VOL. 174, NO. 3 THE AMERICAN NATURALIST SEPTEMBER 2009

Bumblebees Learn to Forage like Bayesians

Jay M. Biernaskie,[*,†] **Steven C. Walker,**[†] **and Robert J. Gegear**[‡]

Department of Ecology and Evolutionary Biology, University of Toronto, 25 Willcocks Street, Toronto, Ontario M5S 3B2, Canada

Submitted March 10, 2009; Accepted April 15, 2009; Electronically published July 24, 2009

Online enhancements: appendixes.

ABSTRACT: Bayesian foraging in patchy environments requires that foragers have information about the distribution of resources among patches (prior information), either set by natural selection or learned from past experience. We test the hypothesis that bumblebee foragers can rapidly learn prior information from past experience in two very different experimental environments. In the high-variance environment (patches of low and high quality), stochastic optimality models predicted that finding rewards should sometimes sharply increase an optimal forager's tendency to stay in a patch (an incremental response), whereas in the uniform environment, finding rewards should always decrease the tendency to stay (a decremental response). We use Cox regression models to show that, in a matter of hours, bees learned to match both predicted responses, resulting in a reward intake rate that averaged 80% of the predicted maximum. Following training in either environment, bees' adaptive behavior carried over to a common test environment, thus confirming the influence of memorized prior information. Although Bayesian foraging by learning is often presumed, this study is the first to clearly isolate the adaptive use of a learned prior expectation. More generally, it highlights the remarkable adaptive plasticity of an important generalist pollinator and agent of selection.

Keywords: Bayesian updating, learning, marginal value theorem, optimal foraging, patch-use model.

Introduction

The history of foraging theory related to patch use shifts from considering when a forager should optimally leave a patch of resources (e.g., Charnov 1976) to considering how the forager should decide when to leave (reviewed by McNamara et al. [2006]). This shift reflects an input of ecological realism to the patch-use problem; in particular, because food items are often encountered stochastically and patches vary in quality, a forager must use incomplete information to decide when a patch should be abandoned. The optimal decision rule, which tells a forager how to infer future patch quality on the basis of its current experience in a patch, can be calculated with Bayes's theorem by an investigator who knows the distribution of food items among patches in the environment (the prior distribution; e.g., Oaten 1977; McNamara 1982; Green 2006). In order to approximate an optimal foraging rule, animals must similarly have some knowledge of the prior distribution. They might learn prior information from their past experience; however, if the distribution of food items is stable over time, that information could instead be hardwired by natural selection, resulting in animals that behave as if they know the prior distribution for their environment (McNamara et al. 2006; Pierre and Green 2008).

Examples of patch-leaving behavior from both vertebrates (reviewed by Valone [2006]) and invertebrates (see Pierre and Green 2008) can be interpreted as Bayesian-like, yet the extent to which prior knowledge of resource distributions is learned or hardwired remains unclear. Even the most impressive examples of Bayesian foraging by birds do not necessarily imply a special learning ability. In many cases, birds encountered an unmodified distribution of food items in nature (e.g., Alonso et al. 1995) or an experimental distribution that might reflect the type of ancestral environment to which the animals are already adapted (e.g., the negative-binomial distribution studied by Van Gils et al. [2003] or the empty/full distributions used by Lima [1984]; see Ydenberg 1998). Furthermore, the birds used in many experimental studies were trained to a particular food distribution for weeks. If learning did take place during this time, it is not clear how long it took for prior knowledge to be replaced by new information.

We hypothesized that an invertebrate forager, the bumblebee, should have the capacity for Bayesian-like foraging and the ability to rapidly learn prior information from very different environment types. As generalists, bumblebees have an evolutionary history of encountering a diversity of nectar resources that can vary greatly among plants of the same species and among flowers within a single plant (Goulson 2003, p. 73 and references therein). Furthermore, the extent of variation among plants will

---

[*] Corresponding author; e-mail: jay.biernaskie@utoronto.ca.

[†] We declare equal contributions from J.M.B. and S.C.W.

[‡] Current address: Department of Neurobiology, University of Massachusetts Medical School, Worcester, Massachusetts 01605.

Am. Nat. 2009. Vol. 174, pp. 413–423.
0003-0147/2009/17403-51136$15.00.
DOI: 10.1086/603629

differ among species that bees visit in a lifetime and may also fluctuate within a population as a result of daily patterns of reward depletion and renewal (e.g., Pleasants and Zimmerman 1979; Thomson 1988). It is exactly in this type of environment that natural selection would favor the ability to learn about the distribution of resources among patches.

When patches are consistently rewarding and of the same quality, bumblebees retain prior information about the mean quality of the environment and adjust their patch-leaving rules accordingly (Biernaskie and Gegear 2007). In contrast, the leaving rules that bees use on plants with variable rewards are often interpreted as fixed rules of thumb (e.g., a threshold nectar volume or number of empty flowers; Pyke 1978; Hodges 1985) or as rules that otherwise depend only on recent success in a patch (e.g., Kadmon and Shmida 1992). Yet given the variance in rewards that can be found within and among patches, bees that follow a fixed departure rule would often leave good patches too soon or stay too long in poor patches. They could avoid such errors by incorporating information about the distribution of rewards in their past experience.

Our study rigorously tests whether bumblebees can learn to adapt to novel environments by adjusting their patch-leaving rules in a Bayesian-like manner. The experiments involve training naive foragers (for hours, rather than weeks) in either a uniform environment (all patches having the same number of rewards) or a high-variance (HV) environment (patches of either low or high quality) and then observing each forager in a common test environment. This design is unique in that it isolates the effect of memorized prior information on foraging behavior in the common environment and thus explicitly tests the central feature of Bayesian foraging via learning (McNamara et al. 2006; Valone 2006). We also take advantage of recent theoretical developments (Olsson and Holmgren 1998; Green 2006) to calculate the optimal patch-leaving rule in each training environment. Bumblebees are certainly not perfect Bayesian decision makers, but optimality models allow us to compare their actual reward intake rates to the expected intake rate of an optimal forager. As a compliment to the optimal foraging models, we use a statistical model (extensions of Cox's proportional hazards model; Therneau and Grambsch 2000) to estimate bees' tendency to remain in a patch as a function of information gained by sampling the patch. In this way, the deterministic optimality predictions can be compared with a more realistic, proximate model of animal behavior (Wajnberg 2006; Pierre and Green 2008). We show a strong qualitative match between the two models, which, together with evidence from the test environment, strongly suggests the use of learned prior information in an adaptive, Bayesian-like way.

Material and Methods

Experimental Methods

Worker bumblebees (*Bombus impatiens* Cresson; supplied as eggs and larvae by Biobest Biological Systems, Leamington, ON) were trained to collect a 30% sucrose solution (herein "nectar") from artificial flowers within a large screened enclosure (3.4 m × 2.0 m × 4.6 m). Bees entered the enclosure through a gated tunnel connected to their nest box. The flowers were clear 1.5-mL polypropylene microcentrifuge tubes with the cap removed and a circular collar (6.5-cm diameter) of blue cardboard fixed around the tube's entrance (Gegear and Laverty 2005). To train for initial flower recognition, four to five filled flowers were presented at a single location within the enclosure. Worker bees from the colony eventually found and foraged on these flowers to sustain the colony when experiments were not in progress. Active foragers were painted with a small unique marking for identification.

For all training and tests, 12 flowers were grouped into rectangular patches (4 × 3) spaced 7.5 cm apart on pieces of rigid poster board. The environment consisted of 10 patches, arranged in two rows of five and spaced 1 m apart. In total, 50 flowers in the environment each contained 5 μL of nectar, and the remaining 70 flowers contained 5 μL of water (details of the distribution of rewards are given below). During training and tests, only a single marked forager was allowed to make repeated foraging bouts to the environment, and during foraging, we controlled bees' movements by covering all but two patches: the one being harvested and the (unvisited) adjacent patch. This ensured that bees' interpatch travel distance was consistent, that patches were never revisited, and that bees experienced the entire distribution of patch types that we intended. Individuals were trained on a particular distribution of rewards for at least 45 patch visits (usually 2–3 h) consisting of multiple foraging bouts (lengths of bouts ranged from four to 10 patch visits). A bout ended when the bee returned to the colony to unload its harvest, and during this interruption, we refilled all rewarding flowers that had been visited in the environment (after checking that visited flowers were in fact emptied). Hence, a fully replenished environment was always available at the start of each new foraging bout. We felt that 45 patch visits would be sufficient to test whether bees updated their pretraining prior distribution, presumably set mainly by evolutionary history, with new information about the distribution of rewards in the following training environments.

Uniform treatment. Each patch contained five rewarding flowers in a random spatial arrangement. After every second foraging bout, the position of patches in the environment was haphazardly shuffled, and after every 15 patch visits (or whenever the bout ended thereafter), a new set of patches with a new random arrangement of rewards was presented. We did this to minimize bees' ability to memorize the location of rewarding flowers but also to minimize the evaporation of rewards and the influence of floral scent marks.

High-variance treatment. Half of the patches contained a single reward, and half contained nine rewards (all randomly distributed). The sequence of patch types in the environment was randomized. Patches were shuffled after every second bout, but in this case, patches were moved to a new randomized sequence. Also as above, a new set of patches was presented after every 15 patch visits.

Once the threshold of 45 training visits was exceeded (excluding the final patch of each bout), we replaced all patches with the test environment, and the test began on the next foraging bout. All patches in the test environment contained five rewarding flowers in a common spatial arrangement. Bees from both training environments experienced the common arrangement of rewards and the same sequence of patches within the test environment. The test lasted for only 10 patch visits so that the observed behavior in the test could reflect any learned information from prior training; in this way, we hoped to minimize the amount of new information that might be learned within the test itself (this is particularly relevant for HV-trained [HV-T] bees, given that the HV environment was very different from the test environment).

Individual bees went through only one sequence of training (uniform or HV) and test. In total, 20 bees (10 of each training type) from three different colonies were observed. We chose newly emerged workers that were active foragers at the time that trials were to begin. Rather than assigning treatments to bees in a random order, we alternated training types in order to distribute them equally among colonies. The behavior of each bee during its last 15 patch visits of training, herein "late training," and all 10 patch visits in the test (as always, excluding the last patch visit of a bout) were videotaped. For the last three bees assigned to each treatment, we also recorded 15 patch visits of early training, which began after a bee's initial 10 patch visits (ensuring that it had learned to efficiently handle flowers and travel between patches). One observer used JWatcher, version 6.1 (Blumstein and Daniel 2007), to measure the number of rewards found, interreward search times, giving-up times (time from obtaining the last reward or entering the patch, whichever was more recent, to leaving), and travel times among patches.

Setup and Parameterization of the Optimal Foraging Model

The model environment corresponds to our experimental training environments but consists of an infinite series of patches. Before a patch is visited, y of 12 flowers contain a nectar reward, and the distribution of y among patches is the model analogue of our experimental manipulation. Each patch visited in the HV environment has probability 1/2 of being either a low-quality ($y = 1$) or a high-quality ($y = 9$) patch. Each patch in the uniform environment initially has five rewards ($y = 5$) with probability 1. These are the prior distributions that enter into the optimal foraging model below. Bees are assumed to know the appropriate distribution with certainty by late training.

A foraging strategy divides bees' time into searching for rewards on a patch or traveling among patches. Real bees also had a nonzero handling time associated with nectar consumption, but handling time has no effect on the choice of an optimal strategy (Oaten 1977). Interpatch travel time is a random variable with mean τ. Interreward search times on patches are assumed to be exponentially distributed random variables with rate parameter $a(y - x)$, where a is a constant called the searching efficiency (sensu Murdoch and Oaten 1975) and x is the number of rewards already found in the patch. Under this exponential distribution, $1/a(y - x)$ is the mean interreward time.

Our foraging model thus requires an estimate of three parameters: the average travel time between patches, τ, and the average searching efficiency in HV and uniform environments. Average travel time, τ, was estimated as the average measured time for bees to travel between patches. Searching-efficiency parameters were estimated by a survival analysis that assumed that interreward times were exponentially distributed with rate $a(y - x)$. We graphically checked that interreward times were indeed approximately exponentially distributed (app. A in the online edition of the *American Naturalist*). For each training environment, a was estimated by maximum likelihood. Estimation of a requires that all giving-up times be treated as censored interreward times (the true time is at least as long, but we cannot know how much longer). We handled these censored observations following Haccou and Meelis (1992, box 4.16).

Optimal Foraging: The Potential Value Model

The optimal foraging model is a slightly modified version of Green's (2006) continuous-time, random-search model. Following other authors, we call it the potential value model, which finds a rule that is known to maximize energy intake from any particular resource distribution (e.g., Olsson and Holmgren 2000). Here, we give an intuitive account of the model and reserve mathematical details for appendix B in the online edition of the *American Naturalist.*

Potential value measures patch quality as the expected number of rewards that a forager will obtain from the remainder of the current patch visit divided by the expected amount of time spent searching the remainder of the patch. Notice that potential value depends on how the forager decides to leave in the future, a unique problem that requires working backward through time with dynamic programming (Mangel and Clark 1989). The model assumes that foragers leave a patch when potential value falls below a threshold, C. Each choice of C leads to a long-term intake rate,

$$R = \frac{\mathrm{E}(G)}{\mathrm{E}(S) + \tau}, \tag{1}$$

where $\mathrm{E}(G)$ is the total number of rewards that a forager following rule C expects to find in a patch and $\mathrm{E}(S)$ is the total time that such a forager expects to spend searching a patch. It can be shown that the threshold rule that maximizes the reward intake rate in a particular environment, C_i^* ($i =$ HV or uniform), is equal to R_i^*, the maximum long-term intake rate for environment i (Green 2006). The optimal thresholds for our parameterized models are given in figure 1*A* and 1*B* (*dashed lines*).

The potential value rule optimally combines current patch information (the number of rewards found and the total time spent searching) with prior knowledge of the resource distribution. In this way, rewards found within a given amount of time can give an optimal forager information about the quality of the patch being exploited and, in some cases, information about the number of rewards remaining. With prior knowledge of the HV environment, an optimal forager is partially informed (uncertain) of patch quality only when finding fewer than two rewards (consistent with being on either a high-quality or a low-quality patch). If no rewards are found, potential value quickly declines to the leaving threshold (fig. 1*A*). Finding one reward, however, indicates the possibility of being in a high-quality patch, and finding a second reward indicates a high-quality patch with certainty; in both of these cases, potential value sharply increases upon finding a reward and so should bees' tendency to stay in the patch (an incremental response; Wagge 1979). After finding two rewards in the HV environment (becoming fully informed and knowing that seven rewards remain), potential value varies only with the number of rewards found (but not with time): each future reward suddenly reduces the potential value (tendency to stay; a decremental response; Driessen et al. 1995). Notice that uniform-trained (U-T) optimal foragers are always fully informed and should behave as if they know that all patches initially contain five rewards (fig. 1*B*). As above, each reward found indicates that the patch is closer to being empty.

Statistical Model of Behavior: Cox's Proportional Hazards

The potential value model is a useful tool for identifying optimal behavior but is less useful for describing the actual behavior of foraging animals. In a given context (number of rewards found and time spent searching), the potential value model predicts that bees with common prior information will behave in exactly the same manner. Hence, optimal behavior is inflexible because it is fully determined by optimal use of prior information. However, we expected (and indeed observed) that even a single bee in identical contexts would exhibit behavioral variation. A more flexible statistical model that makes no assumptions about optimality or the effect of prior information is required to infer bees' actual tendency to stay on a patch. For this purpose, we used a version of Cox's proportional hazards model (following, e.g., Haccou et al. 1991; Wajnberg et al. 1999). Our strategy was to assess the degree of similarity between the potential value model, which is optimized to experimental reward distributions, and the flexible Cox model, which we inferred from experimental data. If the fitted Cox model were to match the potential value model (as depicted in fig. 1), this would imply that bees used prior information to adaptively adjust their patch-leaving tendencies (i.e., that their behavior is Bayesian-like).

Our Cox model shares an important feature with the potential value model that allows the inference of Bayesian-like foraging: both models assume that bees' tendency to stay in the current patch is a function of the time spent searching the patch, t, and the number of rewards, $x(t)$, obtained by time t. Potential value is itself a measure of bees' tendency to stay in a patch; in the Cox model, the analogous measure is given by the reciprocal of the so-called hazard. In survival analyses of patch residence time (e.g., Wajnberg et al. 1999), the hazard function is the probability density that a foraging animal will leave its current patch in the next instant. The decision to leave a patch under the Cox model is thus a random event that becomes more or less likely depending on bees' experience in the patch (i.e., on t and $x(t)$). It follows that, unlike the potential value model, the Cox model does not predict a sharp leaving threshold.

The hazard function was modeled as

$$h(t, x(t)) = h_0(t)\exp(-\beta_{x(t)}), \tag{2}$$

where β_i ($i = 0, \ldots, 9$) is a coefficient measuring the effect of finding the ith reward on the hazard and $h_0(t)$ is the so-called baseline hazard function, which does not depend on the number of rewards obtained. The number of rewards obtained, $x(t)$, was treated as a time-varying categorical factor (see Wajnberg 2006). We set $\beta_0 = 0$ to treat zero rewards as the reference category, meaning that the

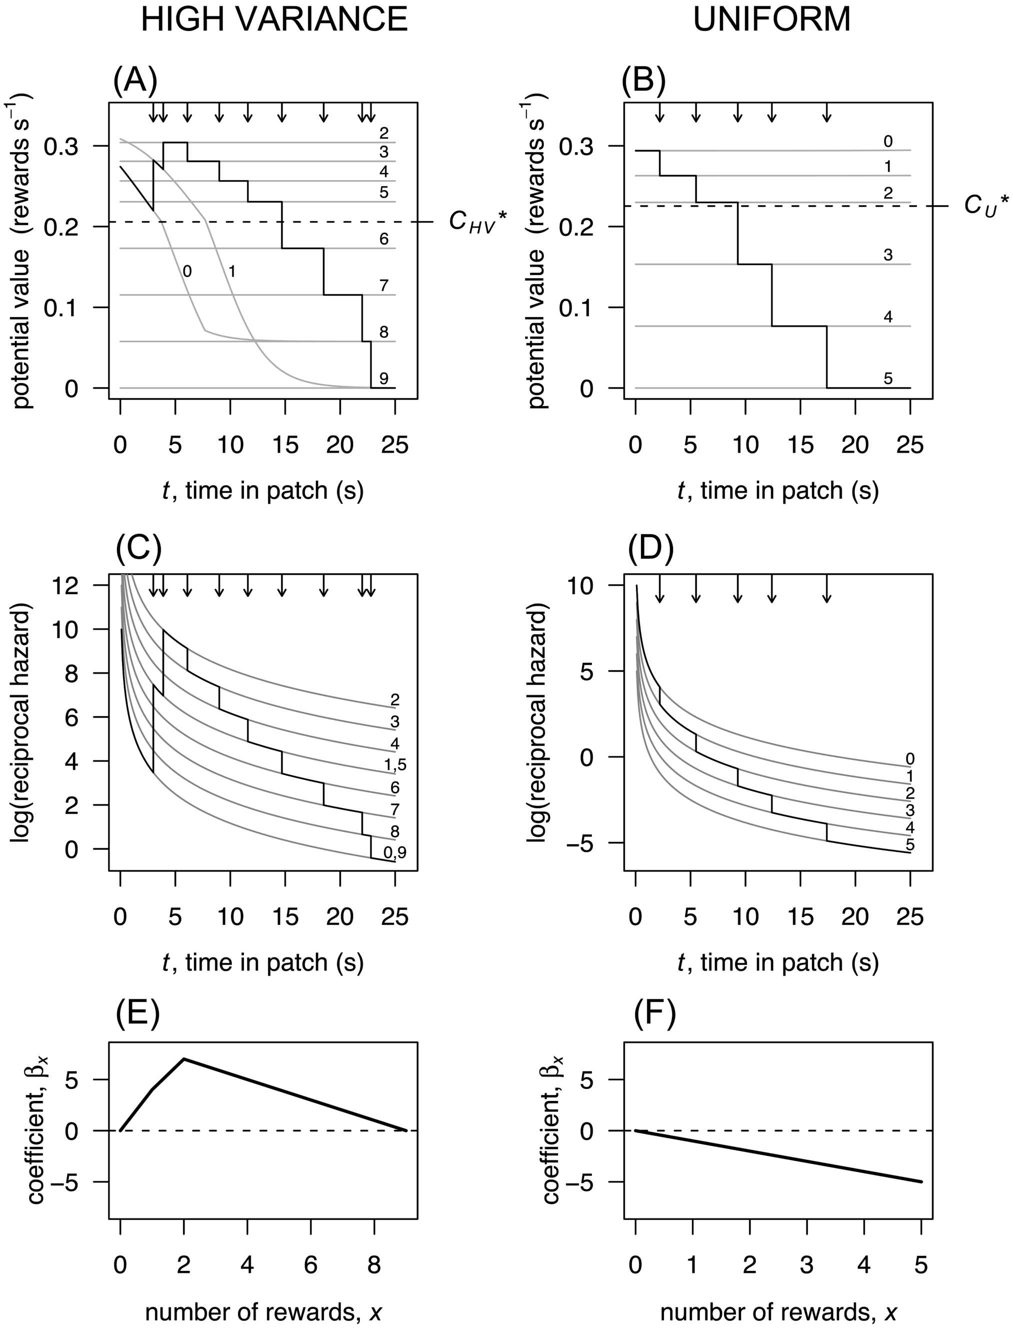

Figure 1: Predicted patch-use behavior of high-variance-trained (*A*, *C*, *E*) and uniform-trained (*B*, *D*, *F*) bees. *A*–*D* measure bees' tendency to stay in a patch (*Y*-axis) as a function of time spent searching the patch (*X*-axis) and the number of rewards found (numbers associated with each solid thin line) under the potential value model (*A*, *B*) and a Cox proportional hazards model (*C*, *D*). Bold thick lines represent changes in the tendency to stay during a hypothetical patch visit, where the times at which rewards are found (*arrows*) are the mean times to finding *x* rewards measured in the late training stage of our experiments. In *A* and *B*, strictly optimal foragers leave the patch when their potential value (i.e., tendency to stay) falls below the threshold value, C^*_{HV} or C^*_{U}, respectively (*dashed lines*). In the Cox model, bees' decision to leave is a random event with a likelihood (i.e., hazard function) that depends on *x* and *t*. Hence, the reciprocal hazard (*Y*-axis; *C*, *D*) measures bees' tendency to stay, and it should mimic the tendencies in *A* and *B* (as shown) if bees approximate the optimal rule. *E* and *F* give the qualitative pattern in estimated Cox model coefficients that we expect to measure if bees do in fact approximate optimal behavior (these predictions correspond to the data in fig. 2). Each coefficient (*Y*-axis) gives a measure of the increased tendency to stay in a patch upon finding *x* rewards (*X*-axis).

hazard for the zero-rewards context is given by the baseline hazard function. Coefficients greater than (less than) 0 indicate an increased (decreased) tendency to stay in a patch, relative to the zero-rewards category. For each of the four data sets, HV training and test and uniform training and test, one categorical-predictor model was fitted.

We used the potential values of Bayesian foragers (fig. 1*A*, 1*B*) to derive qualitative predictions for the reciprocal hazard functions of Bayesian-like foragers (fig. 1*C*, 1*D*). Both models predict an incremental response for HV-T bees that find a first and second reward and a decremental response for U-T bees and for HV-T bees in fully informed contexts. Figure 1*E* and 1*F* translates the predicted adjustments to the reciprocal hazard function in HV and uniform environments, respectively, to the predicted pattern in the β coefficients as rewards are obtained. HV-T bees' coefficients are predicted to increase upon finding the first two rewards (an initial incremental response) and then decrease as more rewards are found (a subsequent decremental response), whereas U-T bees' coefficients always decrease as more rewards are found (decremental response). This pattern of β coefficients is required for qualitative agreement between the Cox model and the potential value model, and it is our primary test of Bayesian-like foraging.

We also fitted Cox models that treated the number of rewards as a continuous predictor:

$$h(t, x(t)) = h_0(t)\exp(-\beta x(t)), \tag{3}$$

where β is a single slope coefficient. In this case, one model was fitted to each of the following six data sets: partially informed HV-T bees (zero to two rewards obtained) in (1) training and (2) test, fully informed HV-T bees (two to nine rewards obtained) in (3) training and (4) test, and U-T bees in (5) training and (6) test. For the uniform and fully informed HV models, we expected negative slopes (a consistent decremental response to each reward), whereas for partially informed models, we expected positive slopes (see fig. 1*E*, 1*F*).

We used the coxph function in the R survival package to fit the 10 Cox models. Data were arranged in the counting process formulation of survival times, which allows coxph to fit models with time-varying covariates (following Therneau and Grambsch 2000). To account for repeated measures on each bee, we used the cluster function, which implements the robust sandwich variance estimator, as recommended by Therneau and Grambsch (2000). We checked for proportionality of the hazard functions among reward categories, an important assumption of the Cox model, using the R cox.zph function, which implements the testing procedure of Grambsch and Therneau (1994). For all models, we failed to reject the null hypothesis that the hazards were proportional (all $P > .05$), indicating that our modeling approach is appropriate. Wald tests were used to assess the overall significance of each model.

Results

Predicted Behavior in the Training Environment

Prediction 1. HV-T bees that find a first and second reward should show an incremental response (positive β coefficients) but all subsequent rewards should decrease their tendency to stay in the patch (fig. 1*E*, 1*F*). Each reward found by a U-T bee should elicit a decremental response (giving increasingly negative β coefficients).

Figure 2*A* and 2*B* shows the estimated β coefficients for HV-T and U-T bees. These figures resemble fairly closely the predictions in figure 1*E* and 1*F*, indicating that bees' behavior at the end of training approximated the predictions of our optimal foraging model. One notable exception is that HV-T bees tended to stay longer than predicted, after exhausting patches with nine rewards (the confidence interval for β_9 does not overlap 0; fig. 2*A*). Because of this exception, we could not detect a significant decremental response of HV-T bees in the fully informed state, using the continuous-predictor model.

Prediction 2. Once in a fully informed state, the potential value model predicts a sharp leaving threshold based on the number of rewards obtained (a fixed number rule). The critical number of rewards is six and three for HV-T and U-T bees, respectively (see the hypothetical patch visits in fig. 1*A*, 1*B*).

There was a large degree of variation in the number of rewards obtained, indicating that bumblebees certainly did not follow a sharp threshold. However, the mean behaviors—6.2 rewards obtained (95% confidence interval: 5.7, 7.7; $n = 10$ bees) by HV-T bees on high-quality patches and 3.5 rewards obtained (3.3, 3.7; $n = 10$) by U-T bees—were very close to the optimal values.

Prediction 3. The predicted long-term reward intake rate of an optimal forager (R^*) in our HV and uniform environments was 0.21 and 0.23 rewards/s, respectively (fig. 1*A*, 1*B*; recall that $C_i^* = R_i^*$).

HV-T bees in late training achieved a mean long-term intake rate of 0.16 rewards/s (0.14, 0.18; $n = 10$ bees), or 78% of the rate of an optimal forager. U-T bees achieved a rate of 0.19 rewards/s (0.16, 0.21; $n = 10$), or 82% of the maximum rate. For comparison, we considered a hypothetical naive strategy of searching each patch for a fixed amount of time, such that the probability of finding every reward would be at least 0.95. The intake rate achieved by this strategy would be only 26% and 35% of the maximum rate in the HV and uniform environments, respectively.

To test whether bees' foraging success in late training

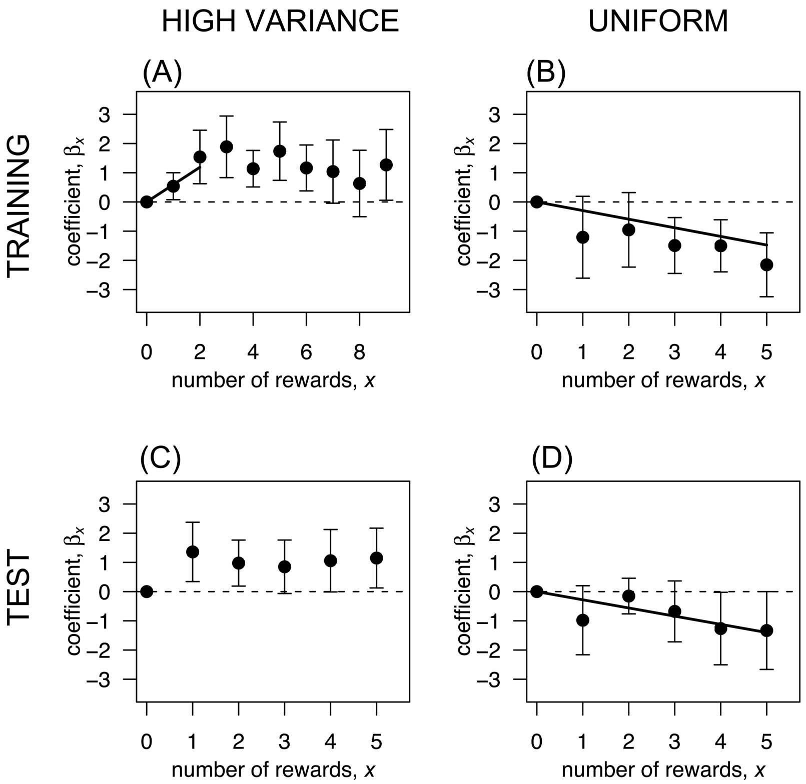

Figure 2: Fitted Cox model coefficients (*Y*-axes) versus number of rewards obtained (*X*-axis) by high-variance-trained (HV-T; *A*, *C*) and uniform-trained (U-T; *B*, *D*) bees in the late training (*A*, *B*) and test (*C*, *D*) environments. Points with error bars (95% confidence intervals) show fits to the models when x is treated as a categorical factor (there is no error at $x = 0$ because this category was used as the baseline). Three of four categorical-predictor models were significantly different from a model in which all coefficients were 0 (HV-training model [*A*], Wald = 1,878, $P < .0001$, df = 9; U-training model [*B*], Wald = 76.5, $P < .0001$, df = 5; U-test model [*D*], Wald = 21.3, $P < .0008$, df = 5), and the model for the HV test (*C*) was very nearly significant (Wald = 10.8, $P = .0544$). Slopes of the lines are given by the significant β coefficients in the continuous-predictor models (they are not least squares fits to the coefficients estimated under the categorical predictor models). Three of six models were significant (partially informed HV-training model [*A*], Wald = 11.2, $P = .0008$, df = 1; U-training model [*B*], Wald = 28.7, $P < .0001$, df = 1; U-test model [*D*], Wald = 7.33, $P = .0068$, df = 1). In general, HV-T bees show an increased tendency to stay on finding the first rewards in a patch (*A*, *C*), whereas U-T bees show a decreased tendency to stay as each reward is found (*B*, *D*).

was due to learning, we also measured the intake rate of the individuals observed during early training. The intake rate of these six bees in early training was significantly lower, on average, than their rate at the end of training (mean difference: −0.044 rewards/s [−0.083, −0.0050]; paired *t*-test, $t_5 = 2.90$, $P = .03$). Another remarkable contrast between early and late training was the difference in time spent searching an exhausted (fully emptied) patch before giving up. By the end of training, U-T bees in particular had much shorter giving-up times than U-T bees during early training (fig. 3; estimated mean in early training: 14.4 s [7.6, 21.1]; $n = 3$ bees; estimated mean in late training: 4.4 s [2.8, 6.0]; $n = 10$ bees).

Patch-Leaving Decisions in the Test Environment

Some differences in the behavior of U-T and HV-T bees in the training environments are confounded with the distribution of rewards in those environments. For example, HV-T bees were much more likely than U-T bees to abandon a patch without finding rewards, as predicted by the potential value model. However, this difference in behavior may not reflect a difference in decision making. It is possible that all bees followed a fixed giving-up time threshold that was simply reached much sooner on low-quality patches in the HV environment than it was in the uniform environment (where rewards were relatively easy to find). The purpose of the test environment, then, is to compare patch-leaving behaviors on a common background. Any differences between HV-T and U-T bees in the test environment must be due to decisions based on alternative prior information learned during training.

Given the same current information about patch quality in the test environment (number of rewards found), HV-T and U-T bees did indeed show very different tendencies

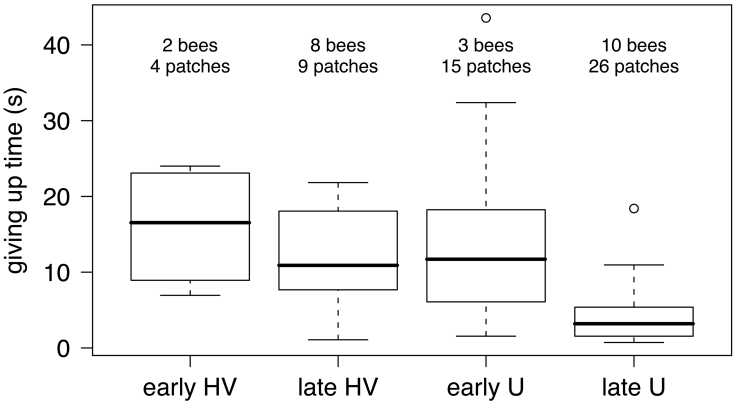

Figure 3: Box plots of observed giving-up times on exhausted patches (time since finding the last reward to leaving the patch) with either five (uniform [U]) or nine (high-variance [HV]) initial rewards in early and late training. The number of patches that were exhausted and the number of bees that exhausted at least one patch are shown. If, by late training, bees had learned the distribution of rewards in their environments (i.e., had become fully informed), they should know that the patch is exhausted and hence search only briefly. This prediction is supported for U-trained bees but not for HV-trained bees.

to stay in a patch (fig. 2*C*, 2*D*). As expected of U-T bees, each reward found tended to decrease their tendency to stay in the patch (a decremental response, as observed in the training environment). In contrast, finding rewards tended to increase HV-T bees' tendency to remain in a patch. This incremental response of HV-T bees in the test environment can be partly explained by their (now confirmed) tendency to abandon patches before finding a reward. Seven HV-T bees left a test patch without finding rewards (for a total of 13 instances), compared to only one instance by a U-T bee (likelihood-ratio test, weighted by the number of instances, $\chi^2 = 14.0$, $P < .0002$). Hence, relative to the baseline tendency for HV-T bees to stay in a patch with no rewards, finding rewards consistently increased their tendency to stay.

Discussion

Testing Optimal Foraging Theory

The first optimal foraging models to include the stochastic element of finding discrete food items in patches (e.g., Oaten 1977; Green 1980; Iwasa et al. 1981; McNamara 1982) revealed that the form of the best patch-leaving rule depends critically on knowing the distribution of food items among patches. More recent techniques (Green 2006) have now allowed us to extend those models to the specific distributions used in our experimental environments, where the theory can be rigorously tested under controlled conditions. By calculating the best rule in both environments, we were able to confirm that Bayesian-like information use is indeed critical to forage successfully (i.e., that a naive forager would perform poorly by ignoring information). We were also able to estimate the overall success of bumblebees relative to an optimal Bayesian decision maker (on average, bees were 80% optimal, according to the potential value model).

Still, even though stochastic optimality models are more realistic than their predecessors (e.g., Charnov 1976), theorists recognize that animal foragers do not actually compute conditional probabilities and cannot realistically follow the deterministic rules of an optimal forager (McNamara et al. 2006). Completing this final link from foraging theory to actual behavior is the first major contribution of our study, by matching detailed optimal foraging models to a plausible, non-Bayesian behavioral mechanism that mimics the best rule (cf. Pierre and Green 2008). In this way, we have shown a remarkable correspondence between the sharp incremental and decremental responses predicted of an optimal forager and the actual adjustments to bumblebees' average tendency to remain in a patch, estimated by the statistical model. As expected of bees with knowledge of a uniform environment, finding rewards almost always decreased their tendency to stay in a patch; as for bees with knowledge of a highly variable environment, finding a reward early in a patch visit sharply increased their tendency to stay.

Bayesian-Like Foraging via Learning and Memory

The use of a proximate rule that mimics a Bayesian estimator of patch quality does not alone imply the use of learning or memory (e.g., Pierre and Green 2008). Numerous examples of either incremental or decremental responses to resource items have been reported in studies of parasitoids (reviewed by Wajnberg [2006]) and a previous study of bees (Lefebvre et al. 2007). In all of these studies, however, it could be argued that the distribution of resources was sufficiently similar to the foragers' ancestral environments, where the observed (incremental or decremental) response to food items may have been hardwired by natural selection (see Wajnberg et al. 1999). The experiments reported here are the first to show that individuals of the same species can adapt to very different environments by using both incremental and decremental responses in the appropriate context (highly variable and uniform distributions, respectively).

Hence, the second major advance of our study is to confirm that bumblebees' tendencies to remain in a patch are indeed shaped by adaptive learning and memory. The adaptive use of a learned prior distribution (the central feature of Bayesian foraging by learning) was clearly demonstrated in our test environment, where the observed behavioral differences between U-T and HV-T bees must have been due to alternative information learned from past experience. Consider, for example, that HV-T bees that were initially unsuccessful in a test patch were very likely to abandon it. This is an adaptive tactic in their previous HV environment, where finding no rewards early in a patch visit indicates a low-quality patch. Abandoning low-quality patches in the HV environment is acceptable if the forager knows that high-quality patches can be found elsewhere. U-T bees with the same average experience in a test patch, however, almost always persisted to find rewards. Persistence is adaptive in their previous uniform environment (and in the test environment, incidentally) because if foragers know that every patch is equal, then finding no rewards early in a visit indicates that five rewards still remain.

The second line of evidence for bees' adaptive learning is the differences in behavior from early to late training. Bees improved their average foraging success (reward intake rate) in a matter of hours (fewer than 45 patch visits), and by the end of training, many aspects of their behavior were consistent with our proximate version of the optimal foraging rule. In particular, it is notable that by the end of training, bees in the uniform environment abandoned exhausted patches only about 4 s (on average) after finding the final reward; in the early training, bees' average giving-up time was more than three times longer. Although the potential value model predicts that U-T bees should never stay on a patch long enough to find all five rewards, the short giving-up times in late training (and the observed decremental response to rewards) do suggest a degree of certainty that five was the maximum number available. This learned giving-up behavior of U-T bees is striking when compared to a well-known study of Bayesian-like birds (Lima 1984; see also Valone 1992). In that study (over a course of 10 days and hundreds of patch visits), woodpeckers consistently oversampled exhausted patches, usually by searching the entire patch, even though only six of 24 units (the equivalent of our flowers) were rewarding.

Admittedly, the bumblebees in our HV environments (both early and late training) similarly oversampled high-quality patches after finding all nine rewards. Accordingly, they did not show a strong decremental response after finding a second reward, implying that, as opposed to optimality assumptions, bees' prior information by the end of training did not fully match the HV distribution. Because bees encountered high-quality patches during only half of their HV training, however, it is possible that their training experience was not sufficient to learn the maximum of nine rewards and the optimal decremental response that U-T bees seemed to learn. Interestingly, a number of authors (e.g., Lima 1984; Valone 1992) have suggested that oversampling is adaptive in foragers' typical (ancestral) environments, where the distribution of resources is rarely so predictable as in a controlled experiment. A highly clumped distribution (e.g., the negative binomial), for example, favors an incremental response to every reward found (Iwasa et al. 1981; Olsson and Brown 2006). Our data suggest the intriguing hypothesis that bumblebees' innate foraging rule, set by evolutionary history, is an incremental response; with enough experience in the appropriate environment, however, they are also able to learn a decremental response. An innate incremental response could be adaptive because clumped distributions have been more common in bees' ancestral environments or perhaps because an incremental response is the best way to learn about a new environment. Our study cannot fully address our hypothesis because bees in early training were first exposed to extremely rewarding flowers during the floral recognition phase of training. Prior information from evolutionary history was therefore briefly tainted to an unknown extent.

Bayesian-Like Foraging in an Ecological Context

Finally, the behavior of nectar foragers has interesting implications in the ecological context of pollination, and our study is relevant in this respect. Some authors have hypothesized that bumblebees' apparently fixed rules of thumb, such as "leave a plant after two consecutive empty flowers," influence the evolution of floral displays and rewards (e.g., Harder et al. 2004), including the phenomenon of nonrewarding (cheating) flowers (Bailey et al. 2007). Our study demonstrates, however, that bumblebees' response to nonrewarding flowers is certainly more dynamic than the proposed rule of thumb. There are at least two situations in which a bee should accept unrewarded search time on a plant: (1) it has experienced a uniform environment in the past and expects to find a certain number of rewards per plant and (2) after experiencing a highly variable environment in the past, the bee finds some rewards on the plant, encouraging it to stay and gain more information about the plant's true quality.

It becomes a greater challenge to determine the outcome of selection on floral traits when pollinator behavior on a plant with a certain trait value depends on the traits of other plants in the environment. It is now clear that bumblebees in particular have a remarkable ability to learn about the resources in their environment. As a result, bees' patch-leaving behavior (and the pollination consequences, from a plant's perspective) should depend on both the mean quality of rewards (Biernaskie and Gegear 2007) and the distribution of rewards (this study) that they learn from past foraging experience.

Acknowledgments

This research was conducted in the laboratory of J. Thomson, and we thank him especially for the use of his resources. We thank P. Abrams and J. Perry for comments on earlier drafts. R. Green, J. van Gils, and E. Wajnberg provided helpful correspondence during the development of this article, and P. Haccou and an anonymous reviewer offered very useful recommendations. J.M.B. and S.C.W. were supported by Discovery Grants from the Natural Sciences and Engineering Council of Canada (NSERC) to P. A. Abrams and D. A. Jackson, respectively, and by University of Toronto graduate fellowships; R.J.G. was funded by an NSERC postdoctoral fellowship.

Literature Cited

Alonso, J. C., J. A. Alonso, L. M. Bautista, and R. Munoz-Pulido. 1995. Patch use in cranes: a field test of optimal foraging predictions. Animal Behavior 49:1367–1379.

Bailey, S. F., A. L. Hargreaves, S. D. Hechtenthal, R. A. Laird, T. M. Latty, T. G. Reid, A. C. Teucher, and J. R. Tindall. 2007. Empty flowers as a pollination-enhancement strategy. Evolutionary Ecology Research 9:1245–1262.

Biernaskie, J. M., and R. J. Gegear. 2007. Habitat assessment ability of bumble-bees implies frequency-dependent selection on floral rewards and display size. Proceedings of the Royal Society B: Biological Sciences 18:125–129.

Blumstein, D. T., and J. C. Daniel. 2007. Quantifying behavior the JWatcher way. Sinauer, Sunderland, MA.

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

Driessen, G., C. Bernstein, and J. J. M. van Alphen. 1995. A countdown mechanism for host search in the parasitoid *Venturia canescens*. Journal of Animal Ecology 64:117–125.

Gegear, R. J., and T. M. Laverty. 2005. Flower constancy in bumblebees: a test of the trait variability hypothesis. Animal Behavior 69:939–949.

Goulson, D. 2003. Bumblebees: their behavior and ecology. Oxford University Press, Oxford.

Grambsch, P. M., and T. M. Therneau. 1994. Proportional hazards tests and diagnostics based on weighted residuals. Biometrika 81: 515–526.

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

———. 2006. A simpler, more general method of finding the optimal foraging strategy for Bayesian birds. Oikos 112:274–284.

Haccou, P., and E. Meelis. 1992. Statistical analysis of behavioural data: an approach based on time-structured models. Oxford University Press, Oxford.

Haccou, P., S. J. De Vlas, J. J. M. van Alphen, and M. E. Visser. 1991. Information processing by foragers: effects of intra-patch experience on the leaving tendency of *Leptopilina heterotoma*. Journal of Animal Ecology 60:93–106.

Harder, L. D., C. Y. Jordan, W. E. Gross, and M. B. Routley. 2004. Beyond floricentrism: the pollination function of inflorescences. Plant Species Biology 19:137–148.

Hodges, C. M. 1985. Bumble bee foraging: the threshold departure rule. Ecology 66:179–187.

Iwasa, Y., M. Higashi, and N. Yamamura. 1981. Prey distribution as a factor determining the choice of an optimal foraging strategy. American Naturalist 117:710–723.

Kadmon, R., and A. Shmida. 1992. Departure rules used by bees foraging for nectar: a field test. Evolutionary Ecology 6:142–151.

Lefebvre, D., J. Pierre, Y. Outreman, and J.-S. Pierre. 2007. Patch departure rules in bumblebees: evidence of a decremental motivational mechanism. Behavioral Ecology and Sociobiology 61: 1707–1715.

Lima, S. L. 1984. Downy woodpecker foraging behavior: efficient sampling in simple stochastic environments. Ecology 65:166–174.

Mangel, M., and C. W. Clark. 1989. Dynamic modeling in behavioral ecology. Princeton University Press, Princeton, NJ.

McNamara, J. M. 1982. Optimal patch use in a stochastic environment. Theoretical Population Biology 21:269–288.

McNamara, J. M., R. F. Green, and O. Olsson. 2006. Bayes' theorem and its applications in animal behavior. Oikos 112:243–251.

Murdoch, W. W., and A. Oaten. 1975. Predation and population stability. Advances in Ecological Research 9:1–31.

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.

———. 2000. Optimal Bayesian foraging policies and prey population dynamics: some comments on Rodriguez-Girones and Vasquez. Theoretical Population Biology 57:369–375.

Pierre, J.-S., and R. F. Green. 2008. A Bayesian approach to optimal foraging in parasitoids. Pages 357–383 *in* E. Wajnberg, C. Bernstein, and J. van Alphen, eds. Behavioral ecology of insect parasitoids. Blackwell, Oxford.

Pleasants, J. M., and M. Zimmerman. 1979. Patchiness in the dispersion of nectar resources: evidence for hot and cold spots. Oecologia (Berlin) 41:283–288.

Pyke, G. H. 1978. Optimal foraging: movement patterns of bumblebees between inflorescences. Theoretical Population Biology 13: 72–98.

Therneau, T. M., and P. M. Grambsch. 2000. Modeling survival data: extending the Cox model. Springer, New York.

Thomson, J. D. 1988. Effects of variation in inflorescence size and floral rewards on the visitation rates of traplining pollinators of *Aralia hispida*. Evolutionary Ecology 2:65–76.

Valone, T. J. 1992. Information for patch assessment: a field investigation with black-chinned hummingbirds. Behavioral Ecology 3: 211–222.

———. 2006. Are animals capable of Bayesian updating? an empirical review. Oikos 112:252–259.

van Gils, J. A., I. W. Schenk, O. Bos, and T. Piersma. 2003. Incompletely informed shorebirds that face a digestive constraint maximize net energy gain when exploiting patches. American Naturalist 161:777–793.

Wagge, J. K. 1979. Foraging for patchily distributed hosts by the parasitoid, *Nemertis canescens*. Journal of Animal Ecology 48:353–371.

Wajnberg, E. 2006. Time allocation strategies in insect parasitoids: from ultimate predictions to proximate behavioral mechanisms. Behavioral Ecology and Sociobiology 60:589–611.

Wajnberg, E., M. C. Rosi, and S. Colazza. 1999. Genetic variation in patch time allocation in a parasitic wasp. Journal of Animal Ecology 68:121–133.

Ydenberg, R. C. 1998. Behavioral decisions about foraging and predator avoidance. Pages 343–378 *in* R. Dukas, ed. Cognitive ecology. University of Chicago Press, Chicago.

Associate Editor: Marc Mangel
Editor: Michael C. Whitlock

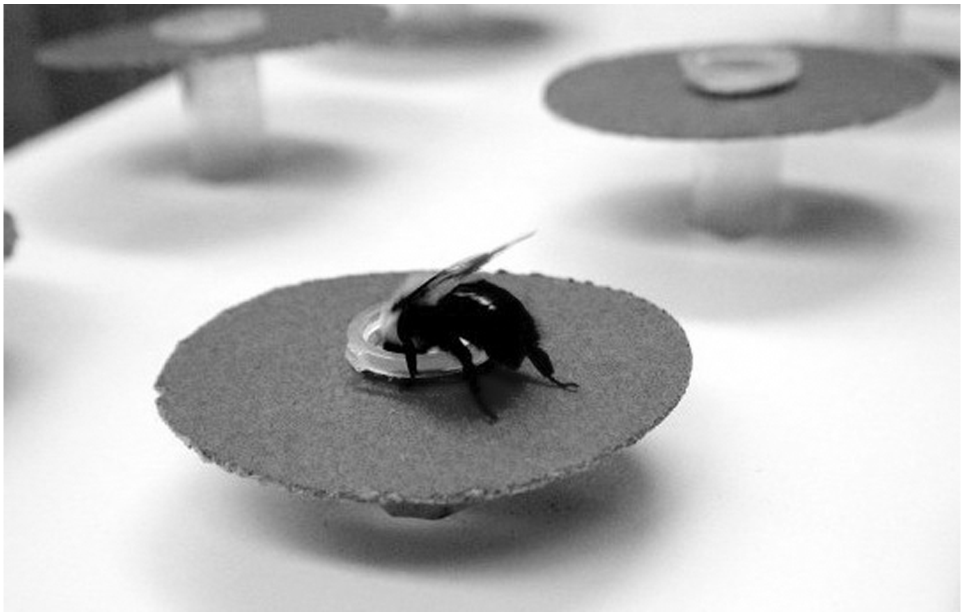

A bumblebee worker finds a rewarding “flower.” Photograph by Jay Biernaskie.