), which, from Equation 2, is $p(P_2)\lambda$. This happens when $\dot{p}(F_{1,t}) = \dot{p}(F_{2,0})$, that is, when $p(P_1)\lambda e^{-\lambda t} = p(P_2)\lambda$, at which time the marginal return from the better patch equals the initial marginal return from the poorer patch. Solving for t yields the predicted point of indifference:

$$t^* = \ln[p(P_1)/p(P_2)]/\lambda, \quad \lambda > 0. \quad (3)$$

As soon as $t > t^*$ the animal should switch; this is the optimal giving-up time. If, for instance, the priors are $p(P_1) = 3/4$, $p(P_2) = 1/4$, and $\lambda = 0.10/\text{s}$, then the searcher should shift to Patch 2 when $t > 10.99$ s. This analysis omits travel time. In the experimental paradigm to be analyzed, travel time is brief and, as we shall see, its omission causes no problems.

Note that the proposed experimental arrangement is different than the traditional “concurrent Schedule of reinforcement” because, unlike traditional concurrents, the probability of reinforcement in a patch does not increase while the animal is in the other patch; that is, the “clock stops running” when the animal is elsewhere. The paradigm provides a model of foraging between two patches at steady states of repletion rather than between patches that are replenished while the searcher is absent. Traditional concurrent schedules are like trap lines; once the prey falls in, it remains until collected. The present “clocked” concurrents are more like a hawk on the wing; by searching the north slope the hawk misses the darting ground squirrel on the south slope, who will not wait for him. Like the hawk, animals in this experiment will not switch patches because things are continuously getting better elsewhere but rather because of the increasing certainty that things are not as good in the current patch as they are likely to be in the other patch when first chosen. Each response yields information, a posteriori, about whether the chosen patch will be fruitful on the current trial. Can animals use such information? If they can, it will lead them to switch at $t = t^*$. Experimental designs similar to this one have been executed by Mazur (1981) and Zeiler (1987); Houston and McNamara (1981) have derived models for other concurrent paradigms. However, only in the present case is the optimal short-term strategy also the optimal long-term strategy, and trade-offs between delay and probability of reinforcement are eliminated. The present design offers a “pure” case in which to test for optimality.

This model provides a second test of the optimality of the subjects' search behavior. From the point at which the slopes of two exponential functions such as Equation 1 are equal, all subsequent parts of the curves are identical. (To see this, cut out the middle curve in Figure 1 after t^* and position it over the first part of the bottom curve. This identity is a unique property of exponential distributions). Spending $t = t^*$ seconds in the better patch brings the posterior probability that that side actually contains the target down to $1/2$. At $t = t^*$ the subjects should become indifferent, and, because the detection functions are thereafter identical—the probabilities of payoff on the two sides are equal—they should thereafter remain generally indifferent. However, it continues to be the case that the longer they spend on one side, the greater the a posteriori probability that food is to be found in the other side. Therefore, they should alternate quickly and evenly between patches. The dwell time in a patch after $t = t^*$ should depend only on the travel time, which in the present case is symmetrical. As travel time increases, dwell time should increase but should remain equal on each side.

It was our strategy, then, to design an experimental paradigm that was isomorphic with this idealized search model, a model whose theoretical import has been thoroughly analyzed (Koopman, 1980; Stone, 1975), one for which there are explicit measurable optimal strategies and one that neither plays off short-term benefits against long-term ones nor introduces stimulus changes such as conditioned reinforcers with their own undetermined reinforcing strength. Optimal search is well defined in this experimental paradigm: Choose the better patch exclusively for t^* seconds and be indifferent thereafter. If pigeons search optimally, then they must behave this way. If they do not behave this way, then they are not searching optimally. If they are not searching optimally, we can ask further questions concerning constraints on learning, memory, and performance that might be responsible for the observed deviations from optimality or questions concerning our assumptions of what is or should be getting optimized (Kamil, Krebs, & Pulliam, 1987; Templeton & Lawlor, 1981). It is not a model of optimality that is being tested here; that is canonical. It is pigeons that are being tested here in their ability to approximate that ideal.

Experiment 1

Method

Subjects—Seven homing pigeons (*Columba livia*), all with previous experimental histories, were maintained at 80% to 85% of their free-feeding weights in a 12-hr photoperiod.

Apparatus—Experiments were conducted in a standard BRS/LVE (Laurel, MD) experimental chamber 33 cm high $\times$ 36 cm wide $\times$ 31 cm deep, beginning approximately 3 hr into the day cycle. Three response keys were centered on the front wall 7.5 cm apart and 20 cm above the floor. A 6 cm wide $\times$ 4 cm high aperture through which reinforcers (2.5-s access to mixed grain) could be delivered was centered on the front wall with the bottom of the aperture 8 cm above the floor. A houselight was centered at the top of the front panel. White masking noise at a level of approximately 75 dB was continuously present.

Procedure—Sessions consisted of 60 trials, on any one of which the reinforcer was available (primed) for responses to only one of the keys. The probability that it could be obtained by responding on the left key was $p(P_1)$ and on the right key, $p(P_2) = 1 - p(P_1)$. These probabilities were arranged by randomly sampling without replacement from a table so that in each session the subjects' relative rate of payoff on the left key was exactly $p(P_1)$.

Each trial started with only the central key lit green. A single response to this key extinguished it and lit the white side keys, initiating the search phase. Reinforcement was scheduled for responses according to Equation 1, with t advancing only while the animal was responding on the primed side. This is a “clocked” version of a “constant-probability variable interval (VI) Schedule.” It guarantees that the rate of reinforcement while responding on a key is λs^{-1} . It models foraging situations in which the targets appear with constant-probability λ every second but will leave or be appropriated by another forager if they appear when the subject is not looking, as often occurs in the search for mates or prey. The particulars of this task satisfy the assumptions of Koopman's (1980) basic search model.

Trials lasted until the reinforcer was obtained with the next trial commencing after a 3-s intertrial interval. A minimum of 30 sessions were devoted to each of the conditions, which are identified in Table 1 by the values of $p(P_i)$ and λ that characterized them. The obtained values of $p(P_i)$ exactly equaled the programmed values. In these experiments, λ , the probability of reinforcement being set up during each second on the primed side, was the same for each of the keys. Data are times spent responding on each key (when that key was chosen first), measured from the first response on that key until the first response on the other side key, collected over the last 15 sessions of each condition, and the number of responses on each key in 1-s bins. All subjects experienced Conditions 1 and 2 and thereafter were assigned to other conditions. The better patch was on the right key under Condition 5 and on the left key under all other conditions.

Results

In Condition 1 the average rate of availability of the prey on the primed side was $\lambda = 1/10$ (i.e., a VI 10-s schedule), and the prior probability of either side being primed was 0.5. The pigeons' initial giving-up time from their preferred side was 3 s, and thereafter they showed little bias, spending approximately 2.6 s on visits to the left key and 2.4 s on visits to the right key.

In Condition 2, $p(P_1) = 0.75$, $\lambda = 1/10$. Figure 2 shows the relative frequency of responses on the better key, averaged over all 7 subjects, as a function of the time into the trial. The optimal behavior, indicated by the step function, requires complete dedication to the better side until 11 s have elapsed and thereafter strict alternation between the sides. None of the individual

subject's average residence profiles resembled a step function (cf. Figure 9), although on individual trials they did. This is because there was variability in the location of the riser from one trial to the next, and that was the major factor in determining the shape of the ogives. During the first 3 s 96% of the responses were to the better side, but thereafter no animal approximated the optimal performance. On the average the animals spent 6.7 s on the better side before giving up; with a standard error of 0.9 s, this is significantly below the optimal duration of 11 s. Not only was there a smooth and premature decrease in the proportion of responses on the better side, but the proportion remained biased toward the better side. Another perspective on this performance is provided by Figure 3, which shows the amount of time spent on each side before a changeover to the other side as a function of the ordinal number of the changeover. After the initial visit to the better patch, the pigeons alternated between the two, spending a relatively constant amount of time in each patch over the next dozen switches. Table 1 shows that the dwell time in the better patch on the second visit was longer than that on the first visit to the nonpreferred patch under all other experimental conditions, indicating a similar residual bias.

In Condition 3, the prior $p(P_1) = 0.75$, and $\lambda = 1/20$, corresponding to a VI 20-s schedule on the side that was primed. The initial giving-up time doubled to just over 13 s but still fell short of the optimal, now 22 s. A residual bias for the better patch was maintained for 15 subsequent alternations between the keys.

In Condition 4, the prior $p(P_1) = 0.75$, and $\lambda = 1/40$, corresponding to a VI 40-s schedule on the side that was primed. Again, there was an increase in the initial visit to the preferred patch, but it too fell short of the optimal, now 44 s. There was a maintained residual bias for the better patch.

Throughout these conditions, the better patch was always assigned to the left key to minimize the hysteresis that occurs when experimental conditions are reversed. Our intention was to place all biases that may have accrued in moving from one experimental condition to another in the service of optimization, and yet the animals fell short. In Condition 5, the prior for the better patch was reduced to $\frac{2}{3}$, and the better patch was programmed for the right key. The rate parameter $\lambda = 1/10$, corresponding to a VI 10-s schedule on the side that was primed. Table 1 shows that the initial giving-up time fell to 4 s, again too brief to satisfy the optimal dwell time of $10 \ln(2/1) = 6.9$ s.

To assess the amount of hysteresis in this performance, in the final condition (Condition 6) the locations of the two patches were again reversed, with the priors and rate constants kept the same as in Condition 5. Table 1 shows that initial dwell time was longer under this arrangement, although still significantly below the optimal 6.9 s.

Discussion

The pigeons did not do badly, achieving some qualitative conformity with the expectancies of optimal search theory and maintaining a good rate of reinforcement in the context. There are three details in which data did depart from optimality: (a) The pigeons leave the better patch too soon (see Figure 2); (b) they maintain a residual bias for the better patch through subsequent alternations between them (see Figures 2 and 3); (c) their relative probability of staying in the better patch is not a step function of time. These aspects are treated in order by examining alternative hypotheses concerning causal mechanisms.

Premature Giving Up

Travel time: The premature departure is clearly nonoptimal under the canonical model of optimal search. It could not be due to the added cost of travel time between the keys because

that should have prolonged the stays on either side rather than abbreviating them. Traditional programming techniques use a delay in reinforcement after the animal changes over to a concurrently available schedule, called a *changeover delay*, to minimize rapid alternation between the schedules. This is necessary because in those concurrent schedules the probability of reinforcement continues to accrue in one schedule while the animal is engaged in the other, thus often reinforcing the first changeover response, unless such a changeover delay is used (see, e.g., Dreyfus, DePorto-Callan, & Pseillo, 1993). Unlike such traditional schedules, however, the contingencies in the present experiment do not simultaneously encourage and discourage animals from switching. The base-rate probability of reinforcement in the first second after a switch to the other key is independent of how long the animals have been away from it. The addition of a changeover delay would have prolonged visits to the patches, but the appropriately revised model would then predict even larger values of t^* . Finite travel times cannot explain the failure to optimize, and procedural modifications to force longer stays would force even larger values for t^* . Success at eventually getting giving-up times to equal redefined values of optimality would speak more to the experimenter's need to optimize than to that of the subjects.

Matching: Perhaps some mechanism led the animals to match their distribution of responses to the rates of reinforcement (Baum, 1981; Davison & McCarthy, 1988). Indeed, the overall proportion of responses to the better key did approximately equal the probability of reinforcement on it. However, that hypothesis explains none of the features of Figures 2 and 3. To see this, we plot the posterior probabilities of reinforcement as a function of time on a patch in Figure 4. The time courses of the ogives are vaguely similar to the data observed, but (a) they start not near 1.0, like the data, but rather at the value of the prior probabilities, (b) they are flatter than the observed data, and (c) the mean of the ogives occurs later in the trial than the observed probabilities. Perhaps a more complicated model that had matching at its core could account for these data, and if history is a guide one will be forthcoming, but there are other problems confronting such matching theorists.

Relative probabilities are not the same as relative rates of reinforcement the way those are measured in the matching literature: There the time base for rates includes the time the animal might have been responding but was occupied on the other alternative. In these experiments the relative probabilities of reinforcement are given by the priors, and the rates of reinforcement while responding are equal to λ/s for each of the alternatives. However, because the animals spend proportionately more time responding on the better alternative, the relative rate of reinforcement for it in real time (not in time spent responding) is greater than given by the priors. In these experiments it equaled the relative value of the priors squared. If the prior for an alternative is 0.75, its relative rate of reinforcement (in real time) was 0.90. This construal of the independent variable would only make things worse for the matching hypothesis. Matching may result from the animal's adaptive response to local probabilities (Davison & Kerr, 1989; Hinson & Staddon, 1983), but it does not follow that matching causes those locally adaptive patterns.

Flat optima: Just how much worse off does the premature departure leave the birds? It depends on what the animals do thereafter. If they immediately go back to the preferred key and stay there until t^* , they lose very little. If they stay on the other side for a lengthy period, they lose quite a bit. Figure 5 shows the rates of reinforcement obtained for various dwell times, assuming the animals switch back and forth evenly thereafter, derived from simulations of the animals' behavior under Condition 2. We see that rate of reinforcement is in fact highest where we expect it to be, for dwells of just over 11 s. The sacrifice of reinforcement incurred by switching at 6 s is not great. However, if nonoptimality is simply a failure to discriminate the peak of this function, why should the pigeons have not been as likely to overstay the optimal on the better key than to quit early? They do even better by staying for 16 s than by staying for only 6 s.

This relatively flat optima should leave us unsurprised that giving-up times were variable but does not prepare us for the animals' uniformly early departures.

Alternative birds-eye views: Perhaps the birds were operating under another model of the environment (Houston, 1987). Perhaps, for instance, they assumed that the prior probabilities of reward being available in either patch, $p(P_i)$, equaled 1.0 but that the detection functions had different rate constants equal to $\lambda p(P_i)$: $\lambda_1 = 0.075$, $\lambda_2 = 0.025$. This “hypothesis” preserves the overall rates of reinforcement on the two keys at the same value. However, under this hypothesis the value of t^* for Condition 2 is 14.4 s, an even longer initial stay on the preferred side. It cannot, therefore, explain the early departures.

Alternatively, even though the detection function was engineered to have a constant probability of payoff, the animals might be predisposed to treat foraging decisions routinely under the assumption of a decreasing probability. This would make sense if animals always depleted the available resources in a patch as they foraged. This is often the case in nature but not in this experiment, in which they received only one feeding and were thereafter required to make a fresh selection of patches. Of course, such a hypothesis (of decreasing returns) might be instinctive and not susceptible to adjustment by the environmental contingencies. If this is the case, it is an example of a global (“ultimate”) maximization that enforces a local (“proximate”) minimum: The window for optimization becomes not the individual's particular foraging history but the species' evolutionary foraging context. Such instinctive hypotheses would be represented by different detection functions (e.g., “pure death functions”) than those imposed by the experimenter, ones recalcitrant to modification. This could be tested by systematically varying the experimental contingencies and searching for the hypothetical detection function that predicted the results without the introduction of a bias parameter or by systematically comparing species from different ecological niches. Simpler tests of the origin of the bias are presented later.

Experience: Perhaps the animals just did not have enough experience to achieve secure estimates of the priors. However, these experiments comprised more than 1,500 trials of homogeneous, consistent alternatives, more than found in many natural scenarios. Sutherland and Gass (1995) showed that hummingbirds could recover from a switch in which of several feeders were baited within 30 trials.

Could it be sampling error that causes the problem? Random variables can wander far from their means in a small sample size. Had the patch to be reinforced been primed by flipping a coin (i.e., by a strictly random “Bernoulli process”), by the time the animals had experienced 1,000 trials the standard error of the proportion of reinforcers delivered in the better patch would be down to $[(0.75 \times 0.25)/1,000]^{0.5} = 0.014$; their experienced priors should have been within 1.4% of the programmed priors. In these experiments, however, the primed patch was determined in such a way that by the end of each session the relative payoff on the better side was exactly $p(P_1)$, with standard error of 0 from one session to the next, further weakening the argument from sampling variability. The pigeons' bias cannot be attributed to Bernoulli variability intrinsic to a sampling process.

Perhaps the problem arose from an extended experimental history with the better patch on the same side. No; if anything, that should have prolonged giving-up times, which fell short of optimal. The decision to avoid hysteresis effects that derive from frequent changes of the location of the best alternative may have resulted in dwell times that were longer than representative. It cannot explain times that were shorter than optimal. It is the latter issue we were testing, not point estimates of dwell times.

Time horizons: This model gives the probability that reinforcement is primed for a patch, given that it has not been found by time t . However, perhaps the decision variable for the animals is the relative probability of finding food for the next response or in the next few seconds or in the next minute. Would these different time horizons change their strategies? No. Because of the way in which the experiment was designed, as long as the time horizons are the same for each patch, the optimal behavior remains the same.

Of course, the time horizons might have been different for the two patches. That hypothesis is one of many ways to introduce bias in the model, to change it from a model of optimal search to a model of how pigeons search. Optimality accounts provide a clear statement of the ideal against which to test models of constraints that cause animals to fall short, and that is their whole justification.

A representativeness heuristic: Perhaps the subjects leave a patch when the probability of reinforcement falls below 50%, given that food is going to be available in that patch. That is, whereas they base their first choice of a patch on the prior (base-rate) probabilities, thereafter they assume the patch definitely contains a target, and they base their giving-up time on the conditional probability of reinforcement. Figure 1 shows that this value is the abscissa corresponding to an ordinate of $p = .5$ on the dashed curve, which equals 6.9 s for $\lambda = 1/10$. This is close to the obtained average giving-up time of 6.7 s. Although there is a kind of logic to this strategy, it is clearly nonoptimal because the subjects do not know that reinforcement is going to be available in that patch; furthermore, if they did know that, they should not leave at all! The value of 6.9 s is representative of the amount of time it takes to get reinforcement in Patch 1 *if* it will be available there; that is, this time is representative if the prior base rates are disregarded. A similar fallacy in human judgment has been called “the representativeness heuristic” and is revealed when people make judgments on the basis of conditional probabilities, completely disregarding the priors. This hypothesis might provide a good account of giving-up times when λ is varied, but because it rules out control of those times by the prior probabilities, $p(P_i)$, it cannot account for the observed changes in behavior when the priors are varied (see Table 1). However, there may be a seed of truth in this hypothesis: Perhaps the priors are discounted without being completely disregarded.

Washing out the priors: What if the animals lacked confidence in the priors despite the thousands of trials on which they were based? Perhaps they “washed out” those estimates through the course of a trial. If so, then at the start of a new trial after a payoff on the poorer patch, the animals should choose that patch again (win-stay). However, the first datum in Figure 2 shows that this did not happen: 97% of the time they started in Patch 1. If we parsed the trials into those after a reward on one patch versus those after a nonreward on that patch, it is likely that we would see some dependency (Killeen, 1970; Staddon & Horner, 1989). However, it is easy to calculate that the choice of the dispreferred alternative after a reward there could increase to no more than 12% to retain the 97% aggregate preference for the better alternative. This is not enough to account for the observed bias. It is possible, however, that it is an important part of the mechanism that causes the priors to be discounted on a continuing basis.

Discounting the priors: Missattribution: Likelihood ratios of 2:1 (rather than the scheduled 3:1) would closely predict the observed first giving-up times in Conditions 2 to 4. Why should the priors be discounted, if this is what is happening? In no case are the priors actually given to the subjects; they must be learned through experience in the task (Green, 1980; McNamara & Houston, 1985; Real, 1987). The observed discounting may occur as a constraint in the acquisition of knowledge about the priors, or it may occur in the service of optimizing other variables not included in the current framework. In the first instance, let us assume that the subjects occasionally missattribute the source of reinforcement received in one patch to the

other patch (Davison & Jones, 1995; Killeen & Smith, 1984; Nevin, 1981). Then the likelihood ratio will become less extreme, a kind of regression to the mean (see the Appendix for the explicit model and parameter estimation). If they missattribute the source of reinforcement 18% of the time, it leads to the giving-up times shown in the last column of Table 2.

Discounting the priors: Sampling: There may be other reasons for discounting the priors. If we simply weight the log-likelihood ratio of the priors less than appropriate (i.e., less than 1.0), we guarantee an increased probability of sampling unlikely alternatives. In particular, if we multiply the log-likelihood ratios by 0.6, the predicted giving-up times are within 0.2 s of those predicted by the missattribution model. Arguments have occasionally been made that such apparently irrational sampling may be rational in the long run (Zeiler, 1987, 1993). What is needed is to rationalize the “long run” in a conditional probability statement (i.e., to “conditionalize” on the long run); until that is done, it is the theorist's conception of rationality, not the subject's, that is uncertain. An example of such an analysis is provided by Krebs, Kacelnik, and Taylor (1978; also see Lima, 1984) for a situation in which patches provided multiple prey at constant probabilities, but the location of the patch with the higher probability varied from one trial to the next. In this case, sampling is obviously necessary at first because the priors are 0.5; once the posteriors for assigning the identity of the better patch reach a criterion (either through a success or after n unrequited responses), animals should choose the (a posteriori) better patch and stay there. Thus, the behavior predicted in this “two-armed bandit” scenario is a mirror image of the behavior predicted in the present experiment.

These alternate rationales for discounting the priors are amenable to experimental test. Both incur one additional parameter—missattribution rates or discount rates—whose values should be a function of experimental contingencies or the ecological niches of the subjects. In experiments not reported here, we attempted to test the missattribution hypothesis by enhancing the salience of the cues, but this did not improve performance. However, such tests are informative only when they achieve a positive result, because the obtained null results may speak more to the impotence of the manipulations than to that of the hypothesis.

Residual Bias—Real (1991) showed that bumblebees do not pool information about the quality of a patch from more than one or two visits to flowers in it (i.e., take into account the amount of time spent and number of successes to achieve an appropriately weighted average; also see McNamara & Houston, 1987). This may also be the case in the present study. Figure 3 suggests that the pigeons did not treat the better response key as the same patch when they revisited it but rather as a different patch. Three dwell times alone give an accurate account of the pigeons' foraging over the first dozen alternations: initial visits to the preferred side, all subsequent visits to the preferred side, and all visits to the nonpreferred side (see Figure 3). The return to the better patch may properly be viewed not as a continuation of a foraging bout but as exploration of a new patch whose statistics are not pooled by the animal with the information derived from the first search.

Such a partitioning of feeding bouts into three dwell times is less efficient than pooling the information from earlier visits; the animals' failure to pool, perhaps because of limits on memory, constrains the best performance that they can achieve. Had the initial giving-up time been optimal, they could have achieved globally optimal performance by calculating and remembering only two things: Search the better patch first for t^* seconds; thereafter, treat both patches as equivalent. Thus, optimal performance would have required them to remember only two things. Describing the machinery necessary for them to figure these two things out, however, is a matter for another article.

Because all the subjects switched too early, they could partially “correct” this deviation from optimality by staying longer on the better side on their next visit to it. An optimal correction

in Condition 2 would have required the pigeons to spend about 6 s in the better patch on their first return to it. However, the duration of the animals' visits to the preferred patch remained consistent at 3.4 s through the remainder of the trial. Given that residual and constant bias, the pigeons finally exhausted the remaining posterior advantage for the better side at about 22 s into the trial. There was scant evidence, even at that point, of their moving toward indifference (see Figure 2). However, most trials terminated before 22 s had elapsed; therefore, most of the conditioning the subjects received reinforced the residual bias toward the better patch. A test of the hypothesis that the subjects treat the better key as a different patch after the first switch and that the residual bias was caused by the failure to fully exploit the posteriors on the first visit is provided in the fourth condition of Experiment 2. However, adequate discussion of asymptotic bias is contingent on our having a model of fallible time perception, to which construction we now turn.

Ogival Residence Profiles—Optimal behavior in these experiments is a step function of residence time on the first visit to the preferred side, “the ‘all-or-none’ theme so common in optimal behaviour” (Lea, 1981, p. 361). However, because temporal discriminations are fallible, we do not expect to find a perfect step function; on some trials the pigeons will leave earlier or later than on others, and this is what makes the average probability of being in the better patch an ogival function of time.

There are many models of time perception, most involving pacemaker-counter components. Such systems accrue pulses from the pacemaker and change state when their number exceeds a criterion. Consistent with the central limit theorem, as the criterial number of counts increases, the distributions of these responses approach the normal. The variance of the distributions will increase with their means: either with the square of the means (e.g., Brunner, Kacelnik, & Gibbon, 1992; Gibbon, 1977; Gibbon & Church, 1981) or proportionally (e.g., Fetterman & Killeen, 1992; Killeen & Fetterman, 1988). In general, they will change as a quadratic function of time (Killeen, 1992), as outlined in the next section.

General Timing Model—Consider a system in which time is measured by counting the number of pulses from a pacemaker, and those pulses occur at random intervals (independent and identically distributed) averaging τ seconds. The variance in the time estimates that is due to the randomness of the pacemaker may be represented as a quadratic function of τ . The counting process may also be imprecise and thereby add variability to the process, which also may be represented as a quadratic function of the number of counts, n . How do these two sources of variance—a random sum of random variables—combine to affect the time estimates? Killeen and Weiss (1987) gave the variance of the estimates of time interval t for such a process, σ_t^2 , as

$$\sigma_t^2 = (at)^2 + bt + c^2. \quad (4)$$

The parameter a is the Weber fraction; it depends only on the counter variance and is the dominant source of error for long intervals, in which the coefficient of variation (the standard deviation divided by the mean) is simply a . The parameter b captures all of the pacemaker error, plus Bernoulli error in the counter; its role is greatest at shorter intervals. The period of the pacemaker, τ , is embedded in b . The parameter c measures the constant error caused by initiating and terminating the timing episode and other variability that is independent of t and n ; it is the dominant source of error for very short intervals.

Figure 6 shows the distribution of estimates of subjective time over real times of 5, 10, 20, 30, and 40 s. To draw this figure, the parameter a in Equation 4 was fixed at 0.25, and the other parameters set to 0. The optimal times for switching out of the better patch for λ of 1/10 and 1/20 are designated by the vertical lines. Notice that as the discriminial dispersions move to the

right, they leave a portion of their tail falling to the left of the optimal giving-up time. Even when 40 seconds have elapsed there is a nonnegligible portion of instances in which the pigeons' subjective time falls below the giving-up time of 22 s that is optimal for Conditions 2 to 4. According to this simple picture, we expect a slow, smooth approach of the residence profiles to asymptote, with the ogives being asymmetrical and skewed to the right, just as shown in Figure 2.

However, the model is not yet complete. The animal must estimate not one but two temporal intervals: the amount of time it has spent in a patch, t , whose variance is given by Equation 4, and the criterion time at which it should leave, t_C . When $t - t_C > 0$, the animal switches. If the animal is optimal, $t_C = t^*$. However, the representation of the criterial time must also have a variance (i.e., the vertical lines in Figure 6 should be represented as distributions). The variance of the statistic $t - t_C$ equals the sum its component variances, each given by Equation 4. Combinations of all the possible resulting models—varying all three parameters in Equation 4 and varying the relative contribution of the criterial variance—were fit to the data, and the simplest to give a good account of them all sets $a = c = 0$ and uses the same parameter b for both t and t_C . That is, we assume Poisson timing with the variance of the underlying dispersions proportional to $t + t_C$. Equation 4 then gives us the standard deviation from these two sources of variance as $\sigma_t + t_C = \sigma = \sqrt{b(t + t_C)}$.

While $t - t_C < 0$, the animal works the better patch. After the initial visit to the alternative patch at $t = t_C$, it should revisit the preferred patch, spending the proportion p of its time there and the rest in the alternative patch. Because of the spread of subjective time around real time, the average probability of being in the better patch will be an ogival function of time. For a small number of counts, the distributions will be positively skewed, resembling gamma distributions, and as the number of counts increases, they will approach normality. We may write the equation for the ogives as

$$p(P_{1,t}) = \phi(t_C - t, \sigma) + p(1 - \phi(t_C - t, \sigma)). \quad (5)$$

The first term to the right of the equal sign gives the probability of not having met the criterial count by time t , during which time the probability of being in the better patch is 1.0; the second parenthetical term gives the probability of having met the criterial count, after which the probability of being in the better patch falls to p . If the animals behave optimally, p should equal 0.50. The logistic distribution provides a convenient approximation to the normal $\Phi(t_C - t, \sigma_t)$ and is used to fit these data. The variance of the distributions is $b(t + t_C)$. This is the model that draws the curves through the data in Figures 2, 7, and 8. For the data in Figure 2, $t_C = 5.2$ s, $b = 0.08$ s, and $p = .61$.

Of course, this is not the only model of the timing process that would accommodate these data. Models such as scalar expectancy theory (Gibbon, 1977), among others, could do just as well. The point is not to test or develop a particular theory of timing. Killeen and Weiss (1987) provided a framework for many types of timing models, of which the one chosen here is among the simplest that is adequate for these data. The point is to get some use out of these timing models in addressing other substantive issues. The criterial time, t_C , provides an efficient measure of the initial commitment to a patch because it is based on all the data in the residence profiles, and it is not so greatly affected by probe responses to the alternate key. It unconfounds initial visit time from key bias p . It delineates the residence time profiles, an alternative perspective on the foraging behavior. It rules out some timing mechanisms.

We now have assembled the tools—Bayesian models of optimal performance and timing models for fallibility in estimating the Bayesian optimal—that enable us to examine these three types of deviation from optimality. The subsequent experiments use the tools in a more detailed analysis of search behavior.

Experiment 2

This experiment tests two hypotheses mentioned previously: The ogival shape of the data in Figure 2 was due to inevitable imprecision in timing, and the residual bias was a kind of “catch-up” behavior, capitalizing on the surplus probability of a payoff that was left in the better patch because of the subject's early departure from it. Conditions 1 to 3 replicate those conditions from the previous experiment, and in Condition 4 the subjects are given a cue to help them discriminate t^* . Condition 5 is a recovery of Condition 3.

Method

Subjects—Four common pigeons (*Columba livia*), all with previous experimental histories but none with experience in search tasks, were maintained at 80% to 85% of their free-feeding weights.

Apparatus—Experiments were conducted in a BRS/LVE experimental chamber. The interior of the chamber was painted black but was otherwise identical to that used in Experiment 1. The reinforcer was 2.8-s access to mixed grain followed by a 3-s blackout. White masking noise (approximately 75 dB) was continuously present.

Procedure—Sessions consisted of 60 trials, on any one of which the reinforcer was available (primed) for responses to only one of the keys. The probability that it could be obtained by responding on the better key was $p(P_1)$ and on the other key $p(P_2) = 1 - p(P_1)$. For half the subjects the better key was on the left, and for the other half it was on the right. The probabilities were arranged by randomly sampling without replacement from a table so that in each session the subjects' relative rate of payoff on the left key was exactly $p(P_1)$.

The center key was not used in this experiment. Each trial started with both side keys lit green. A probability gate was queried every second, and with probability λ set reinforcement for the next response to the primed key. Reinforcement remained set until the animal collected it or responded on the nonprimed side, in which case it was canceled. In the latter case the probability gate would again be continually queried until reinforcement was reset, and this process continued until the trial ended with reinforcement. There were no other consequences for responding on the nonprimed side. After reinforcement the chamber was darkened for a 5-s intertrial interval.

In Condition 4, the two side keys were illuminated with green light until the optimal time to switch (22 s) and then changed to red. All other aspects of this procedure were the same as in Condition 3.

Approximately 26 sessions were devoted to each of the conditions except Condition 4, which ended after 16 sessions. The conditions are identified in Table 3 by the values of $p(P_1)$ and λ that characterized them. The probability of reinforcement being set up during each second on the primed side (λ) was the same for each of the keys. Data are the probability of a response on the better key in 1-s bins, averaged over the last 14 sessions of each condition (except Condition 4, in which they were averaged over the last 10 sessions).

Results

In Condition 1, the average rate of availability of reinforcement on the primed side was $\lambda = 1/10$ (i.e., a VI 10-s schedule), and the prior probability of either side being primed was .5. The pigeons' initial giving-up time ranged from 1.5 to 3.6 s, with a mean of 2.9 s and a between-subjects standard deviation of 0.92 s (Table 3).

In Condition 2, $p(P_1) = 0.75$, $\lambda = 1/10$. In the top panel of Figure 7, the relative frequency of responses on the better key, averaged over all 4 subjects, is displayed as a function of the time into the trial. As in the first experiment, optimal behavior requires complete dedication to the better side until 11 s have elapsed and strict alternation between the sides thereafter. During the first 5 s, more than 90% of the responses were to the better side, but only one pigeon (P58) waited 11 s, on the average, to leave; overall, the animals remained 8.4 s on the better side before giving up. The subjects maintained a bias for the preferred patch throughout the rest of the trial.

In Condition 3, the prior $p(P_1) = 0.75$, and $\lambda = 1/20$, corresponding to a VI 20-s schedule on the side that was primed. In the bottom left panel of Figure 7, the relative frequency of responses on the better key, averaged over all 4 subjects, is displayed as a function of the time into the trial. The average initial giving-up time doubled to 14 s but fell short of the optimal 22 s. The individual mean giving-up times ranged from 11 to 16 s, with a standard deviation of 2.1 s. The subjects maintained a bias for the better patch throughout the rest of the trial. Note that in most cases the means of the residence profiles (t_C) are slightly greater than the mean giving-up times (see Table 3). This is because *all* responses on the better patch drive the ogives up toward 1.0, even when they are returns to that patch, whereas only the very first departure on each trial contributes to the average giving-up time and its distribution.

In Condition 4 the parameters were the same as for Condition 3 except that the keys changed colors when the optimal wait time had elapsed. In the top right panel of Figure 7, the relative frequency of responses on the better key, averaged over all 4 subjects, is displayed as a function of the time into the trial. The behavior of 3 of the 4 pigeons very closely approximated the ideal, shown by the step function. Despite a slow decrease in the probability of staying on the better patch, that probability remained above 90% through the first 21 s. On change of keylights, the subjects immediately switched to the alternate key and then soon switched back, falling into a pattern of regular alternation thereafter, spending just 51% of their time on the preferred key. The behavior of Subject P11 was not entrained by the keylight change. That pigeon had had the earliest average departure time in the previous condition (11 s), which increased to only 13.3 s in this condition.

Having been shaped to the optimal behavior, will the pigeons persist in it when the signals are removed? The bottom right panel of Figure 7 shows that in general they did not. Of the 3 pigeons that behaved optimally in the signaled condition, Subject P9 remained closest to optimal, dropping from 22.0 s to 19.9 s. The other two shifted to mean departure times of 15 s, about 1 s longer than in Condition 3. Subject P11 dropped from 13.3 s under the signal to 11.8 s in this condition. The subjects maintained a bias throughout the rest of the trial. The giving-up times were not significantly different than those found in Condition 2, $t(3) = 2.20$, $p < .12$, and remained significantly below the optimal, $t(3) = -3.97$, $p < .029$. The values of b in Equation 4 were 0.12, 0.04, 0.001, and 0.04 for the curves drawn in Figure 7.

This analysis is reinforced by Figure 8, which shows the median (across-subjects) distributions of giving-up times, along with the best fitting gamma densities. These times are measured as the average of the time of the last response on the better patch (the giving-in time) and the first response on the alternate patch (the moving-on time). For Condition 1, the data are organized with respect to the side preferred on the average, which varied idiosyncratically among subjects. The shape of the data for individual subjects looks similar to these median curves, whereas the mean and variance were characteristically different for each subject. The data were equally well fit with Gaussian densities. These curves are essentially the negative derivatives of the residence profiles shown in the previous figure.

Discussion

The ogival residence profile displayed in the top left panel of Figure 7 is similar to that shown in Figure 2. As in the first experiment, most of the animals left the better patch prematurely, averaging 8.5 s residence before switching (see Table 3). The smooth curves are from Equations 4 and 5; in all panels except the top right one, $p = .57$. The top right panel shows that the change in keylight color strongly controlled switching, with 3 of the 4 pigeons averaging a 21.1-s residence in the better patch before leaving. These pigeons could learn the local properties of these schedules; they just could not time the shift in those properties accurately without the help of an external clock. Even the clock was of little help to Subject P11, which contributed most of the “ramp”, leading down to 22 s, to the average data. Subject P11's giving-up time did increase from 8 to 13.3 s, but this is still far short of optimality.

These results were impressive even though not unexpected. (In an earlier unreported experiment, we achieved a similar elimination of residual bias in the subjects from Experiment 1 by turning off the keylight on the inferior patch until t^* .) Three of the 4 pigeons agreed that our definition of t^* suited their purposes. It is an interesting question whether we could have brought Subject P11 under stimulus control if we had started the duration of the cues at its criterial value and then slowly lengthened them.

This is not the first experiment to show that stimuli that clarify contingencies can have such strong control of behavior that they create the conditions for their own maintained control, an experimental version of a self-fulfilling prophecy (e.g., Killeen, 1989; Killeen & Snowberry, 1982). It exemplifies in learned behavior the near-inexorable control of arbitrary sign stimuli, such as those that dominate sexual selection.

Notice that in Condition 4 after the immediate sampling of the alternate patch the animals settled down to indifference between the patches. This indifference suggests that the continued apparent bias for the better key shown in the other panels of this figure was a by-product of imprecise timing of t^* ; when given a precise cue, animals became precisely indifferent in the long run (the right tread of the step in that panel is drawn at $p = .50$).

In Condition 5, we reverted to continuous green lights, which in Condition 4 governed exclusive preference of the better alternative in 3 of the 4 pigeons. Our hope was to maintain that control through this replication. Despite the therapy effected by these cues, however, when returned to the continuous green lights in Condition 5, the pigeons reverted to their uniform early departures; Condition 4 conferred an average improvement of only 1.6 s on subsequent performance.

As to the particular timing model we used, a few others, vesting all of the pigeons' temporal uncertainty in either the criterial time or in the real time or using a timing model in which the variance grew as the square of the criterial time (only a in Equation 4 greater than zero), would have sufficed and would have told similar stories. These alternate versions are within a few percentage points of accuracy in fitting the data. In one case (only a greater than 0), p may be set to .50, which suggests that all of the apparent residual bias is due to the animal's extended confusion as to the real time (see Figure 6). However, we rejected that interpretation because graphs of average dwell time per visit appear similar to Figure 3, indicating the residual bias is real and continuing, not a by-product of continued skew in the ogives.

Experiment 3

This experiment tests the generality of these findings by changing the contingencies of reinforcement from time based to response based.

Method

Subjects and Apparatus—The subjects and apparatus were the same as in Experiment 2. The keylights were both green.

Procedure—All contingencies were the same as in Conditions 3 and 5 from Experiment 2, except that now each response in the primed patch had a probability of .05 of being reinforced. In the previous experiments, a time-varying probability was resampled with each peck. In the present experiment, the probability is time invariant, with 1 in 20 responses in the primed patch reinforced (a “random ratio 20” schedule). A 3-s blackout followed each reinforcer. No responses in the unprimed patch were reinforced. The purpose of this manipulation was to determine whether the change in currency from search time to search responses would have a substantial impact on the animals' strategy. Now the optimal strategy is to stay on the better patch until 22 responses have been made and then switch to the other patch, alternating regularly thereafter.

Results

The pigeons quickly adapted to the new contingencies, but there was a much wider spread in the giving-up times than under previous conditions. Figure 9 shows the residence profiles for each of the subjects. In these experiments the x axis is now number of responses, not time.

We may still use the timing model for this experiment, assuming that it is key pecks that are accumulated by the counter, not pulses from a pacemaker. All timing-counting models were within a few points of one another. Most subjects showed residual bias for one or the other patch, no matter which model of counting was used. Figure 9 shows the fit of the model used in the previous experiments, with only $b > 0$. The values for b are smaller than those found for the time-based currencies, averaging $0.02 r^{-1}$.

Two of the subjects were very close to the optimal initial residence run ($r_c = 21.5$ and 22.2 responses), and 2 fell short (13.5 and 7.1 responses); Subject P11 again was the most impulsive. The average was closer to the ideal than was the average initial residence time under the time-based currencies of the previous condition. The large individual differences and continuing bias somewhat undermine the significance of this result: Although the average run length moved closer to the ideal, the variance of individual averages was greater than for time-based performances.

It could be argued that the ratio contingencies forced faster responding, but the pigeons, out of a long history in this experiment, still based their first departure on time spent responding. If the average of the residence profiles of Figure 9 were graphed on a time axis, it would approximately superimpose with those of the top left panel of Figure 7, supporting this argument. The spread between the ogives of individual subjects becomes the same in both conditions when plotted as a function of elapsed time. The chief difference is that the redrawn ogives from the present condition are somewhat steeper, and on the average asymptote closer to 0.5 than they do for the time-based contingencies of the previous experiment. The hypothesis that time, not responses, is the intrinsic currency could be settled with experiments that block responding by darkening the keys briefly; but that is not assayed here, where the issue is left unsettled.

General Discussion

The optimality approach is a curious one because, unlike a model of behavior, it is impossible to prove it false: The pigeon's lack of optimality in the present experiment is no threat to the canonical theory of optimal search, which guided the experiment's design. Theories of

optimality are not theories about organisms; they are theories about contexts and systems. All organisms—from pigeons and pilots to trout fisherman—would be best off behaving according to Equation 3. The present analysis demonstrates that under some very general assumptions the search behavior of pigeons is nonoptimal; it rules out numerous explanations of that nonoptimality; it provides a Bayesian framework to evaluate models of biased optimization and a general model of timing and counting to account for variability. These permit us to distinguish different sources and signatures of bias; as those become certainly identified, we may move from this qualified (biased) optimization model to an unqualified one in which the constraints become conditionalized as parameters or boundary conditions. In the larger view, however, “optimal |” should always be written with a vertical bar, because it is not a theory of behavior until one begins to list the “givens” under which optimality is to be tested, the constraints of memory, accuracy in timing, and histories of reinforcement. It is in just such a clarification of the operative constraints that the optimality approach finds its value.

There are other approaches, ones that begin with known behavioral processes and parameters and apply those to real or simulated foraging scenarios (e.g., Brunner et al., 1992; Fantino & Abarca, 1985; Shettleworth, 1988). The task then becomes one of relating systematic differences between species to differences in their niches in ways that demonstrate improved fitness. This approach focuses on optimization as a process of adaptation rather than on the optimal, the hypothetical endpoint of that process. For instance, R. A. Johnson, Rissing, and Killeen (1994) showed that seed-harvesting ants whose foraging pattern exposed them to greater variability in seed types had faster learning and forgetting functions than another species that experienced less variability. Both this bottom-up approach and the top-down approach of optimality are useful; it is when they converge, showing how system parameters (rather than particular performances) are adapted to move organisms toward an optimal, that they provide the most satisfying story (Kamil & Mauldin, 1988).

Acknowledgements

This research was supported in part by National Science Foundation Grants BNS 9021562 and IBN 94-08022 and National Institute of Mental Health Grant R01 MH 48359. Experiment 1 was Gina-Marie Palombo's honors thesis.

Appendix

Derivation of the Optimal

Bayesian Analysis

The abbreviated proof of Equation 3 in the text permitted a simple rule of thumb: Switch when the marginal rates of return from the two patches become equal, at which point the probabilities of obtaining food are identical for the two patches for all future time. This rule of thumb yields the optimal tactic—exclusive preference until t^* and rapid alternation thereafter—only when the detection function is exponential. A more general solution is given here. Bayes's theorem is the analytic model of choice because it permits continuous updating of probabilities as a function of new information (Green, 1980; McNamara & Houston, 1980). In the present case, the new information is the failure to detect the prey in a patch with continued searching.

Problem: Find the posterior probability that food is in fact available in Patch i , given that it has not been found there by time t : $p(P_i|F_{i,t})$. This is the statistic on which the optimal forager must base its decision to leave the patch.

Given the following:

$p(P_i) \equiv$ The prior probability of food being primed for Patch i at the start of any trial; the base rates.

$p(F_{i,t}) \equiv$ The probability of finding food in Patch i by time t .

$p(F_{i,t}|P_i) = 1 - e^{-\lambda t} \equiv$ The detection (gain) function: the conditional probability of finding food in Patch i by time t , given that it has been primed for Patch i . Lambda (λ) is the rate constant of the detection function, and its reciprocal is the mean time between detections

Then we may find the unknown (posterior) by use of Bayes's theorem:

$$p\left(P_i|\tilde{F}_{i,t}\right) = p\left(\tilde{F}_{i,t}|P_i\right) \times \frac{p(P_i)}{p\left(\tilde{F}_{i,t}\right)}. \quad (\text{A1})$$

To evaluate Equation A1, we must write the equations for the knowns:

1. The probability of not finding food in Patch i by time t , given that it has been primed for that patch, is simply the complement of the detection function:

$$p\left(\tilde{F}_{i,t}|P_i\right) = 1 - p\left(F_{i,t}|P_i\right) = e^{-\lambda t}. \quad (\text{A2})$$

2. The numerator of Equation A1 is simply the prior for that patch.
3. To obtain the denominator of Equation A1, express the unconditional probability of finding food in terms of the conditional probability and the priors:

$$p(F_{i,t}) = p(P_i) p(F_{i,t}|P_i) = p(P_i) (1 - e^{-\lambda t}). \quad (\text{A3})$$

The probability of not finding food is the complement of Equation A3:

$$p\left(\tilde{F}_{i,t}\right) = 1 - p(F_{i,t}) = 1 - p(P_i) (1 - e^{-\lambda t}). \quad (\text{A4})$$

Inserting Equations A2 and A4 into Equation A1, we rearrange to arrive at the key prediction:

$$p\left(P_i|\tilde{F}_{i,t}\right) = \frac{p(P_i)}{(p(P_i) + (1 - p(P_i)) e^{\lambda t})}. \quad (\text{A5})$$

Note that the boundary conditions of this equation accord with common sense:

When $t = 0$, $p(P_i|\tilde{F}_{i,t}) = p(P_i)$; the probability of finding food in the patch just as you start to search is simply the prior probability, $p(P_i)$.

As $t \rightarrow \infty$, $p(P_i|\tilde{F}_{i,t}) \rightarrow 0$; if you have not found it after searching forever, it simply is not there.

Optimal behavior consists in switching the instant after the posteriors for the two patches become equal. Because there are two patches and the food must be in one and only one of them—the priors sum to 1.0, so that $1 - p(P_1) = p(P_2)$ —this happens when $p(P_i|\tilde{F}_{i,t}) = 0.50$. Solving Equation A5 for the time t^* at which this posterior for the better patch has fallen to 0.50 gives

$$t^* = \ln [p(P_1) / p(P_2)] / \lambda. \quad (\text{A6})$$

The optimal time to switch is the log-likelihood ratio of the priors multiplied by the average time between reinforcers ($1/\lambda$) and appears as Equation 3 in the text.

Logistic Distribution

If we divide the right-hand side of Equation A5 by the priors, it starts to become recognizable as a logistic equation:

$$p\left(P_i|\tilde{F}_{i,t}\right) = \left(1 + [(1 - p(P_i)) / p(P_i)] e^{\lambda t}\right)^{-1}. \quad (\text{A7})$$

Look at the coefficient of the exponential. It is the likelihood ratio of the priors. Because in this experiment the probabilities of one or the other of the patches being primed sums to one, $1 - p(P_1) = p(P_2)$ and we may write the logarithm of that ratio as $-\ln[p(P_1)/p(P_2)]$. From Equation A6 we see that this is just λt^* . Therefore, we may rewrite Equation A7 as

$$p\left(\tilde{P}_i | \tilde{F}_{i,t}\right) = \left(1 + e^{\lambda(t-t^*)}\right)^{-1}. \quad (\text{A8})$$

This is the classic form of the logistic distribution. It describes the changing posterior probabilities of knowing that food is primed in the patch and takes as its mean t^* the optimal time to switch. It is drawn in Figure 4 for two values of lambda. Notice that these ogives are much flatter than the residence profiles, and they intercept the y-axis at the prior for the better patch; this discrepancy shows that animals are not matching the distribution of their responses to the relative probabilities of reinforcement through the trial.

Discounting the Priors

Missattribution of Reinforcement

Assume that organisms correctly attribute the source of reinforcement on less than 100% of the trials, so that the fraction m of reinforcers on Patch i will be missattributed to an alternate patch, whereas the fraction $(1 - m)$ of reinforcers on the alternate patch will be missattributed to Patch i . Then the *subjective* prior for Patch i , p_i , is

$$p_i = (1 - m) p(P_i) + m(1 - p(P_i)). \quad (\text{A9})$$

We can see from Equation A9 that the better patch has more to lose by this type of error and the poorer patch more to gain. This blurs the distinctiveness of the priors, effectively discounting them.

Some algebra gives us the likelihood ratio of subjective priors for Patch i :

$$\frac{p_i}{1 - p_i} = \frac{p(P_i) + m/(1 - 2m)}{1 - p(P_i) + m/(1 - 2m)}, m \neq 0.5. \quad (\text{A10})$$

When there is no missattribution ($m = 0$), the subjective priors equal the objective priors; as $m \rightarrow 0.5$, the likelihood ratio is discounted, approaching 1.0 in the limit. Still smaller values of gamma betoken a situation similar to the case of negative d 's in signal-detection theory: The animal has its cues reversed. Equation A10 may be inserted into Equation A6 to derive expected giving-up times (as was done to generate the last column of Tables 2 to 4) or into Equation A7 or Equation A8 for the complete residence profile. For Experiment 1, $m = 0.18$ generated the predictions in Table 2; for Experiment 2, $m = 0.12$. These were the values of m that minimized the sum of squares deviation between the obtained data and the predictions.

Discounting

A more generic manner of discounting the priors is to raise the likelihood ratio to some weight w less than 1. This is equivalent to multiplying the log-likelihood ratios by w , $L = w \ln[p(P_1)/p(P_2)]$, so that Equation A6 becomes

$$t^* = w \ln[p(P_1)/p(P_2)] / \lambda. \quad (\text{A11})$$

In the discussion we referred to the tactic of assuming all priors equal as an example of the representativeness fallacy. These equations give us a more general model for that bias, in which the priors may be discounted by any proportion w . Whereas values of $w < 1$ discount the priors, values of $w > 1$ place a disproportionate emphasis on them (i.e., discount the new information gained from the unrequited time in a patch on each trial), as might be the case if the rates of return of the patches, w , are themselves unreliable.

These alternative modes of discounting the priors give quantitatively very similar results; deciding between them will require qualitative experimental manipulations and a more revealing interpretation of how the discount fraction in the generic model (Equation A11) serves specific proximate goals.

Generalizations

Other Detection Functions

The use of such exponential detection functions was motivated by our desire to provide a paradigm whose optimal solution (dwell until t^* and then be indifferent) was as simple as possible and for which simple rules of thumb were available to guide selection of t^* (switch when marginal rates of return in the patches are equal). It is convenient that the exponential detection function also characterizes the popular constant-probability variable interval Schedule of reinforcement because it supports the simplest theoretical analysis.

Nature or experimenters may provide different detection functions: depleting patches, periodically replenished patches, patches with temporally clustered prey (see, e.g., Cassini, Lichtenstein, Ongay, & Kacelnik, 1993; D. F. Johnson, Tribblehorn, & Collier, 1993). In such cases one may substitute appropriate detection functions for $p(F_{i,t}|P_i)$ in Equation A1 and solve for the predictions of optimal behavior. In all cases the optimal behavior will be to dwell in the better patch until some time t^* ; but subsequently optimal behavior will depend on the exact form of the detection function.

Kamil, Yoerg, and Clements (1988) gave blue jays the choice between a nondepleting patch with a constant probability of finding prey of $1/4$ and a “sudden-death” patch, in which the probability of finding prey was constant at $1/2$ until N prey had been captured and then went to 0. Optimal behavior would, of course, be to stay until the N th capture and then switch, because the probability on the depleting patch goes from $.5$ to 0 at that point. If the animals could count, we would expect a step function. However, animals' ability to count is less than perfect, changing as a function of the number of counts in a similar way to their ability to time (Equation 4). Therefore, useful information is also contributed by the length of runs since a capture. To model this, we conditionalize on both total prey captured on a side (N) and on the number of trials without a capture (n), setting up Equation A1 to evaluate $p(P_i|F_{i,n} \cdot (N = j))$.

As Kamil et al. (1988) noted, “although both these factors have been proposed as rules-of-thumb for patch departure, they have often been treated as mutually exclusive. The data reported here are the first evidence of which we are aware for the joint use of both factors” (p. 851; similar results are reported by Roberts, 1993). Equation A1 may complement this empirical work by providing a coherent theoretical treatment of such information integration.

Other Rates of Return

In the present experiment, the mean times to detection on the primed side ($1/\lambda$) are equal for the two patches. All proofs may be generalized by appropriately subscripting those parameters in the derivations found here.

References

- Baum WM. Optimization and the matching law as accounts of instrumental behavior. *Journal of the Experimental Analysis of Behavior* 1981;36:387–403. [PubMed: 16812255]
- Bell, WJ. Searching behaviour: The behavioural ecology of finding resources. Chapman & Hall; New York: 1991.
- Brunner D, Kacelnik A, Gibbon J. Optimal foraging and timing processes in the starling, *Sturnus vulgaris*: Effect of inter-capture interval. *Animal Behavior* 1992;44:597–613.

- Cassini MH, Lichtenstein G, Ongay JP, Kacelnik A. Foraging behavior in guinea pigs: Further tests of the marginal value theorem. *Behavioural Processes* 1993;29:99–112.
- Charnov EL. Optimal foraging: The marginal value theorem. *Theoretical Population Biology* 1976;9:129–136. [PubMed: 1273796]
- Cialdini, RB. Full-cycle social psychological research. In: Beckman, L., editor. *Applied social psychology annual*. 1. Sage; Beverly Hills, CA: 1980. p. 21-47.
- Cialdini, RB. A full-cycle approach to social psychology. In: Brannigan, GC.; Merrens, MR., editors. *The social psychologists*. McGraw-Hill; New York: 1995. p. 52-72.
- Commons, ML.; Kacelnik, A.; Shettleworth, SJ. Quantitative analysis of behavior. VI: Foraging. Erlbaum; Hillsdale, NJ: 1987.
- Davison M, Jones BM. A quantitative-analysis of extreme choice. *Journal of the Experimental Analysis of Behavior* 1995;64:147–162. [PubMed: 16812765]
- Davison M, Kerr A. Sensitivity of time allocation to an overall reinforcer rate feedback function in concurrent interval schedules. *Journal of the Experimental Analysis of Behavior* 1989;51:215–231. [PubMed: 16812577]
- Davison, M.; McCarthy, D. The matching law: A research review. Erlbaum; Hillsdale, NJ: 1988.
- Dreyfus LR, DePorto-Callan D, Pseillo SA. Changeover contingencies and choice on concurrent schedules. *Animal Learning & Behavior* 1993;21:203–213.
- Fantino E, Abarca N. Choice, optimal foraging, and the delay-reduction hypothesis. *Behavioral and Brain Sciences* 1985;8:315–362.
- Fetterman JG, Killeen PR. Time discrimination in *Columba livia* and *Homo sapiens*. *Journal of Experimental Psychology: Animal Behavior Processes* 1992;18:80–94. [PubMed: 1578202]
- Gibbon J. Scalar expectancy theory and Weber's law in animal timing. *Psychological Review* 1977;84:279–325.
- Gibbon J, Church RM. Time left: Linear versus logarithmic subjective timing. *Journal of Experimental Psychology: Animal Behavior Processes* 1981;7:87–108. [PubMed: 7241054]
- Green RF. Bayesian birds: A simple example of Oaten's stochastic model of optimal foraging. *Theoretical Population Biology* 1980;18:244–256.
- Hinson JM, Staddon JER. Matching, maximizing, and hill-climbing. *Journal of the Experimental Analysis of Behavior* 1983;40:321–331. [PubMed: 16812350]
- Houston, AI. The control of foraging decisions. In: Commons, ML.; Kacelnik, A.; Shettleworth, SJ., editors. *Quantitative analysis of behavior. VI: Foraging*. Erlbaum; Hillsdale, NJ: 1987. p. 41-61.
- Houston AI, McNamara J. How to maximize reward rate on two variable-interval paradigms. *Journal of the Experimental Analysis of Behavior* 1981;35:367–396. [PubMed: 16812223]
- Johnson DF, Triplehorn J, Collier G. The effect of patch depletion on meal patterns in rats. *Animal Behaviour* 1993;46:55–62.
- Johnson RA, Rissing SW, Killeen PR. Differential learning and memory by co-occurring ant species. *Insectes Sociaux* 1994;41:165–177.
- Kamil, AC.; Krebs, JR.; Pulliam, HR., editors. *Foraging behavior*. Plenum Press; New York: 1987.
- Kamil, AC.; Mauldin, JE. A comparative-ecological approach to the study of learning. In: Bolles, RC.; Beecher, MD., editors. *Evolution and learning*. Erlbaum; Hillsdale, NJ: 1988. p. 117-133.
- Kamil AC, Yoerg SI, Clements KC. Rules to leave by: Patch departure in foraging blue jays. *Animal Behaviour* 1988;36:843–853.
- Killeen PR. Preference for fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior* 1970;14:127–131. [PubMed: 16811462]
- Killeen, PR. Behavior as a trajectory through a field of attractors. In: Brink, JR.; Haden, CR., editors. *The computer and the brain: Perspectives on human and artificial intelligence*. Elsevier; Amsterdam: 1989. p. 53-82.
- Killeen, PR. Counting the minutes. In: Macar, F.; Pouthas, V.; Friedman, WJ., editors. *Time, action and cognition*. Kluwer Academic; Norwell, MA: 1992. p. 203-214.
- Killeen PR. Economics, ecologies, and mechanics: The dynamics of responding under conditions of changing motivation. *Journal of the Experimental Analysis of Behavior* 1995;64:405–431. [PubMed: 16812776]

- Killeen PR, Fetterman JG. A behavioral theory of timing. *Psychological Review* 1988;95:274–295. [PubMed: 3375401]
- Killeen PR, Smith JP. Perception of contingency in conditioning: Scalar timing, response bias, and the erasure of memory by reinforcement. *Journal of Experimental Psychology: Animal Behavior Processes* 1984;10:333–345.
- Killeen PR, Snowberry K. Information and cooperative behavior. *Behaviour Analysis Letters* 1982;2:353–360.
- Killeen PR, Weiss N. Optimal timing and the Weber function. *Psychological Review* 1987;94:455–468. [PubMed: 3317471]
- Koopman BO. The theory of search: III. The optimum distribution of searching effort. *Operations Research* 1957;5:613–626.
- Koopman, BO. Search and screening. Pergamon Press; New York: 1980.
- Krebs, JR.; Davies, NB., editors. *Behavioural ecology: An evolutionary approach*. Blackwell Scientific; Oxford: 1978.
- Krebs, JR.; Davies, NB., editors. *Behavioural ecology: An evolutionary approach*. 3rd ed.. Blackwell; Oxford: 1991.
- Krebs JR, Kacelnik A, Taylor P. Optimal sampling by foraging birds: An experiment with great tits (*Parus major*). *Nature* 1978;275:27–31.
- Lea, SEG. Correlation and contiguity in foraging behavior. In: Harzem, P.; Zeiler, MD., editors. *Advances in the analysis of behaviour 2: Predictability, correlation, and contiguity*. Wiley; New York: 1981. p. 355–406.
- Lima SL. Downy woodpecker foraging behavior: Efficient sampling in simple stochastic environments. *Ecology* 1984;65:166–174.
- Mazur JE. Optimization theory fails to predict performance of pigeons in a two-response situation. *Science* 1981;214:823–825. [PubMed: 7292017]
- McNamara JM, Houston AI. The application of statistical decision theory to animal behavior. *Journal of Theoretical Biology* 1980;85:673–690. [PubMed: 7442286]
- McNamara JM, Houston AI. Optimal foraging and learning. *Journal of Theoretical Biology* 1985;117:231–249.
- McNamara JM, Houston AI. Memory and the efficient use of information. *Journal of Theoretical Biology* 1987;125:385–395. [PubMed: 3657218]
- Nevin, JA. Psychophysics and reinforcement schedules: An integration. In: Commons, ML.; Nevin, JA., editors. *Quantitative analysis of behavior. 1: Discriminative properties of reinforcement schedules*. Ballingen; Cambridge, MA: 1981. p. 3–27.
- Real LA. Objective benefit versus subjective perception in the theory of risk-sensitive foraging. *The American Naturalist* 1987;130:399–411.
- Real LA. Animal choice behavior and the evolution of cognitive architecture. *Science* 1991;253:980–986. [PubMed: 1887231]
- Rescher, N. *Peirce's philosophy of science*. University of Notre Dame Press; Notre Dame, IN: 1978.
- Roberts WA. Testing a stochastic foraging model in an operant simulation: Agreement with qualitative but not quantitative predictions. *Journal of the Experimental Analysis of Behavior* 1993;59:323–331. [PubMed: 16812688]
- Shettleworth, SJ. Foraging as operant behavior and operant behavior as foraging: What have we learned?. In: Bower, GH., editor. *The psychology of learning and motivation*. Academic Press; New York: 1988. p. 1–49.
- Shettleworth SJ. Animals foraging in the lab: Problems and promises. *Journal of Experimental Psychology: Animal Behavior Processes* 1989;15:81–87.
- Staddon JER, Hinson JM, Kram R. Optimal choice. *Journal of the Experimental Analysis of Behavior* 1981;35:397–412. [PubMed: 16812224]
- Staddon JER, Horner JM. Stochastic choice models: A comparison between Bush-Mosteller and a source-independent reward-following model. *Journal of the Experimental Analysis of Behavior* 1989;52:57–64. [PubMed: 16812590]

- Stephens DW, Dunbar SR. Dimensional analysis in behavioral ecology. *Behavioral Ecology* 1993;4:172–183.
- Stephens, DW.; Krebs, JR. Foraging theory. Princeton University Press; Princeton, NJ: 1986.
- Stone, LD. Theory of optimal search. Academic Press; New York: 1975.
- Sutherland GD, Gass CL. Learning and remembering of spatial patterns by hummingbirds. *Animal Behaviour* 1995;50:1273–1286.
- Templeton AR, Lawlor LR. The fallacy of averages in ecological optimization theory. *American Naturalist* 1981;117:390–393.
- Williams, BA. Reinforcement, choice, and response strength. In: Atkinson, RC.; Herrnstein, RJ.; Lindzey, G.; Luce, RD., editors. *Stevens' handbook of experimental psychology*. Wiley; New York: 1988. p. 167-244.
- Williams BA. Choice as a function of local versus molar reinforcement contingencies. *Journal of the Experimental Analysis of Behavior* 1991;56:455–473. [PubMed: 1774539]
- Zeiler MD. On optimal choice strategies. *Journal of Experimental Psychology: Animal Behavior Processes* 1987;13:31–39.
- Zeiler MD. To wait or to respond. *Journal of the Experimental Analysis of Behavior* 1993;59:433–444. [PubMed: 16812692]

Figure 1.

The probability of reinforcement as a function of time. The dashed curve shows the conditional probability of reinforcement as a function of time in either patch, given that reinforcement is scheduled for that patch. The middle and bottom curves show the unconditional probability of reinforcement in Patches 1 and 2, in which the priors are 0.75 and 0.25, respectively. Note that if an animal has not received reinforcement in Patch 1 by 11 s, the residual probability of reinforcement (25%; the distance from 0.5 to 0.75) exactly equals that available from Patch 2. Furthermore, at that point the distributions are congruent: The curve for Patch 1 between the ordinates of 0.5 and 0.75 is of precisely the same form and scale as that for Patch 2 between the ordinates of 0.0 and 0.25. All future prospects are identical. Therefore, after exploring Patch 1 for 11 s, the forager should become indifferent and thereafter treat the two patches as identical.

Figure 2.

The proportion of responses in the better patch as a function of time through the trial in Condition 2. The circles show the average data from 7 pigeons, and the step function shows the optimal behavior. The smooth curve is drawn by Equations 4 and 5, a Poisson model of the timing process described later in the text. Residence profiles from individual subjects resembled the average (see, e.g., Figure 9).

Figure 3.

The duration of responding to a key as a function of the ordinal number of the visit to that key. The data are averaged over 7 pigeons in Condition 2 and correspond to the data shown in Figure 2. The first datum shows the initial giving-up time for the first visit to the better (75%) key. Optimally the first visit should last for 11 s, corresponding to the abscissa of the riser on the step function shown in Figure 2, and thereafter the visits should be of equal and minimal duration. The error bars are standard errors of the mean; because of the large database, they primarily reflect small differences in dwell times characteristic of different subjects.

Figure 4.

The posterior probabilities that food is primed for the a priori better patch as a function of time spent foraging in it, for discovery rates of $\lambda = 1/10$ and $\lambda = 1/20$.

Figure 5.

The rates of reinforcement obtained by dwelling in the preferred patch for various durations before switching to unbiased sampling of the patches. The data are from simulations of responding, averaged over 10,000 trials.

Figure 6.

Hypothetical dispersions of subjective time around 5, 10, 20, 30, and 40 s of real time. The distributions assume scalar timing; the standard deviations are proportional to the means of the distributions. The vertical bars mark the optimal switch points in the major conditions of this study.

Figure 7.

Residence profiles for Conditions 2 (top left panel), 3 (bottom left panel), 4 (top right panel), and 5 (bottom right panel), averaged over the 4 subjects. The ogives are drawn by Equations 4 and 5.

Figure 8.

The probability of leaving a patch as a function of the time spent foraging during that trial. The data are the medians across subjects. (VI = variable interval.)

Figure 9.

Residence profiles for each of the subjects in Experiment 3. The ogives are the same model as used for Figures 2 and 7 but now are based on the counting of responses rather than units of time. P = pigeon.

Table 1

Conditions of Experiment 1

Condition	λ	$p(P_1)$	N	t_1	σ_{t_1}	t_2	t_1 (2nd visit)
1	0.100	0.50	7	3.05	1.38	2.52	3.00
2	0.100	0.75	7	6.71	0.92	1.63	3.59
3	0.050	0.75	3	13.2	2.49	1.54	5.29
4	0.025	0.75	2	22.0	3.44	2.73	11.30
5	0.106	0.33	4	4.00	0.75	1.85	2.54
6	0.100	0.67	4	5.50	0.58	1.44	3.98

Note. λ = the probability of reinforcement during each second of searching; $p(P_1)$ = the prior probability of reinforcement in Patch 1; N = the number of subjects per condition; t_1 = the initial giving-up times; t_2 = the second giving-up times; σ_{t_1} = their standard deviations over subjects, and the subsequent visit durations.

Table 2
Optimal and Obtained Giving-Up Times and the Predictions of the Bayesian Model With Discounted Priors

Condition	λ	$p(P_1)$	t^*_{Opt}	t_{Opt}	t^*_{Dis}
1	0.100	0.50	1.0 ^a	3.09	1.00 ^a
2	0.100	0.75	11.0	6.71	6.64
3	0.050	0.75	22.0	13.20	13.30
4	0.025	0.75	43.9	22.00	26.50
5,6	0.010	0.67	6.9	4.74	4.35

Note. λ = the probability of reinforcement during each second of searching; $p(P_1)$ = the prior probability of reinforcement in Patch 1, Opt = optimal; Obt = obtained; Dis = discounted.

^a All models predict minimal dwell times on each side in this condition.

Table 3

Conditions of Experiment 2

Condition	λ	$p(P_1)$	t_I	t_C	t_{Dis}^a
1	0.10	0.50	2.90	2.57	1.0 ^a
2	0.10	0.75	8.45	8.44	8.0
3	0.05	0.75	14.00	15.70	16.0
4	0.05	0.75	19.40	21.50	—
5	0.05	0.75	15.50	16.70	16.0

Note. λ = the probability of reinforcement during each second of searching; $p(P_1)$ = the prior probability of reinforcement in Patch 1; t_I = the mean initial giving-up time; t_C = the mean of the residence profiles (see Figures 2 and 7) and the predicted mean with discounted (Dis) priors (missattribution error of 12%).

^a All models predict minimal dwell times on each side in this condition.