

Quantifying the Adaptive Value of Learning in Foraging Behavior

Sigrunn Eliassen,^{1,*} Christian Jørgensen,¹ Marc Mangel,² and Jarl Giske¹

1. Department of Biology, University of Bergen, P.O. Box 7803, N-5020 Bergen, Norway; 2. Department of Applied Mathematics and Statistics, University of California, Santa Cruz, California 95064

Submitted March 2, 2009; Accepted April 16, 2009; Electronically published August 20, 2009

Online enhancement: appendix.

ABSTRACT: The value of acquiring environmental information depends on the costs of collecting it and its utility. Foragers that search for patchily distributed resources may use experiences in previous patches to learn the habitat quality and adjust their behavior. We map the ecological landscape for the evolution of learning under a range of conditions, including both spatial and temporal heterogeneity. We compare the learning strategy with genetically fixed patch-leaving rules and with strategies of foragers that have free and perfect information about their environment. The model reveals that the efficiency of learning is highest when low encounter stochasticity results in reliable estimates of patch quality, when there is no or little temporal change, and when there is little spatial variability. This partially contrasts with the value of learning, which is highest when there is temporal change, because flexible strategies may track the environmental trend, and when there is spatial variability, because there is a need to distinguish between good and bad patches. Learning rules with short-term memory are beneficial when patch information is accurate and when there is temporal change, whereas learning rules that update slowly are generally more robust to spatial variability.

Keywords: learning, patch-time allocation, foraging, information value.

Introduction

Foraging conditions vary in both time and space as a result of fluctuations in physical and biotic factors. Such changes may follow predictable patterns—for example, as they vary throughout the annual cycle—or fluctuations can be more stochastic. To illustrate this, imagine an arctic fox searching the snow-covered tundra for food in early spring. If lucky, it may encounter a reindeer that has starved to death. A reindeer carcass represents a large but rare food package, and to the fox such an encounter may be a highly stochastic event. There might be periods and areas where reindeer

are more likely to die from starvation, but would we expect that the fox try to predict locations where starved reindeer fall dead? Contrast this with the situation a few months later, when the same fox feeds on eggs and chicks on a bird cliff. Some parts of the cliff may be easier to hunt, and different seabird species may vary in their antipredator behavior and their timing of egg laying and fledging. Successful encounters may now be relatively frequent and may vary in a more systematic manner between areas and over time. Would we predict that the fox try to learn these patterns? And would it adjust its behavior as more birds populate the cliff in early spring or when the chicks fledge some months later? In this article, we present a method for quantifying the adaptive value of learning, which we use to map the general ecological conditions under which learning is favored.

There are several theoretical models that describe how resource patches are most efficiently exploited, but they often do not explain how the forager can arrive at such optimal behavior. Among the classics are the marginal-value theorem, which addresses how long a forager stays in a gradually depleting patch (Charnov 1976), and the ideal free distribution model, which predicts how foragers should distribute in a heterogeneous habitat (Fretwell and Lucas 1970). Common to these two models is that they assume that foragers are able to perfectly assess the spatial heterogeneity and alter their behavior accordingly. For example, the marginal-value theorem predicts that a forager should leave a patch when the instantaneous intake rate in the patch drops to the average intake rate in the habitat (Charnov 1976). But how can a forager know what the average foraging rate in the habitat is? And how does it measure its own intake rate if the food arrives through stochastic encounters? For this, either learning or genetically coded preconceptions of the environment are needed.

The marginal-value theorem highlights that foragers need information on two hierarchical levels to decide when to leave a patch: information about the current intake in

* Corresponding author; e-mail: sigrunn.eliassen@bio.uib.no.

the patch must be compared with information about the average intake rate in the habitat as a whole. In this article, we study how information about the patch and information about the habitat are linked and their influence on the evolutionary landscape of learning. Under most natural conditions, both patch information and habitat information are associated with uncertainty, and they pose different but interlinked challenges. In general, models and empirical studies of patch allocation behavior have considered only one of these information problems at a time. The main focus has been on the way foragers use search time and resource encounters to estimate the quality of a patch (Oaten 1977; Iwasa et al. 1981; McNamara 1982; Green 1984; Valone 1991; Olsson and Holmgren 1998; van Alphen et al. 2003; Olsson and Brown 2006). Other models have focused on the additional challenge that arises when foragers use information on patch quality to learn about large-scale changes in habitat conditions (McNamara and Houston 1985, 1987; Mangel 1990; Rodríguez-Gironés and Vásquez 1997; Hirvonen et al. 1999).

Acquiring information is valuable only when it leads to behavioral changes that enhance fitness; this is the utility of learning (Gould 1974; Stephens 1989; Dall et al. 2005). The information that is conveyed by resource encounters may depend on how the resources are distributed in space and time. This is illustrated by observations of the parasitoid specialist *Cotesia rubecula* and its generalist sibling species *Cotesia glomerata*, which both use a simple, non-responsive strategy when feeding on a uniformly distributed host (Vos et al. 1998). When the generalist parasitoid forages on a highly aggregated host species, it adjusts its behavior on the basis of recent host encounters because encounters convey valuable information on local resource conditions. Information collection and use, therefore, must be interpreted in an ecological context, by quantifying the fitness value this information represents to the forager (see Yoccoz et al. 1993).

In our model, the ecological costs of learning include direct costs, such as increased travel to sample more patches, but also opportunity costs: during the extra time the forager travels, it does not forage, and it may risk leaving a good patch prematurely. Sampling costs typically increase when the forager needs to travel far between patches or when patch quality varies so much that the forager needs to visit many patches to obtain a reliable estimate of the resource level in the habitat. Despite being simple mechanisms, such ecological costs of gathering information are seldom considered explicitly in foraging models (but see Eliassen et al. 2007; Stephens 2007). Most models have instead reflected a fixed investment in cognitive and sensory machinery by assuming that learning costs are independent of the ecological conditions or by

assuming a constant perception error (e.g., Richards and de Roos 2001; Collins et al. 2002).

The emphasis given to new experiences and the rate of forgetting older ones may represent adaptations to characteristics of the natural environment (McNamara and Houston 1987; Mangel 1990; Hirvonen et al. 1999; Schacter 1999). The ability to modify behavior on the basis of experience shows large genetic variability between populations and species (Dukas 2004). In a recent study on bumblebees, learning speed varied between colonies by a factor of nearly five, and this variation directly affected foraging performance (Raine and Chittka 2008). In *Drosophila*, differences in larval foraging mode and learning abilities are linked to a natural polymorphism in a single nucleotide (de Belle et al. 1989). “Rover” larvae are more mobile than the “sitter” type and have superior short-term learning performance but poorer long-term memory. These characteristics may be adaptations to variable and more stable environments, respectively (Mery et al. 2007). Evolutionary changes in memory characteristics and learning rates may also occur within a few generations in laboratory populations (Mery and Kawecki 2002).

The dynamic structure of the costs and benefits of learning and the presence of genetic variation in learning abilities beg the question, Under what ecological conditions should enhanced learning abilities evolve? The value of learning is generally assumed to be highest in variable environments, whereas the process of information acquisition relies on some persistence in order to be efficient (Eliassen et al. 2007; Stephens 2007). We present a framework in which the costs of learning emerge from active sampling behavior and fitness benefits may result if the forager can utilize the acquired information to alter its behavior. We compare a set of strategies that vary in the patch and habitat information that they use. Under different environmental conditions, we can thereby quantify the strength of selection in favor of learning and the fitness potential for improving patch or habitat assessment rules. We ask, Under what conditions would we predict that a forager should try to learn differences among patches or trends over time? How much emphasis should learners put on recent, as opposed to past, experience? In short, we try to map the ecological landscape for the evolution of learning in foraging behavior.

Model Description

We model a forager that searches for resources distributed in well-defined patches within the habitat. Time is divided into discrete time steps in which individuals either search for resources in a patch or move between foraging patches. By harvesting a patch, the forager reduces the local resource level, thus decreasing the probability of further re-

source encounters. Individuals do not interact, so resource depletion in a patch is unaffected by other foragers. Behavior determines the long-term intake rate, and we use the energy gain during a season of constant length as the fitness measure. With this model, we investigate the value of learning by studying patch-leaving strategies that differ in their use of information.

We focus on foragers that use resource encounters to estimate patch quality and integrate this information to track changes in mean resource conditions in the habitat. Following the logic of the marginal-value theorem, foragers should leave a patch when they have depleted it to a state where resource encounters would be higher elsewhere. Patch-leaving decisions are based on the time n it takes to accumulate one unit of resources. When n exceeds a patch-leaving threshold τ , the forager leaves the patch:

$$\begin{cases} \text{stay} & \text{if } n < \tau \\ \text{leave} & \text{if } n \geq \tau \end{cases}$$

When there is one random and independent search per time step, this is identical to a giving-up-time (GUT) rule (e.g., Iwasa et al. 1981; Green 1984). When food items are large, the time it takes to accumulate a given amount of resources will vary, and n will be highly stochastic. A simple GUT rule may therefore arrive at suboptimal patch-leaving decisions in highly stochastic environments (Oaten 1977; Iwasa et al. 1981; Green 1984). When food items become infinitely small, food intake approaches a constant flow, with little or no variance in n . Between these two extremes is a continuum of foraging rates from highly stochastic to deterministic and rate based. Stochasticity thus affects the information quality of n and thereby influences a forager's ability to assess the current resource level in a patch (fig. 1).

Behaviors such as patch residence times and travel frequencies emerge from the patch-leaving strategies. There are no metabolic or predation costs (for effects of predation on learning, see Eliassen et al. 2007). For every parameter combination, results based on the simulated behavior of 2,000 individuals over a foraging season of 4,000 time steps are shown. We study strategies that vary in the way they use information and differ in their access to free information. By comparing the performance of these strategies, we investigate the emergent parameters of the learning rule and quantify the potential value of information and learning in different ecological settings.

The Foraging Habitat

We study foraging in different ecological settings by letting patches vary along three ecological dimensions.

1. *Temporal change.* In many natural systems, resource




Figure 1: The patch assessment problem: resource level as a function of a forager's search time in a patch (solid line). Habitats differ in encounter stochasticity: $\sigma = 1.0$ for A, 0.5 for B, and 0.1 for C. The dashed line shows the resource level r^* at which a prescient forager, having complete and free patch information, leaves the patch. The dotted line shows the optimal search time t^* after which a constant-time forager leaves (representing evolved static behavior). An assessment forager will leave the patch whenever the time to accumulate one resource unit (indicated by the alternating gray and black line along the X-axis) exceeds the patch-leaving threshold τ .

conditions change in a consistent manner as seasons progress, for instance, because of prey growth, temperature, or precipitation. We assume that the average resource level $\bar{R}(t)$ changes continuously throughout the foraging period by an increment g every time step t :

$$\bar{R}(t) = R_0 + gt,$$

where R_0 is the mean resource level at the start of a season. The linear trend lasts throughout the whole season, and we vary the rate and direction of the temporal trend.

2. *Spatial heterogeneity.* Patches in the habitat vary in resource level according to a normal distribution. With a constant variance-to-mean relationship ϵ , the resource level in an unexploited patch i at a time t is

$$R_i(t) = \bar{R}(t) + Z_i \sqrt{\epsilon \bar{R}(t)},$$

where Z_i is a random variable with standard normal distribution (picked only once and characteristic of each patch i).

3. *Stochasticity in resource encounters.* A forager's prey intake may grade from almost constant flow (e.g., an aphid feeding on nectar) to highly stochastic resource encounters (such as the fox searching for reindeer carcasses). We consider different resource environments along this continuum by letting the resource be distributed from few and large items to many and small. The total amount of energy in the environment remains the same, and in a given habitat all items are of equal value q . We assume handling time to be proportional to energy content q and that a forager can handle a maximum of $N = 1/q$ food items per time step. Each time step is therefore divided into N independent searches, and the maximum energy that can be obtained in a time step is 1.

The probability that the forager will encounter a resource in each of its N searches depends on the current resource level $r_i(t)$ and the search efficiency a (Mangel 2006):

$$\begin{aligned} p(r_i(t)) &= \Pr(\text{encounter}; r_i(t)) \\ &= 1 - e^{-ar_i(t)/Nq} = 1 - e^{-ar_i(t)}. \end{aligned}$$

The resource level in patch i at time t is $r_i(t) = R_i(t) - k$ after a forager has encountered k units of resource. To retain the numerical approximation of the discrete-time model, the encounter probability is the same for all N searches within a time step. The time to accumulate a certain amount of resources will be more variable when foragers search for larger items (q close to 1; fig. 1). This is because the number of resources x encountered during a time step follows a binomial distribution, with N trials and probability of success $p(r_i(t))$. The variance in the number of encounters is hence $Np(r_i(t))[1 - p(r_i(t))]$. For a given resource level $r_i(t)$, the variance in resource intake I will be proportional to q :

$$\begin{aligned} \text{Var}(I) &= q^2 \left(\frac{1}{q} \right) p(r_i(t)) [1 - p(r_i(t))] \\ &= qp(r_i(t)) [1 - p(r_i(t))]. \end{aligned}$$

Hereafter, we denote encounter stochasticity $\sigma = q^{1/2}$.

Obtaining Patch Information

Foragers that adjust their patch residence time to the resource level in the patch need to update their information. We consider two types of patch information.

1. *From encounters.* Assessment foragers use the time n it takes to accumulate one resource unit (qN) to estimate patch quality. Since foragers randomly search for food, n will, on average, decline when there are more resources in a patch. In environments with large food items of one resource unit ($q = 1$), n equals the time between resource encounters. For $q < 1$, the search time n is summed over N successful encounters (fig. 1).

2. *Known.* Prescient foragers (sensu Valone 1991) have full and freely available information about the current resource level.

Obtaining Habitat Information

To make a decision whether to leave a patch, a forager compares its patch information with a leaving threshold. We consider three ways in which foragers arrive at this threshold.

1. *Genetically fixed.* Foragers may use a genetically inherited patch-leaving threshold that remains constant throughout the whole foraging season. For every parameter combination, we calculated the fixed patch-leaving threshold τ_T that resulted in the best average performance, using stochastic dynamic programming (Houston and McNamara 1999; Clark and Mangel 2000; details in the appendix in the online edition of the *American Naturalist*).

2. *Learned.* Learners may estimate changes in resource conditions by accumulating patch quality information. A linear-operator rule (McNamara and Houston 1987; Mangel 1990) determines how past and present information are weighted relative to each another. The new estimate (τ_{p+1}) is a weighted average of the previous estimate (τ_p) and the average time between encounters in the recent patch, including travel time v :

$$\tau_{p+1} = (1 - \gamma)\tau_p + \gamma \frac{1}{k} \left(v + \sum_{j=1}^k n_j \right).$$

The learning factor γ determines the relative weight given to new experiences, k is the total resource intake in the

patch, and n_j is the search time to accumulate each resource unit. Information gathering and foraging overlap, but learners often emphasize patch sampling early in the season and apply this information later to enhance foraging performance (Eliassen et al. 2007).

The initial patch-leaving threshold τ_0 and the learning factor γ are assumed to be genetically inherited traits adapted to the prevailing environmental conditions. We determined the parameter combination of τ_0 and γ that maximized lifetime resource intake by numerical simulations. Note that with no learning ($\gamma = 0$), the strategy is identical to a fixed-threshold rule with $\tau_T = \tau_0$.

3. *Known.* Farsighted foragers do not estimate the resource level in the habitat but receive free information about the best patch-leaving threshold $\tau_F(t)$ when they enter a new patch at time t . We call them farsighted because they have free long-range information about their environment. We find the optimal $\tau_F(t)$, which maximizes the forager's expected future resource intake, by using stochastic dynamic programming (details in the appendix).

The Behavioral Strategies

We quantify the value of learning as the difference in foraging performance between (1) a fixed-threshold forager that compares the required time to encounter one resource unit with an evolved, static patch-leaving threshold and (2) learners that perform the same local patch assessment but learn from past experiences and estimate habitat conditions (see Stephens 1989). Foragers may potentially improve the estimation method for both patch quality and habitat averages, and we assess the scope for improvement by comparing the foraging performance of learners with that of foragers that have access to free information (table 1 and, in the online edition of the *American Naturalist*, table A1). As an upper limit, the omniscient strategy, which has free information on both patch and habitat conditions, represents the best possible behavior in a given environment. Strategies with other combinations of patch and habitat information are given in table 1. Computational details on how to calculate the different strategies are presented in the appendix.

Results

The costs and benefits of information acquisition are not imposed in the model but emerge from ecological processes. We first evaluate how different patch-leaving strategies perform in environments that vary in levels of encounter stochasticity and spatial heterogeneity. Of special interest is the patch assessment strategy that can acquire information on resource conditions through regular foraging behavior. When average resource conditions vary over the foraging season, patch information can be used to estimate changing habitat conditions. This learning challenge is our next focus. Finally, we explore when learning strategies may evolve along the general ecological dimensions of spatial variability, encounter stochasticity, and temporal trends.

Patch Information

For high encounter stochasticity, the time it takes to accumulate a given amount of resources may deviate considerably from expectations (fig. 1A). This reduces the quality of patch information that foragers may obtain. When $\sigma = 1$, the assessment rule equals a giving-up-time (GUT) rule and has relatively poor performance (fig. 2). In contrast, when stochasticity is low ($\sigma \ll 1$), patch assessment approaches a rate-based strategy (fig. 1C). Under these conditions, a forager's performance is close to that of a prescient forager (fig. 2), which knows the exact resource level in the patch. Between these two extremes, the performance of the patch assessment rule decreases toward higher encounter stochasticity (fig. 2). It might seem surprising that a strategy that uses information (assessment foragers) performs worse than an ignorant strategy (constant-time foragers). Flexible strategies that use information may, however, make mistakes. Under high stochasticity, the efficacy of assessment is low and mistakes are frequent (fig. 3A).

The value of patch assessment not only is influenced by the risk of making mistakes (fig. 3A) but also depends on the utility of the acquired information. The more variation there is among patches, the more important is it to discriminate between good and poor patches (fig. 3B). As a consequence, the assessment strategy is more profitable in heterogeneous environments (high spatial variability)

Table 1: Foraging strategies with different levels of patch and habitat information

Patch information	Habitat information		
	Genetically fixed	Learned	Known
Ignored	Constant time (C)		
From encounters	Fixed threshold (T)	Learner (L)	Farsighted (F)
Known	Fixed prescient (PT)	Prescient learner (PL)	Omniscient (O)




Figure 2: Effect of encounter stochasticity: foraging performance of fixed-prescient (dashed line), fixed-threshold (solid line), and constant-time-rule (dotted line) foragers in a temporally stable habitat ($\bar{R} = 30$, $\epsilon = 2$). All strategies use optimal parameter values, so the difference between them is due to information quality and information use. The value of patch assessment is indicated by the dark gray area, while the light gray area shows the potential benefit of improving the patch assessment rule.

and when resources are encountered at a regular rate (low encounter stochasticity; fig. 3C).

Habitat Information

We now turn to the second property of the patch-leaving decision: the forager’s perception of conditions in the hab-

itat. Because the model environment is constant between generations, genetically fixed strategies may represent perfect adaptations. We therefore allow the resource levels within the habitat to show a seasonal trend. It is here that learning from past patch experiences may play a role. Our focus is foragers that use the assessment strategy to estimate the resource level in a patch and then compare this estimate with different sources of habitat information. The best such strategy is that of the farsighted forager, which has free and exact information about the average habitat quality. As patches become better over time, this forager will reduce its patch-leaving threshold because it then pays to stay only for the initial harvest period and then leave for an unexploited patch (fig. 4). In contrast, learning foragers estimate these long-term changes from foraging experiences and update their strategy accordingly; their patch-leaving thresholds are similar to those of farsighted foragers but deviate notably early in the season, when patch-sampling effort is most pronounced (Eliassen et al. 2007). Learners thus invest travel time (lost opportunity) when patches are less good to optimize behavior later in the season, when foraging rates are higher (improved utility). This also explains why learning is not adaptive in this model when resources decline over time: early patch sampling then is costly, and the linear-operator rule is outperformed by fixed-threshold strategies.

With more stochasticity (σ), each resource encounter provides less-accurate patch information, which necessitates repeated sampling. It becomes advantageous to stay

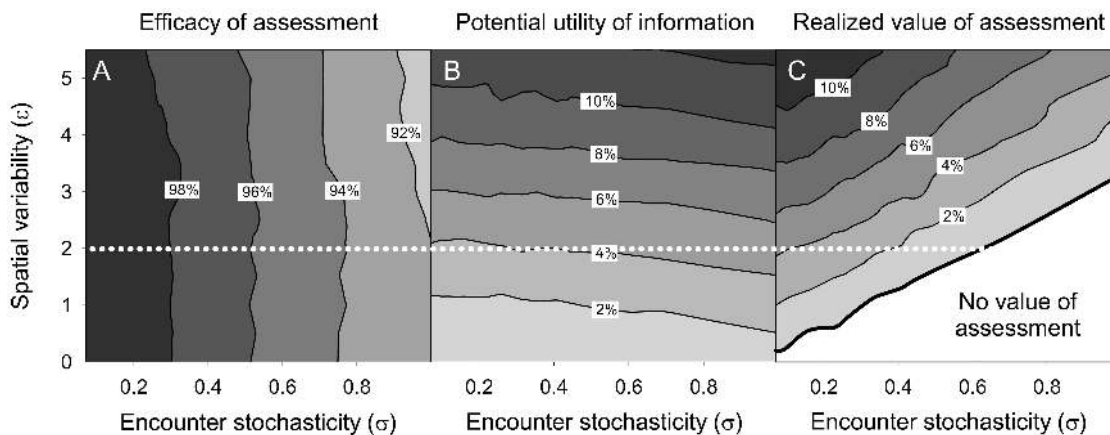


Figure 3: Role of patch assessment in foraging efficiency. The value of obtaining patch information in a temporally stable environment is a function of encounter stochasticity σ and spatial heterogeneity ϵ ($\bar{R} = 30$). In A, the efficacy of assessment is given as the performance of fixed-threshold foragers relative to that of fixed-prescient foragers with full patch information. In B, the potential utility of full patch information is the difference between fixed-prescient and constant-time foragers. In C, the realized value of patch assessment is the performance of fixed-threshold foragers relative to that of constant-time foragers. Patch assessment could potentially evolve for combinations of encounter stochasticity and spatial variability above the zero isocline (C, thick line). The white dotted line marks the spatial variability used for the results presented in figure 2. The contour plots are based on average performance of 2,000 individuals tested for each of 182 parameter combinations of ϵ and σ . For details on strategies and concepts, see table 1 and, in the online edition of the *American Naturalist*, table A1.




Figure 4: The habitat information problem: patch-leaving thresholds of learning (solid line), fixed-threshold (dotted line), and farsighted foragers (dashed line) as functions of time in the foraging season. The three panels represent habitats with different resource encounter stochasticity σ . The average resource level in the habitat increases with time ($g = 0.01$, $\epsilon = 0$, $R_0 = 10$, $a = 0.01$, $v = 25$, $T = 4,000$). On the contour plot in the background, the intensity of the gray shading increases with long-term intake rate of foragers utilizing that patch-leaving threshold at that time of the season (fitness landscape).

a bit longer (fig. 4), as opposed to leaving a potentially good patch too early. Low encounter stochasticity potentially yields higher intake rates, but in this environment intake rates drop more quickly if the patch-leaving strategy deviates from the optimal threshold (fig. 4).

*Integrating Patch and Habitat Information:
The Role of Learning*

The value of learning depends on (1) the accuracy of patch assessment, (2) the opportunity cost related to information sampling, (3) the proficiency of the learning rule, and (4) the utility of updated information. The importance of these elements can be interpreted from comparisons with a nonlearning strategy that assesses patch quality but uses a fixed patch-leaving threshold and with the omniscient strategy that represents the optimal behavior unconstrained by information.

1. When the accuracy of patch assessment is high (low encounter stochasticity), the performance of the learning strategy will approach that of omniscient foragers (fig. 5). With increasing encounter stochasticity the value of learning decreases, mainly because of poor patch assessment, at least when spatial variability is low (figs. 5, 6).

2. The opportunity cost increases with spatial variability because more patches must be sampled to obtain a reliable estimate of the average habitat conditions (fig. 6A). There is hence a trade-off between the quality of the learned habitat information and the time spent traveling to collect it (Eliassen et al. 2007). In much the same way, poor patch estimates select for longer memory records, so that experiences are averaged over more patches (fig. 6B). The opportunity cost is also influenced by the temporal change

in resource conditions. It will be low when sampling can be done during low resource availability early in the season to improve foraging rates later, when foraging conditions are better (fig. 6C).

3. The proficiency of the learning rule is high when the integration of patch experiences into a representative estimate of the habitat quality is straightforward. That is the

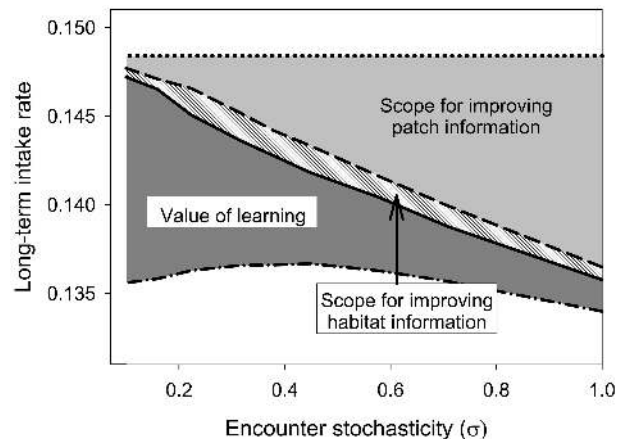


Figure 5: Value of learning and scope for improvement: long-term intake rates of omniscient (dotted line), farsighted (dashed line), learning (solid line), and fixed-threshold (dash-dotted line) foraging strategies as functions of encounter stochasticity. The value of learning (dark gray area) and the potential value of improving patch (light gray area) and habitat (hatched area) information quality are found as the difference in resource intake between foragers utilizing various levels of environmental information (table A1 in the online edition of the *American Naturalist*). Resource conditions change temporally through the foraging season ($g = 0.01$, $R_0 = 10$), but there is no spatial variance among patches ($\epsilon = 0$).

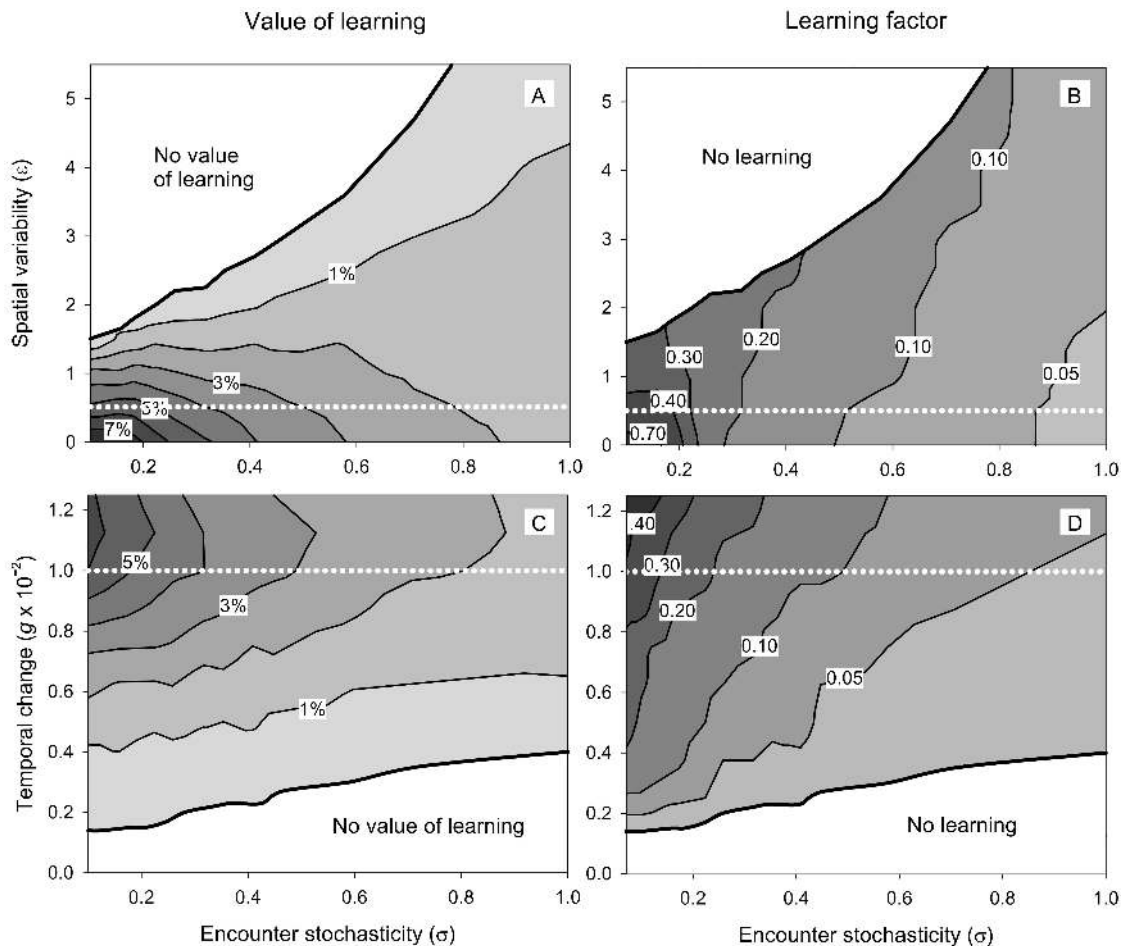


Figure 6: An adaptive landscape for the evolution of learning. The value of learning (A, C) and the weight given to recent information, that is, the learning factor γ (B, D), as functions of three ecological variables: (1) the level of encounter stochasticity σ (X-axis), (2) the spatial variance ϵ among resource patches in the habitat (Y-axis in A, B), and (3) the level of temporal change g through the foraging season (Y-axis in C, D). Learning has no positive effect on fitness for combinations of the environmental parameters in the white areas, and foragers adopt an innate fixed-threshold strategy. In A and C, the labels indicate the relative difference in long-term intake rate between learners and fixed-threshold foragers; the value of learning increases with the intensity of the gray shading. In B and D, adaptive-learning factors are given, with dark gray corresponding to high learning factors. Each contour plot is based on the average performance of 2,000 individuals per strategy tested for 182 parameter combinations ($v = 25$, $a = 0.01$, $g = 0.01$ in A, B; $\epsilon = 0.5$ in C, D). Along the white dotted line, the spatial and temporal variances are the same.

case when encounter stochasticity is low (each patch assessment is close to the true value for that patch) and spatial variability is low (each patch estimate is close to the average habitat value). Under these conditions, the learning factor is high and foragers quickly adjust their habitat estimates (fig. 6B). The proficiency of the learning rule is also high when there is little temporal change (the habitat estimate requires slower updating). However, when there is no temporal change, a fixed strategy can do better, and there is no value of learning (fig. 6D).

4. The utility of updated habitat information generally increases with the level of temporal change (fig. 6C). The utility of information can be interpreted as the gradient

in intake rates that updated information will allow a forager to climb (compare the fixed threshold with the other two strategies in fig. 4). Possessing accurate information is most valuable when the optimal patch-leaving strategy is constantly changing, and deviating from the best strategy increases as the temporal change becomes stronger and more important to keep track of.

Combining these four elements, we arrive at several predictions.

1. The integration from patch to habitat information (the learning process itself) is most proficient when the utility of information is low, which means that the habitat

information can readily be replaced by a fixed-threshold strategy. That the learning process is efficient is therefore not a good predictor of when the learning process is important or likely to evolve.

2. Learning factors generally increase when the temporal change is more rapid, because only the most recent estimates will reflect the current environmental conditions (fig. 6D). Spatial variability and encounter stochasticity will, on the other hand, introduce noise into the learning process and generally favor lower learning factors to average over more experiences (fig. 6B).

3. The major difficulty in integrating from patch to habitat information depends on the quality of patch information, as can be seen from the areas for potential improvement in figure 5. The performance of the patch assessment rule grades from inaccurate in highly stochastic environments to relatively accurate in environments with low stochasticity. When patch assessment is difficult, foragers need to sample more patches (increased opportunity costs) and possibly reduce their rate of information updating (reduced proficiency). The utility of updated habitat estimates also decreases, as illustrated by the performance of the farsighted foragers (fig. 5), since foragers are unable to take advantage of accurate habitat information when encounter stochasticity is high. As a corollary, one would expect to find a range of adaptations that improves the mechanisms by which foragers estimate the quality of a patch, whereas there is less scope for improvement in the integration from patch to habitat information (see "Discussion").

4. Learning is most valuable when it is optimal to alter the patch-leaving strategy because of temporal changes in resource conditions (fig. 6C). In environments with more spatial variability, foragers obtain less accurate estimates and consequently make more suboptimal patch-leaving decisions. The costs of departing from the optimal patch-leaving threshold is especially high in the least stochastic environments (fig. 4). As spatial variance increases, foragers consequently adopt a fixed strategy in these environments (fig. 6B). This illustrates that it is not the quality of information per se that determines the value of utilizing it but rather the fitness consequence of responding to that information.

Discussion

The learning strategy must first assess the resource level in the patch and then combine several patch assessments into one expectation value for the habitat as a whole. There are two advantages of modeling these steps mechanistically. First, the indirect costs of learning, such as more travel and lost foraging time, emerge from individual behavioral decisions, depending on environmental conditions. This

contrasts with earlier models, which have commonly imposed fixed information costs or used nonadaptive learning rules to study how forager distribution or performance may be affected by limited knowledge (e.g., Green 1984; Bernstein et al. 1988; Collins et al. 2002). Second, we have shown that learning does not necessarily evolve when it is easy, cheap, or fast; learning also must provide an ultimate fitness benefit that more than outweighs the costs. This extends earlier conceptual models that focused on how foragers most efficiently estimate environmental parameters but without considering the utility of that information or the optimal level of uncertainty, given sampling costs (McNamara and Houston 1985, 1987; Rodríguez-Gironés and Vásquez 1997; Hirvonen et al. 1999). By directly modeling the mechanisms, our approach can combine the efficiency and utility aspects of learning.

The indirect costs and the ultimate benefits of information acquisition change in complex ways with different types of environmental variability. Whether learning is adaptive depends on all three ecological dimensions we tested: stochasticity in resource encounters, spatial variability between patches, and temporal trends in the habitat as a whole. Our modeling framework also highlights what part of a foraging strategy has the greatest potential for improvement; in other words, theory can show us where we should particularly look for evolutionary adaptations.

The Value of Information

By optimizing foraging strategies that had some or all of the information for free, we could quantify the value of that information: How much did foraging improve if the forager had accurate patch or habitat information? In the case with no spatial variability, the most valuable information was the patch assessment (fig. 5). If patch information is better, it would also be easier to estimate the correct habitat value. In principle, there are two ways in which patch assessments can be improved. A forager may utilize information on search times and resource encounters in a more efficient way, for instance, by using Bayesian updating rules (Iwasa et al. 1981; McNamara 1982; Olsson and Holmgren 1998; Olsson and Brown 2006). Animals may also use other sensory cues to assemble more-accurate information from several sources (Persons and Uetz 1996; van Alphen et al. 2003). These cues can be environmental characteristics that are correlated with patch quality, such as habitat features or co-occurring species, or social cues based on the behavior of conspecific foragers (Valone 1989; van Alphen et al. 2003). Integrating information from different sources could be especially rewarding at high levels of encounter stochasticity.

The small difference in foraging performance between farsighted and learning foragers represents the fitness ben-

efit that can potentially drive evolution toward better learning strategies. Although the linear-operator rule is very simple, there was thus little scope for replacing it with a more sophisticated learning rule. This partially contrasts the recent findings of Groß et al. (2008), who showed that the performance of linear-operator rules may deteriorate over time in a multiarmed-bandit decision problem. In their model, the drop in performance was due to reduced frequency of sampling alternative options. We do not expect the same effect in our model, however, because foragers directly affect resource levels in their patches and will switch between locations as resources become depleted.

Learning Strategies and Memory

Studies in cognitive psychology illustrate how the rate of information loss may adaptively vary in response to the temporal characteristics of the problem an individual needs to solve (Anderson et al. 1997; Jones and Sieck 2003), and our modeling shows that different ecological conditions select for different learning rates. Rapid temporal change makes recent patch visits more informative and therefore favors short-term memory (see also McNamara and Houston 1987; Hirvonen et al. 1999). Spatial heterogeneity breaks down this temporal correlation and selects for longer memory records (see also McNamara and Houston 1987; Nishimura 1994). Learning incurs lower costs when foragers can easily access information of high quality but has higher value when simple, inflexible strategies perform poorly. As Stephens (2007) notes, learning is an adaptation to predictability as much as it is an adaptation to change.

In general, our results indicate that updating rules rapidly, which puts an emphasis on recent experiences, may be most prevalent in nature. We arrive at this conclusion because learning rules that are updated slowly perform best in relatively stable environments (see also Hirvonen et al. 1999; Fortin 2002), where the value of learning is generally low. In such cases, keeping a long-term memory may impose costs that outweigh the benefits of learning altogether (Mery and Kawecki 2005). Empirical studies show that patch time allocation is often influenced by recent experiences of travel time (e.g., Kacelnik and Todd 1992; Cuthill et al. 1994; Thiel and Hoffmeister 2004) and patch quality (e.g., Wildhaber et al. 1994; Schilman and Rocas 2003; Outreman et al. 2005; Thiel and Hoffmeister 2006). In natural environments, the distribution of resources may vary on multiple scales (Fauchald 1999); hence, heavy reliance on recent experiences may be an adaptation to an environment with strong temporal correlation between similar foraging events (Real 1991; Jones and Sieck 2003). When foragers adjust only slowly to tem-

poral change, behavioral alterations are more subtle and harder to observe (Cuthill et al. 1994; Thiel and Hoffmeister 2006). In such cases, it requires intensive study over longer time periods to achieve the statistical power needed to differentiate slowly updating behavioral trends from fixed strategies. This also contributes to explaining why slow learning is poorly represented in the literature.

Assessment and Errors

According to Weber's law, a forager will show larger variance in memorized searching times as time intervals increase (Shettleworth 1998). This would enhance the profitability of mechanisms relying on estimates over shorter time periods, because the forager can more accurately discriminate time differences (Stephens 2002). In the model, we assumed that foragers were able to assess search time perfectly. This is an obvious simplification that may affect the relative profitability of the different patch assessment mechanisms. In particular, the constant-time rule that relies on the total residence time in the patch would be more prone to errors than the patch-leaving rule, which instead uses time between resource encounters.

All strategies, even the constant-time rule, are adapted to some aspects of the resource conditions in the habitat. Our modeling environment has seasonal variation in resource levels but no temporal change or variability on longer timescales. In this way, the inflexible constant-time foragers and the nonlearning fixed-threshold strategy will be perfectly adapted to the average habitat conditions within a season, given enough time for adaptation to take place. Assuming that all innate strategies are perfectly adapted to an environment that does not vary between years places the cost of information on the evolutionary timescale. Experimental studies, however, indicate that innate preferences may be hard to obtain, in which case the advantage of learning may be more permanent (Mery and Kawecki 2004). Learning may also be profitable when resource conditions change between generations and may have a transient advantage under directional selection toward new optimum strategies (Stephens 1993; Papaj 1994). Together, these observations suggest that our estimated value of learning is conservative.

To summarize, our analysis illustrates that even though a forager is capable of learning, it may not be beneficial to explore all types of resource patterns or learn under all ecological conditions (see also Stephens 1991, 1993; Nishimura 1994). We have also shown the importance of including the mechanisms by which information is gathered, interpreted, and remembered. Understanding how these mechanisms operate and their adaptive value may give further insights into how and when to learn from ecologically emerging information.

Acknowledgments


We thank P. Lundberg and E. J. Milner-Gulland for constructive comments on the manuscript and the Research Council of Norway for funding for S.E. and C.J. and travel money for M.M. We also thank the Center for Stock Assessment Research for helping facilitate the visit of S.E. to the University of California, Santa Cruz.

Literature Cited

- Anderson, R. B., R. D. Tweney, M. Rivardo, and S. Duncan. 1997. Need probability affects retention: a direct demonstration. *Memory and Cognition* 25:867–872.
- Bernstein, C., A. Kacelnik, and J. R. Krebs. 1988. Individual decisions and the distribution of predators in a patchy environment. *Journal of Animal Ecology* 57:1007–1026.
- Charnov, E. L. 1976. Optimal foraging, the marginal value theorem. *Theoretical Population Biology* 9:129–136.
- Clark, C. W., and M. Mangel. 2000. *Dynamic state variable models in ecology*. Oxford University Press, Oxford.
- Collins, E. J., A. I. Houston, and A. Lang. 2002. The ideal free distribution: an analysis of the perceptual limit model. *Evolutionary Ecology Research* 4:471–493.
- Cuthill, I. C., P. Haccou, and A. Kacelnik. 1994. Starlings (*Sturnus vulgaris*) exploiting patches: response to long-term changes in travel time. *Behavioral Ecology* 5:81–90.
- Dall, S. R. X., L.-A. Giraldeau, O. Olsson, J. M. McNamara, and D. W. Stephens. 2005. Information and its use by animals in evolutionary ecology. *Trends in Ecology & Evolution* 20:187–193.
- de Belle, J. S., A. J. Hilliker, and M. B. Sokolowski. 1989. Genetic localization of *foraging (for)*: a major gene for larval behavior in *Drosophila melanogaster*. *Genetics* 123:157–163.
- Dukas, R. 2004. Evolutionary biology of animal cognition. *Annual Review of Ecology, Evolution, and Systematics* 35:347–374.
- Eliassen, S., C. Jørgensen, M. Mangel, and J. Giske. 2007. Exploration or exploitation: life expectancy changes the value of learning in foraging strategies. *Oikos* 116:513–523.
- Fauchald, P. 1999. Foraging in a hierarchical patch system. *American Naturalist* 153:603–613.
- Fortin, D. 2002. Optimal searching behaviour: the value of sampling information. *Ecological Modelling* 153:279–290.
- Fretwell, S. D., and H. L. J. Lucas. 1970. On territorial behaviour and other factors influencing habitat distribution in birds. *Acta Biotheoretica* 19:16–36.
- Gould, J. P. 1974. Risk, stochastic preference, and the value of information. *Journal of Economic Theory* 8:64–84.
- Green, R. F. 1984. Stopping rules for optimal foragers. *American Naturalist* 123:30–43.
- Groß, R., A. I. Houston, E. J. Collins, J. M. McNamara, F.-X. Dechaume-Moncharmont, and N. R. Franks. 2008. Simple learning rules to cope with changing environments. *Journal of the Royal Society Interface* 5:1193–1202.
- Hirvonen, H., E. Ranta, H. Rita, and N. Peuhkuri. 1999. Significance of memory properties in prey choice decisions. *Ecological Modelling* 115:177–189.
- Houston, A. I., and J. M. McNamara. 1999. *Models of adaptive behaviour: an approach based on state*. Cambridge University Press, Cambridge.
- Iwasa, Y., M. Higashi, and N. Yamamura. 1981. Prey distribution as a factor determining the choice of optimal foraging strategy. *American Naturalist* 117:710–723.
- Jones, M., and W. R. Sieck. 2003. Learning myopia: an adaptive recency effect in category learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 29:626–640.
- Kacelnik, A., and I. A. Todd. 1992. Psychological mechanisms and the marginal value theorem: effect of variability in travel time on patch exploitation. *Animal Behaviour* 43:313–322.
- Mangel, M. 1990. Dynamic information in uncertain and changing worlds. *Journal of Theoretical Biology* 146:317–332.
- . 2006. *The theoretical biologist's toolbox: quantitative methods for ecology and evolutionary biology*. Cambridge University Press, New York.
- McNamara, J. M. 1982. Optimal patch use in a stochastic environment. *Theoretical Population Biology* 21:269–288.
- McNamara, J. M., and A. I. Houston. 1985. Optimal foraging and learning. *Journal of Theoretical Biology* 117:231–249.
- . 1987. Memory and the efficient use of information. *Journal of Theoretical Biology* 125:385–395.
- Mery, F., and T. J. Kawecki. 2002. Experimental evolution of learning ability in fruit flies. *Proceedings of the National Academy of Sciences of the USA* 99:14274–14279.
- . 2004. The effect of learning on experimental evolution of resource preference in *Drosophila melanogaster*. *Evolution* 58:757–767.
- . 2005. A cost of long-term memory in *Drosophila*. *Science* 308:1148.
- Mery, F., A. T. Belay, K.-C. S. Anthony, M. B. Sokolowski, and T. J. Kawecki. 2007. Natural polymorphism affecting learning and memory in *Drosophila*. *Proceedings of the National Academy of Sciences of the USA* 104:13051–13055.
- Nishimura, K. 1994. Decision making of a sit-and-wait forager in an uncertain environment: learning and memory load. *American Naturalist* 143:656–676.
- Oaten, A. 1977. Optimal foraging in patches: a case for stochasticity. *Theoretical Population Biology* 12:263–285.
- Olsson, O., and J. S. Brown. 2006. The foraging benefits of information and the penalty of ignorance. *Oikos* 112:260–273.
- Olsson, O., and N. M. A. Holmgren. 1998. The survival-rate-maximizing policy for Bayesian foragers: wait for good news. *Behavioral Ecology* 9:345–353.
- Outreman, Y., A. Le Ralec, E. Wajnberg, and J. S. Pierre. 2005. Effects of within- and among-patch experiences on the patch-leaving decision rules in an insect parasitoid. *Behavioral Ecology and Sociobiology* 58:208–217.
- Papaj, D. R. 1994. Optimizing learning and its effect on evolutionary change in behavior. Pages 133–154 *in* L. A. Real, ed. *Behavioral mechanisms in evolutionary biology*. University of Chicago Press, Chicago.
- Persons, M. H., and G. W. Uetz. 1996. The influence of sensory information on patch residence time in wolf spiders (*Araneae: Lycosidae*). *Animal Behaviour* 51:1285–1293.
- Raine, N. E., and L. Chittka. 2008. The correlation of learning speed and natural foraging success in bumble-bees. *Proceedings of the Royal Society B: Biological Sciences* 275:803–808.
- Real, L. A. 1991. Animal choice behavior and the evolution of cognitive architecture. *Science* 253:980–986.
- Richards, S. A., and A. M. de Roos. 2001. When is habitat assessment an advantage when foraging? *Animal Behaviour* 61:1101–1112.
- Rodríguez-Gironés, M. A., and R. A. Vásquez. 1997. Density-depen-

- dent patch exploitation and acquisition of environmental information. *Theoretical Population Biology* 52:32–42.
- Schacter, D. L. 1999. The seven sins of memory: insights from psychology and cognitive neuroscience. *American Psychologist* 54: 182–203.
- Schilman, P. E., and F. Roces. 2003. Assessment of nectar flow rate and memory for patch quality in the ant *Camponotus rufipes*. *Animal Behaviour* 66:687–693.
- Shettleworth, S. J. 1998. *Cognition, evolution and behavior*. Oxford University Press, Oxford.
- Stephens, D. W. 1989. Variance and the value of information. *American Naturalist* 134:128–140.
- . 1991. Change, regularity, and value in the evolution of animal learning. *Behavioral Ecology* 2:77–89.
- . 1993. Learning and behavioral ecology: incomplete information and environmental predictability. Pages 195–218 in D. R. Papaj and A. C. Lewis, eds. *Insect learning: ecology and evolutionary perspectives*. Chapman & Hall, London.
- . 2002. Discrimination, discounting and impulsivity: a role for an informational constraint. *Philosophical Transactions of the Royal Society B: Biological Sciences* 357:1527–1537.
- . 2007. Models of information use. Pages 31–58 in D. W. Stephens, J. S. Brown, and R. C. Ydenberg, eds. *Foraging: behavior and ecology*. University of Chicago Press, Chicago.
- Thiel, A., and T. S. Hoffmeister. 2004. Knowing your habitat: linking patch-encounter rate and patch exploitation in parasitoids. *Behavioral Ecology* 15:419–425.
- . 2006. Selective information use in parasitoid wasps. *Animal Biology* 56:233–245.
- Valone, T. J. 1989. Group foraging, public information, and patch estimation. *Oikos* 56:357–363.
- . 1991. Bayesian and prescient assessment: foraging with pre-harvest information. *Animal Behaviour* 41:569–578.
- van Alphen, J. J. M., C. Bernstein, and G. Driessen. 2003. Information acquisition and time allocation in insect parasitoids. *Trends in Ecology & Evolution* 18:81–87.
- Vos, M., L. Hemerik, and L. E. M. Vet. 1998. Patch exploitation by the parasitoids *Cotesia rubecula* and *Cotesia glomerata* in multi-patch environments with different host distributions. *Journal of Animal Ecology* 67:774–783.
- Wildhaber, M. L., R. F. Green, and L. B. Crowder. 1994. Bluegills continuously update patch giving-up times based on foraging experience. *Animal Behaviour* 47:501–513.
- Yoccoz, N. G., S. Engen, and N. C. Stenseth. 1993. Optimal foraging: the importance of environmental stochasticity and accuracy in parameter estimation. *American Naturalist* 141:139–157.

Associate Editor and Editor: Michael C. Whitlock



Left, *Veretillum stimpsonii*; center, single sea pen; right, whole colony of sea pens *Kophobelemnon clavatum* V.; from the review of “Synopsis of the Polyps and Corals of the North Pacific Exploring Expedition, under Commodore C. Ringgold and Capt. John Rodgers, U.S.N., from 1853 to 1856, collected by Dr. Wm. Stimpson” (*American Naturalist*, 1870, 4:488–491).