Behavioral Ecology Vol. 8 No. 2: 199–210

Trapline foraging by bumble bees: II. Definition and detection from sequence data

James D. Thomson,[a,b] Montgomery Slatkin,[c] and Barbara A. Thomson[a]
[a]Department of Ecology and Evolution, State University of New York, Stony Brook, NY 11794-5245, USA, [b]Rocky Mountain Biological Laboratory, Crested Butte, CO 81224-0519, USA, and [c]Department of Integrative Biology, University of California, Berkeley, CA 94720, USA

Trapline foraging—repeated sequential visits to a series of feeding locations—presents interesting problems seldom treated in foraging models. Work on traplining is hampered by the lack of statistical, operational approaches for detecting its existence and measuring its strength. We propose several statistical procedures, illustrating them with records of interplant flight sequences by bumble bees visiting penstemon flowers. An asymmetry test detects deviations from binomial expectation in the directionality of visits between pairs of plants. Several tests compare data from one bee to another: frequencies of visits to plants and frequencies of departures to particular destinations are compared using contingency tables; similarities of repeated sequences within bees are compared to those between bees by means of sequence alignment and Mantel tests. We also compared observed movement patterns to those generated by null models designed to represent realistic foraging by non-traplining bees, examining: temporal patterns of the bee's spatial displacement from its starting point using spectral analysis; the variance of return times to particular plants; and the sequence alignment of repeated cycles within sequences. We discuss the different indications and the relative strengths of these approaches. *Key words:* asymmetry test, *Bombus*, foraging, Mantel test, null model, *Penstemon*, sequence, trapline. *[Behav Ecol 8:199–210 (1997)]*

If a foraging animal repeatedly visits a series of fixed resource points or "stations" in a fixed order, we describe the behavior as "traplining" (e.g., Anderson, 1983). Such behavior provokes fascinating questions regarding its genesis, maintenance, and utility. Maintenance and utility are treated in companion papers (Thomson, 1996, Thomson J, Williams N, in preparation). In this paper we confront some definitional problems. Suppose that the set of visited stations expands or contracts on repeated passes through the array. Suppose that the order of visitation is imperfectly replicated. Traplining is easy to define only in its perfected, ideal state. No one has proposed methods to detect traplining statistically or to measure its intensity. This paper explores several remedies, illustrated by observations of bumble bees (*Bombus flavifrons*; Apidae) visiting a designed array of plants of *Penstemon strictus* (Scrophulariaceae). Although we provide some description of what the bees are doing, the paper is essentially a search for an operational definition of traplining.

The term "traplining" was apparently coined by D. H. Janzen to describe a pattern of regularly repeated flower visits by female euglossine bees (see Proctor et al., 1996: 135; Heinrich, 1979:177), although the term does not appear in the most frequently cited reference (Janzen, 1971). The analogy is to a trapper checking traps on a regular basis, and the term has become widely used.

Well before Janzen coined the term, several naturalists studied what we would now call traplining, including Darwin (Freeman, 1968) and Tinbergen (1968). One of Tinbergen's students, Manning (1956), produced the most detailed account of these studies. More recently, traplining by *Bombus* has been reported by Heinrich (1976), Thomson et al. (1982, 1987; Thomson, 1996), and R. A. Johnson (unpublished; see Thomson et al., 1982), while Gill (1988) has done extensive work on hermit hummingbirds. Ackerman et al. (1982) showed that male euglossines, which were previously thought to be widely ranging vagabonds, may sometimes trapline plants much as females do. Aside from studies on flowerfeeders, work on primates has evoked suggestions of traplining as well (Garber, 1988; Janson, 1996; Milton, 1981).

Different criteria have been used to conclude that animals are traplining. Darwin and Janzen mostly drew their inferences from the regular appearance of unmarked individuals at particular stations (although Janzen marked some bees and could identify others by individual characteristics). Darwin noted that he regretted not marking individual bees. Janzen presented one schematic flight map representing a "perfect" trapline, but he did not indicate repeated flights. Without presenting any flight maps, Manning simply stated that marked bees retraced particular pathways. Heinrich (1976) showed sketch maps of small numbers of flights by marked individuals; Thomson et al. (1982) presented similar maps based on more observations. Because hummingbirds cannot easily be followed through tropical forest, Gill's (1988) inferences about traplining behavior derive from the regularity of the reappearance of marked individuals at stations.

None of these authors tested movement patterns for significance. Testing is warranted, however, because any animal using a finite number of stations will occasionally retrace an earlier path by chance alone, and a human observer might subconsciously assign undue weight to such coincidences (Pyke G, personal communication). Without objective, quantitative procedures for trapline detection, observations of this behavior will always seem soft and anecdotal, if not downright dubious, but we currently have no accepted methods for distinguishing traplining from non-traplining behavior or for comparing the strength of traplining among different animals.

Here we try to remedy this lack of methods. We are concerned only with inferring traplining from sequences of interplant movements, not from reappearance schedules. As we will show, traplining is subtle and difficult to demonstrate statistically. It is easy to subject movement data to tests that can reject various hypotheses, but it is harder to devise a hypothesis whose rejection can be considered a sufficient and general demonstration of traplining. For example, we could easily

Received 26 November 1995; accepted 9 June 1996.

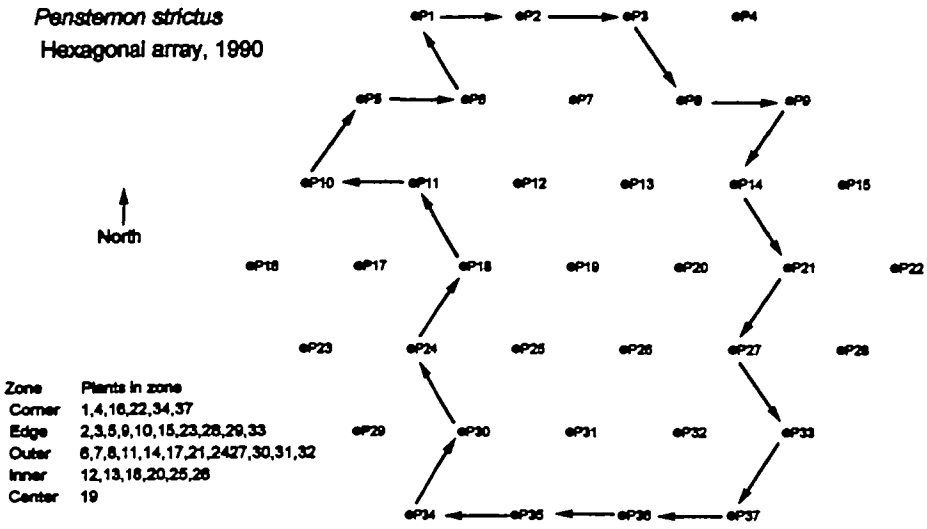

Figure 1
Layout of the hexagonal array of flowering *Penstemon strictus* plants. Interplant spacing is 1.6 m. The plants are divided into several zones that are used in generating null movement sequences. The arrows indicate the model sequence used for a simulation study of measures of traplining strength (details in text).

reject the hypothesis that successively visited plants are randomly drawn from all the plants in a population; no one would claim that such a rejection demonstrated traplining, however. For one thing, the definition of "population" would necessarily be arbitrary: bees that randomly visited plants within a subunit of the study area would appear nonrandom if the entire area were considered. For another, bees that simply moved in such a way as to reduce interplant flight distances (see Pyke, 1978) would fail to match this overly simplistic null model, but they would not show the repetition of flight sequences that (we argue) is an essential characteristic of traplining. Even if a bee were shown to repeat some visit sequences, this need not be a certain indication of traplining. A perfectly randomly foraging bee would occasionally repeat some sequences. The important question is how much repetition constitutes evidence of significant traplining. As soon as one discards true randomness as a meaningful null hypothesis, it becomes virtually impossible to frame a null hypothesis that is neither complicated nor ad hoc.

We pursue several different approaches, each of which illuminates one aspect of the structure of interplant moves. The obvious approach is to construct null models of non-trapline foraging, use these models to generate non-traplining sequences, and compare various properties of our observed sequences to those of the model. This strategy has a clear Achilles' heel: the null model will always be arbitrary, subject to no authority more reliable than common sense. Therefore, we also adopt a less arbitrary procedure that instead asks whether repeated sequences by individual bees differ from those of other bees. We first present several different analyses based on the frequencies of particular categories of movements, then develop a strategy based on pairwise similarities among sequences. This culminates in a Mantel test that establishes simultaneously whether the repeated sequences flown by a bee are significantly individualistic and self-similar in comparison to sequences flown by all bees in a data set. Testing the individuality of bees within a group is not, of course, the same as testing a single bee against randomness, but behavioral ecologists would find both tests relevant to characterizing traplining behavior.

GENERAL METHODS

To obtain data on which to test statistical methods, we sought a plant–bee system in which we could record long sequences of plant visitation by numerous bees. We chose *Penstemon strictus* because of its high natural vistation rate. This species is locally abundant in disturbed roadside habitats below about 3000 m elevation in the West Elk Mountains of Colorado. Plants produce numerous large flowers with a strong nectar flow and a gradual release of pollen that promote very high visitation rates by bumble bees. Our experiments were done above the natural altitudinal range of the plant, in a flower-rich subalpine meadow at Irwin, Colorado, USA (107°06′00″ W, 38°52′35″ N, elevation 3140 m). To our knowledge, no *P. strictus* occurs within several kilometers of the study site. In early June 1990, we planted 37 potted plants in a hexagonal array (Figure 1). The interplant spacing of 1.6 m was chosen to be close enough to allow bees' flights to be easily followed, but far enough that neighboring plants were clearly distinct. The surrounding vegetation was neither weeded nor mowed, so other species of flowers were also available in the array, but these are not shown in Figure 1.

As the plants came into bloom in late July, we used hobby paints (Floquil Corp., New York) to place unique marks on the thoraxes of visiting bumble bees. To stress the bees as little as possible, we caught them in an insect net, immobilized them in a fold of the net, and quickly applied paint through the net's mesh. Most bees were released within 60 s, and a substantial fraction of them returned immediately to foraging.

Over the next several days, two to five observers with tape recorders followed marked bees, noting the times and the sequences of plants visited. Of approximately 30 bees marked, at least 20 were seen again in the array. However, we concentrated on a few bees that appeared to spend most of their time in the study plot, and ultimately chose three *Bombus flavifrons* workers for intense scrutiny. Observers would follow any marked bee if one of the three focal bees were not present, but when a focal bee was seen, an observer would switch to it. From these three bees, we compiled four data sets that are analyzed here. Blue early, red-blue, and pink data sets include all data for the indicated bees from 23 to 28 July; blue late includes all data for bee blue on 5 August. (Neither red-blue nor pink were still working on 5 August.) For some purposes, we further subdivide the blue early data set into two subsets, 23–26 July and 27–28 July. Observations covered the entire activity period of the bees, roughly 0800 to 1800 h.

Because we were interested in flight-path geometry rather than pollination, we recorded a plant visit if the bee approached within 5 cm of an open flower, even if it rejected the flower without alighting. Such rejections were compara-

tively rare; more often, the bee landed for at least a brief inspection before rejecting a plant.

Even with close observation, we occasionally lost sight of a bee in mid-sequence. Unless it was sighted again within a few s, we terminated the sequence, beginning a new sequence when the bee was rediscovered. All the bees occasionally visited flowers other than *Penstemon strictus*. Such visits were recorded on tape, and the locations of the non-*Penstemon* plants were mapped, but the analyses presented here exclude non-*Penstemon* visits. Thus a sequence of *Penstemon* 3 to *Helenium* to *Penstemon* 5 is here recorded as *Penstemon* 3 to *Penstemon* 5. Almost all the non-*Penstemon* visits were short breaks in what was overwhelmingly *Penstemon*-dominated foraging; however, pink was an exception, in that she began to regularly include several *Lupinus* sp. plants in her foraging, especially on 27–28 July.

The various statistical treatments of the data are described separately below. Where our purpose is primarily to illustrate a methodology, we may show the analysis only for selected data sets.

RESULTS

Foraging time budgets

The three focal bees spent most of their foraging time in the array. This was best documented for blue on 5 August. A single observer tried to follow this bee (and no others) continuously from about 0900 until about 1530 h, taking brief breaks when the bee appeared to be leaving for the hive. The total observation period, including breaks, was 22,485 s; blue was in sight for 20,482 s, or 91.1% of the total. As the observer never saw the bee enter the array, but found her only after she had begun foraging, the 91.1% figure underestimates the time she truly spent in the array. We infer that she probably foraged entirely within the array, spending the remaining 5% or so of her time flying to and from the nest and depositing or picking up food there.

This bee had a characteristic way of leaving the array: from one of several *Penstemon strictus* plants in the northwest corner, she would fly in a straight line west–northwest, at about 0.7 m off the ground, passing through a small gap between a large spruce tree and a 2-m spruce sapling. After leaving in this manner, she never reappeared in the array before 3–4 min had passed. We inferred that this was her route home. She tended to make such flights at regular intervals. From 26 through 28 July, we observed enough of these departures to establish 12 bout lengths (Figure 2). These are calculated as the times between successive departures for home, so they include time commuting and time in the hive. There is a strong peak in the frequency histogram, indicating relatively constant trip times. Based on the minimum times between departures and reappearances in the array, we estimate about 29 min foraging and about 4 min travel plus nest time for a bout by this bee. By 5 August, blue's trips had become significantly longer (Figure 2; Mann-Whitney U test, $U = 70$, $N = 12, 5$, $p < .001$).

The bees would make several passes through the array on one trip from the hive. At the plants that blue visited most often, for example, she would appear approximately every 10 min. Therefore, she would typically visit such a plant three to four times on one trip. Considering all bees, large plants received visits roughly every 100 s. Elsewhere, we consider the temporal patterning of visits, using other data sets (Thomson J and Williams N, in preparation).

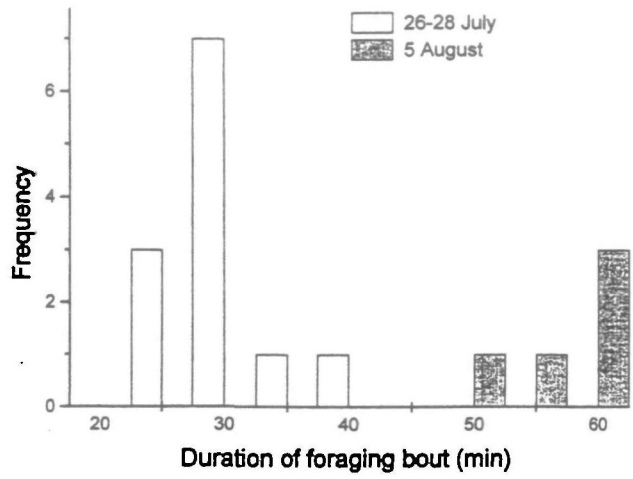

Figure 2
Frequency histogram for the lengths of 12 foraging bouts (including nest time and commuting) by bee blue on 26–28 July (open bars) and for 5 bouts on 5 August (filled bars).

Rewards

Bumble bees showed two behaviors at flowers of *Penstemon strictus*. The most abundant species, *B. flavifrons*, typically enters the flower right-side-up and probes for nectar. There are two lateral nectaries deep in the tubular flower, and one can often see the bee shift position slightly when switching from one to the other. Such bees do not actively collect pollen in the sense of manipulating the anthers (which lie on the roof of the tubular flower); nevertheless, they usually do accumulate small corbicular loads of *Penstemon* pollen, through grooming movements that eventually bring passively applied pollen to the hind legs. An alternative behavior is most often shown by *B. bifarius* (and possibly the similar *B. sylvicola*): these bees enter the flower upside down, grasp the anthers with their legs, and vibrate the flower with a characteristic buzz of their flight muscles. Occasional individuals of *B. flavifrons* mix the two behaviors, but all three of the focal bees in this study showed only nectar-probing behavior at *Penstemon strictus*. Nevertheless, they varied in the extent to which they accumulated pollen loads. Red-blue always carried small loads (i.e., the corbicular surfaces of her tibiae were never more than half covered). Pink always collected large loads, her corbiculae full and bulging. Pink's predilection for visiting lupines as well as penstemons may have contributed strongly to her greater pollen acquisition. Blue tended to accumulate intermediate loads during the earlier observation period, but on 5 August, she collected no visible loads at all.

Flight patterns

Figure 3a–e shows all of the inter-*Penstemon* flights for the five selected data sets (blue early is subdivided for clarity). Each bee visited plants throughout the patch, but there is a clear tendency to concentrate visits within a portion of the array, especially for blue and red-blue, who both preferred the northern half of the array. In contrast, pink ranged more widely. (At the same time, other bees were showing similar preferences for other parts of the array.) Blue's preference for the northern half crystallized with time (compare early and late data sets). These maps show that none of the bees showed strong traplining, in the sense of always going to the same destination from one point of departure. On the other hand, closer inspection of the maps shows that moves from

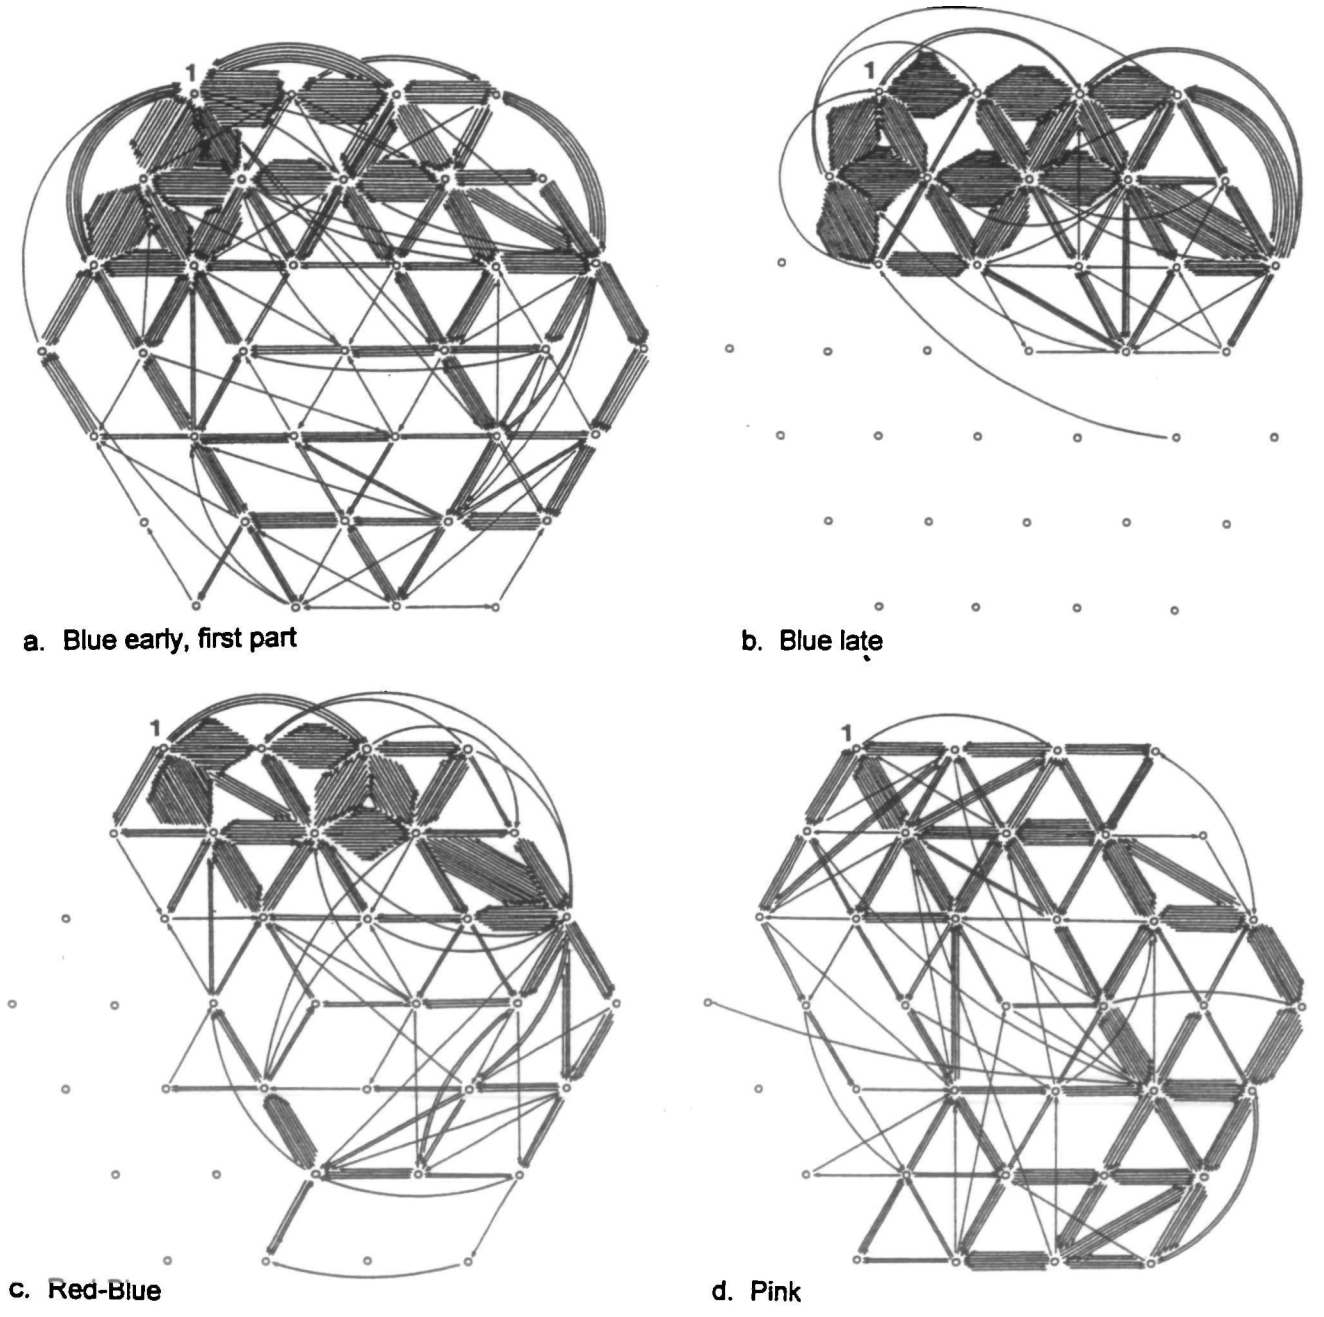

Figure 3
The hexagonal array, with interplant moves shown for selected data sets. The later half of the blue early daty set is not shown. Arrowheads indicate directions of flights.

certain plants were indeed highly directional. At other decision points, the same bee's departure directions might be much more haphazard. The maps also suggest that the edges of the arrays were important in channeling the bees' movements; however, some internal transtions were highly directional.

On 25 and 27 July, we counted the number of open flowers on each *Penstemon* plant. There was little change in these values over this short period. The total number of visits paid to each plant was weakly but significantly correlated with the number of open flowers (Table 1).

Although bees tended to visit more floriferous plants more frequently, this did not cause them to converge in their use of plants. For example, although blue and red-blue did overlap extensively in their foraging areas and even in their general directionality (see below), they distributed their visits differently across plants (2×37 contingency table, with plants lumped where expected values were less than 5; $\chi^2 = 101.3$, 24 df, $p < 10^{-11}$). All other pairwise comparisons of the four data sets were significant at $p < 10^{-7}$.

All bees overwhelmingly tended to fly between immediately neighboring plants (Table 2). There is a suggestion that the three bees differ in their distributions of distances flown (3×4 contingency table, $G = 12.48$, $p = .052$), but this might

Table 1
Pearson correlations of the number of visits to a plant and the number of flowers open on that plant for various data sets (n = 37 for each)

Data set	Flower census	
	25 July	27 July
Blue, 25–26 July	0.452**	0.502**
Blue, 27–28 July	0.474**	0.522**
Red-blue, 25–28 July	0.304 (ns)	0.351*
Pink, 25–28 July	0.505**	0.524**

* $p \leq .05$; ** $p \leq .01$.

reflect differences in opportunity rather than differences in flight behavior: bees that foraged more often at the edges of the array, as opposed to the center, faced a somewhat different distribution of potential neighbor categories.

Transition matrices and an asymmetry test

The next series of tests are based on the transition matrix, where each interplant movement is cast into a 37 × 37 table. Rows indicate the plant the bee moved from, columns indicate the plant a bee moved to. If a traplining bee follows a particular path each time it passes through an array of plants, we would expect the transition matrix to be asymmetrical about the main diagonal; for example, the number of moves from plant 1 to 2 would differ from the number from 2 to 1. Our test followed a suggestion of Oden's, previously used by Sokal (1991): for each pair of plants, we calculated the binomial probability of the observed departure from a 1:1 expectation. To test the entire transition matrix, we used Fisher's method of combining probabilities (Sokal and Rohlf, 1995: 794). We eliminated from consideration any plant pairs for which the data set contained fewer than six transitions, because such pairs give little information on directionality and therefore dilute the test's effectiveness. By this procedure, all data sets for blue and red-blue showed highly significant asymmetry ($p < .01$), but pink's movements did not ($p > .05$). Further trimming the data set to eliminate pairs with <10 transitions did not change these results. Note that this technique detects only unidirectional traplining: if a bee commuted back and forth along a particular sequence of plants, transitions would be symmetrical, even though we would consider such commuting to constitute a special case of traplining.

Trapline skeleton diagrams

We then used the asymmetry analyses to produce a graphic representation of trapline structure. In Figure 4, symbols describing the frequency and directionality of interplant moves are superimposed on the map of plant locations. These maps emphasize those portions of a bee's movements that show structure consistent with noncommuting traplining. It is apparent that bee pink showed little such structure, in contrast to the other two focal bees, which showed a general tendency to pass through the array in a clockwise fashion.

Individuality of movements

Above, we showed that the bees with the most similar use of plants (red-blue and blue) still used their plants with different frequencies. Now we test a second aspect of individuality: when two focal bees have arrived at the same plant, do they move to the same set of plants or do their destinations differ? We tested this by constructing for each plant and for each pair of bees (actually, data sets) a 2 × n contingency table [(bee data set A versus bee data set B) × (the n nearest neighbors of the plant in question)]. Thus, for each plant, we could evaluate the probability that the departure distribution was independent of bee identity; then we combined the probabilities across plants, using Fisher's test as before. We eliminated any plants from which either bee of the tested pair departed less than twice. For each of the six possible pairs of data sets, there is at least one plant at which the bees differ in their departure distributions at $p < .006$. Combining data from all plants, all pairs differ at $p < 10^{-6}$. Thus, all bees make individualistic departure decisions when they are at the same plants.

Table 2
Classification of interplant visits by neighbor status

Data set	Neighbor category			
	Nearest	Second	Third	Farther
Blue early	741	56	35	19
Red-blue	255	33	16	12
Pink	285	26	7	12
Total	1281	115	58	43
Fraction	0.856	0.077	0.039	0.029

"Nearest neighbors" are at adjacent vertices in the hexagonal array, 1.6 m apart. "Second nearest" neighbors are >1.6 but <3 m apart. "Third nearest neighbors" are two vertices (3 m) apart, and "farther neighbors" are >3 m distant.

Direct comparison of sequence data

The above analyses depended only on transitions between successive pairs of plants. To complement this approach, we considered longer sequences. So that we could compare sequences within and between bees, we scanned all data sets for all sequences that began and ended with the same plant (henceforth, the "terminal plant"). Sequences of length 3 (e.g., 1 to 2 to 1) were eliminated because they often represented a bee that had left a plant before it "intended" to—sometimes because of disturbance by observers, other insects, or a gust of wind—and then returned after a brief inspection at a nearby plant. We then subjectively chose a terminal plant for which several sequences were available in each of the bee data sets; for the analysis presented here, plant 8 was chosen. This produced a number of sequences of varying length.

Next, an index of similarity was calculated for all pairs of sequences using a simple technique used for aligning DNA sequences (Waterman, 1989). Although alignment methods are often complex and controversial, in our case where the endpoints are fixed, it was easy to derive an index that counts how many insertions or deletions are necessary to render two sequences the same. The algorithm is best understood by envisioning the two sequences written out as the row and column headings of an $n \times m$ matrix where n and m are the lengths of the two sequences. The elements of the matrix are scored as 1 if the row and column headings match or as 0 if the headings differ. Then dummy rows and columns are inserted to put as many of the 1's as possible on the principal diagonal. As our index of similarity, we divided the number of matches on the diagonal of this expanded matrix by the total number of cells along the diagonal. We did not count the terminal cells, which were forced to match. For n sequences, this produced a symmetric $n \times n$ similarity matrix.

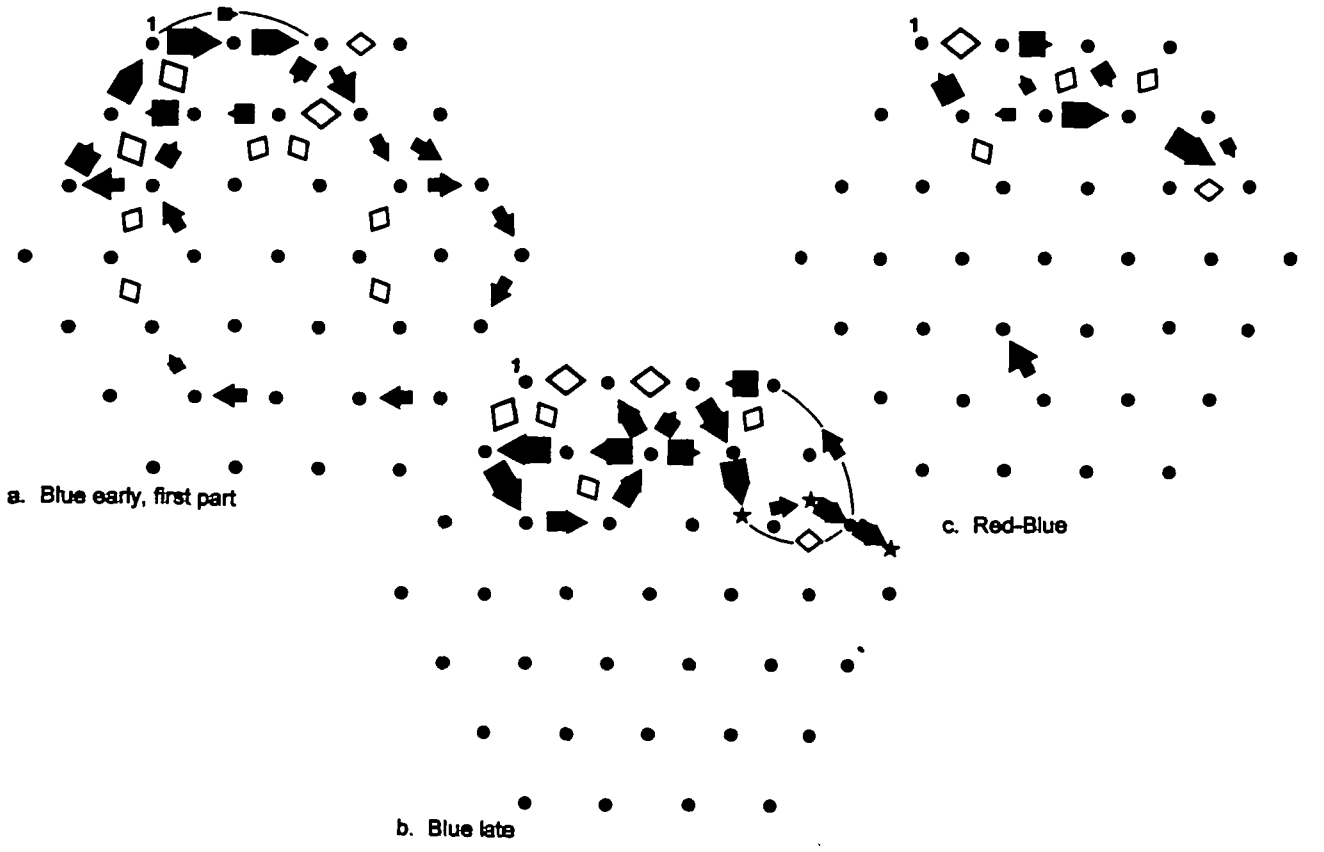

Figure 4
Trapline skeleton maps summarizing the raw maps shown in Figure 3. Arrows indicate interplant transitions for which the bee in question showed directionality. The size of the arrowhead indicates the significance of the directionality: large, medium, and small arrowheads denote p levels of .01, .05, and .1, respectively. The size of the shaft of the arrow indicates the total number of transitions observed: large, ≥15; medium, 9–14; small, 6–8. Where directionality was insignificant even though traffic was high ($p > .1$), diamonds denote sample size (large diamonds, ≥15 transitions; small diamonds, 6–14). Stars indicate non-*Penstemon* plants (not included in numerical analyses). No map is shown for the pink data set because so few transitions showed significant directionality.

The similarity matrix could be subjected to any of the standard ordination or clustering techniques that employ such matrices. Because we were more interested in a significance test for traplining, we instead used a Mantel test (Mantel, 1967; Sokal, 1979; Sokal and Rohlf, 1995), as follows. We produced an $n \times n$ design matrix that contained 1's in the cells where the similarity matrix contained sequences produced by the same bee and 0's in the cells where the similarity matrix contained sequences produced by different bees. That is, the design matrix represents an ideal case in which all sequences generated by one bee show perfect similarity, and all sequences generated by different bees are completely different. The Mantel procedure computes a correlationlike statistic, r (the standardized Mantel statistic), that shows how closely the observed similarity matrix resembles the ideal design matrix. To assess the significance of r, a randomization process then permutes the values in the similarity matrix by redistributing the values across the rows and columns. We produced 999 permutations, calculating r for each one and comparing the true observed value of r to the distribution of the 999 randomized versions. The observed matrix was more similar to the design matrix ($r = .162$) than any of the 999 randomizations, i.e., $p < .001$. Therefore, sequences within bees (or, more precisely, within data sets, because the blue early and blue late sets are here being contrasted as if they were from different bees) resemble each other significantly more closely than they resemble sequences from different bees.

Generating null sequences

Rather than comparing bees with each other, we now compare each data set to neutral expectations. As noted in the introduction, the simplest hypotheses for non-traplining movement (e.g., random plant choices) are too simplistic to be informative about traplining. We know that bees tend to move toward near neighbors in many circumstances (Morse, 1982), including our array (Table 2). Our data would emphatically reject the hypothesis of random moves, but we would not consider this to demonstrate traplining. Instead, we take the observed tendency to make short moves as a fundamental component of foraging that we would expect to see in non-traplining and traplining bees alike. It forms a basic constraint in our null model. As an additional constraint, we also regard a bee's aversion to returning to just-visited plants as another fundamental component. For a bee to be traplining, in this view, its movements must contain a higher-order repetitive structure that cannot be produced by the simple action of these two constraints.

We could not simply use the observed probabilities of first, second, etc., nearest-neighbor moves to condition the null

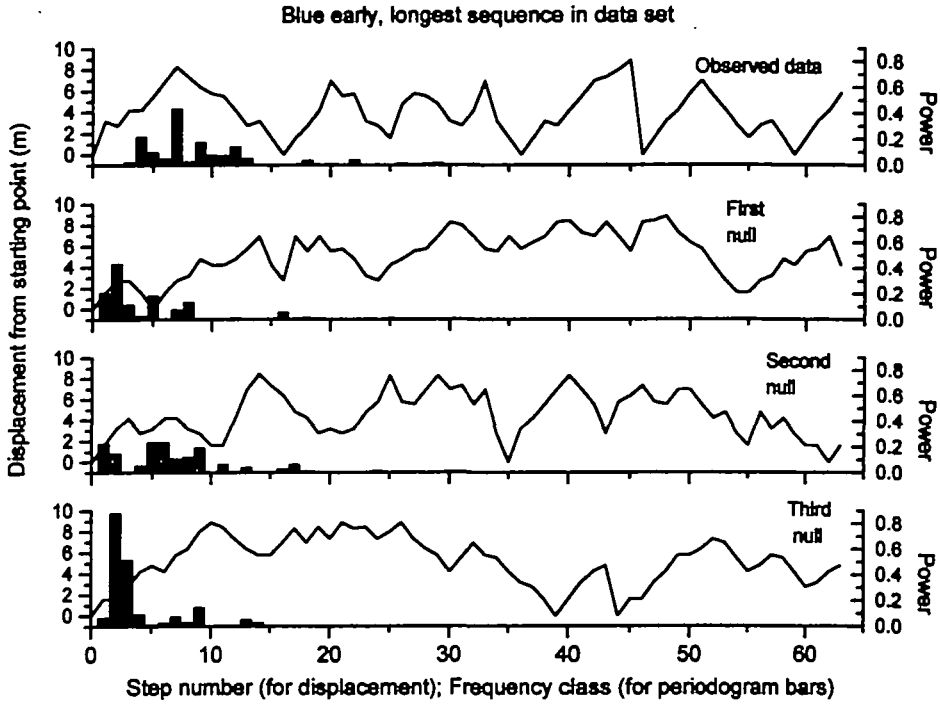

Figure 5
Spectral analysis of spatial displacements from the starting point for the longest sequence in the first data set. The first three null sequences are also shown. The solid line, based on the left y-axis, indicates the bee's straight-line distance from the starting plant as a function of the number of interplant moves or steps that the bee has taken since leaving that plant. The shaded bars, based on the right y-axis, are the periodogram (i.e., the squared Fourier amplitudes), indicating the power of the series at the frequency indicated.

model because the numbers of neighbors in each class vary with a plant's position in the array. For example, central plants have six equally near neighbors, but corner plants have only three. Therefore, we divided the array into five zones (Figure 1). From the data, we calculated the zone-specific frequencies of first, second, third, fourth, and farther nearest-neighbor moves ("neighbor classes") for each bee moving from plants in each zone. Casting these data into a three-way contingency table (bee × zone × neighbor class) revealed a significant three-way interaction; i.e., the bees had idiosyncratic probabilities of moving different distances, depending on where they were in the array. We therefore did not pool the data to calculate an overall set of zone-specific neighbor-class movement probabilities, but rather used the individual probabilities for each "bee" (actually, each data set) to generate null sequences for comparison to that bee's observed moves.

Similarly, we examined the data for bee-specific and zone-specific probabilities of a bee retracing its steps. We looked for one-step retraces (*a* to *b* to *a*) and two-step retraces (*a* to *b* to *c* to *a*). Both were rare, occurring in 1.11% and 1.92% of the cases, respectively. Neither the bee nor the zone of origin significantly affected these probabilities, so we used the pooled estimates in all cases.

In practice, we generated a set of 999 null sequences for each of the observed sequences longer than 19 moves. Our algorithm started the null bee at the same plant at which the observed sequence started and produced a sequence of the same length. Obviously, the first move could not be a retrace, and the second move could not be a two-step retrace, but whenever a retrace was possible, we drew a random number to decide whether a retrace would occur, using the observed probabilities. If we drew a retrace, we moved the bee accordingly and went on to decide on the next move. If we did not draw a retrace, we selected a second random number and used it to select a neighbor class (i.e., nearest neighbor, second nearest, etc.). Having decided on a neighbor class, we determined which plants qualified as possible, non-retrace destinations, then randomly chose one of them to be the destination.

Spatial displacement analysis

We compared observed and null sequences in several ways. We looked not at sequence data per se but at the x–y coordinates of the plants visited. We could then calculate the Euclidean distance of the bee from its starting position as a function of the number of moves it had made ("step number"). These displacements are equivalent to time series and are highly periodic (i.e., sinusoidal) if the bee is faithfully repeating a circuit. Null sequences should be less so. We followed the recommendations of Wilkinson (1990) to compute spectral analyses of the periodicity of the displacement series using the SYSTAT package. Choosing only the longer sequences, we padded the end of each sequence with zeroes to extend its length to a power of two. We then conditioned the data by first subtracting the mean from each value, then applying a split-cosine-bell-tapering function to downweight the ends of the sequence. (Deviating from these recommendations produced only minor changes in results.) We then performed a Fourier decomposition of the conditioned series. We used 52 frequency classes, regardless of the length of the series. The squared magnitudes of the Fourier coefficients are graphed against the frequency classes to produce periodograms. These are displayed, along with the displacements themselves, for the longest sequences in the blue early data set, along with the first three of the matching nulls that we generated (Figure 5).

The observed displacement series appears more regularly periodic than the nulls, which show more of the character of constrained random walks. Graphs (not shown) for blue late had a similar character; red-blue and pink had sequences that were too short to show such clear cyclicity, although the nulls still looked less cyclic than the real data. In the periodograms, a simple sinusoidal function would produce a single sharp peak at the frequency of the sinusoid; more irregular patterning would cause the peak to be less defined or absent. The real sequences do in fact show prominent peaks. Most often, the periodograms for the null series either lack a distinct peak or have a peak nearer the origin than the observed data. Un-

Table 3
Analysis of distributions of return lengths in observed and null sequences

Data set	Sequences analyzed	Ratio of observed variance to index: Weighted null variance (min, med, max)	Ratio of observed variance to index: Mean ratio	Significance for all sequences combined: Sign test	Significance for all sequences combined: Fisher's combined test
Blue early	20	0.034, 0.668, 3.584	0.846	$p < .05$	$.05 < p < .1$
Pink	4	0.611, 1.091, 2.785	1.244	—[a]	$.5 < p < .1$ (ns)
Red-blue	6	0.282, 0.620, 0.867	0.568	$p < .05$	$.5 < p < .1$ (ns)
Blue late	5	0.150, 0.393, 0.446	0.367	—[a]	$p < .001$

To be included, observed sequences had to be 20 or more steps long, with at least one return.
[a] Too few sequences.

fortunately for our goal of testing for traplining, the periodograms include too much information to be well summarized by a single parameter. This prevents a randomization-based significance test.

Variance in return lengths

We used sequence data, looking at the distribution of "return lengths," i.e., the number of steps taken before returning to the same plant. Our algorithm counted all returns to all plants, so one sequence typically contained a number of returns. We calculated the variance of these returns for each observed sequence and for its 999 null equivalents. Perfect traplining should produce a variance of zero; imperfect but significant traplining should yield a variance among the lowest 5% of the nulls. Some null sequences contained no returns or only one, leaving the variance of return lengths undefined. A total lack of returns is equivalent to a total lack of traplining, so for ranking the observed sequence with respect to the null sequences, we scored all nulls that lacked returns as if their variances had been higher than the observed sequences.

The results (Table 3) are roughly concordant with the results of the asymmetry tests. None of the four sequences from pink showed a significantly lower variance than the matching null distributions, suggesting again that this bee was not traplining. The blue early, blue late, and red-blue data sets contained some sequences for which the observed variance was significantly lower than for the null equivalents but also contained sequences for which this was not true. To combine data from different sequences, we used both sign tests and Fisher's combined probabilities test. For all bees, seven sequences had higher variances than the null expectation, while 28 had lower variances ($p < .01$, sign test). However, considering the bees individually, only blue late had a significant Fisher's combined test ($p < .001$). Blue early is significant at $p < .05$ by the sign test (5 greater, 15 less) but not quite significant by the Fisher's test ($.1 < p < .05$). Similarly, red-blue is significant by the sign test (0 greater, 6 less), but not by the Fisher's test. With only four sequences, neither pink nor blue late can be subjected to the sign test.

We calculated an index of return variability for each data set. For each sequence, we divided the variance of the return lengths for the observed sequence by the mean variance of return lengths in the matching set of 999 nulls. We then took as our index the weighted mean of these ratios for all the sequences within a data set, using sequence length as a weighting factor. Our rationale for weighting was that longer sequences generally contain more returns and therefore give more reliable information about the distribution of returns. The weighted mean ratio should be zero for a perfectly traplining bee and 1.0 for a non-trapliner that exactly matches the null model; thus, the index of return variability captures an important aspect of traplining. As expected, pink is the most variable, blue late the least (Table 3).

Finally, we examined the internal similarity of sub-sequences within sequences. If the bee returned to the starting plant twice or more in the same sequence, we computed the similarity measure based on maximum alignment (as used above for the Mantel test) for all pairs of sub-sequences that started and ended with that terminal plant. This provided a measure of how consistently the bee repeated a path through the array; sequences by traplining bees should have high internal similarity. Again, we computed the same statistic for each of the 999 nulls, ranked the observed sequence among the nulls to produce a p value for each observed sequence, then used Fisher's test to arrive at a combined probability across all sequences (Table 4). Because this test requires at least two complete cycles within each sequence, it can be computed only for longer sequences. Only 16 sequences, across the four data sets, meet this criterion. For only one of these, a sequence from blue late, is the individual-sequence p value significant at the .05 level. However, when p values are combined across sequences, the blue early set becomes significant at $p < .025$, and when the four data sets are pooled, $p < .01$. For the sequences that meet the test's criterion, there is a weak but significant tendency for repeated cycles in the observed sequences to resemble each other more than cycles in the null sequences do.

Table 4
Summary of internal similarity (by maximum alignment) of repeated sub-sequences in observed and null sequences

Data set	Sequences analyzed	Mean p value	Overall significance (by Fisher's test)
Blue early	8	.184	<.025
Pink	1	.284	>.1 (ns)
Red-blue	3	.172	<.1 (ns)
Blue late	4	.375	<.1 (ns)
Pooled	16	.236	<.01

Statistics were defined only for observed sequences in which the bee returned to the initial plant twice, so only longer sequences could be analyzed.

Comparison of index performance

To gain additional insight into the behavior of the various approaches for measuring the intensity of traplining, we created simplified synthetic sequences that embodied various strengths of repetition, as follows. Using the hexagonal array, we produced a model sequence that began at plant 1 and returned there after describing a simple noncrossing loop of

20 nearest-neighbor steps with no first-or second-step retraces. We then simulated sequences of 401 steps in which "bees" started at plant 1, then repeated this path with varying degrees of fidelity determined by a trapline-strength parameter, t. When the bee was at any plant in the model sequence, it moved to the next plant in the model sequence with probability t. With probability $(1 - t)$, the bee moved randomly to any nearest neighbor. When the bee strayed from the model sequence, it moved to randomly chosen nearest neighbors until it returned to a plant in the model sequence. Thus, for $t = 1.0$, the bee repeated the model path exactly, returning to plant 1 after 20 identical circuits of perfect traplining. For $t = 0$, the bee moved in a random nearest-neighbor walk, subject to the constraint on retraces: no traplining. For intermediate values of t, the bee would tend to follow the model path but would be subject to getting lost more or less often.

We generated sequences for values of t ranging from 0 to 0.9 by tenths, then subjected them to three of the analyses described above: asymmetry test, variance of return lengths, and within-sequence alignment of subsequences. To summarize the results of the asymmetry test, for each value of t, we calculated a simple asymmetry index, based on the Fisher's combined test, by dividing the sum of the logs of the binomial transition probabilities by the number of i,j pairs in the test. This is essentially equivalent to dividing a G statistic by its degrees of freedom. We then scaled these asymmetry indices by dividing each one by the asymmetry index for $t = 0$.

For the variance of return lengths and the within-sequence alignment statistics, we generated 999 runs per t value, then used the $t = 0$ sequences as the null baseline to which we compared the others for significance testing (as well as for scaling). (Using the null-sequence algorithms that contained bee-specific transition probabilities would have been inappropriate.) In addition to calculating the variance and alignment values for a single simulated sequence of length 401, as we did for the asymmetry index, we also calculated the mean of the 999 iterated sequences for each t value.

The results (Figure 6) suggest that all three indices vary systematically with t, as expected. However, values of $t < 0.5$ produce sequences that are virtually indistinguishable from random walks. At low t values, apparently, bees tend to wander off the model sequence more frequently than they return to it. Only at relatively high values of t did we produce sequences of length 401 that showed "significant traplining" according to the tests described above. Sequence similarity becomes insignificant below $t = 0.7$, although significant asymmetry remains at $t = 0.5$. This makes sense; the algorithm used will still generate some level of asymmetric transitions whenever the bee wanders back into the model sequence. If real traplining follows this model closely, a weak tendency to trapline (i.e., small t value) will be difficult to detect. In this small study, the variance of return lengths method was the most problematic for significance testing, in that some sequences with lower t values were more significant than some with higher values.

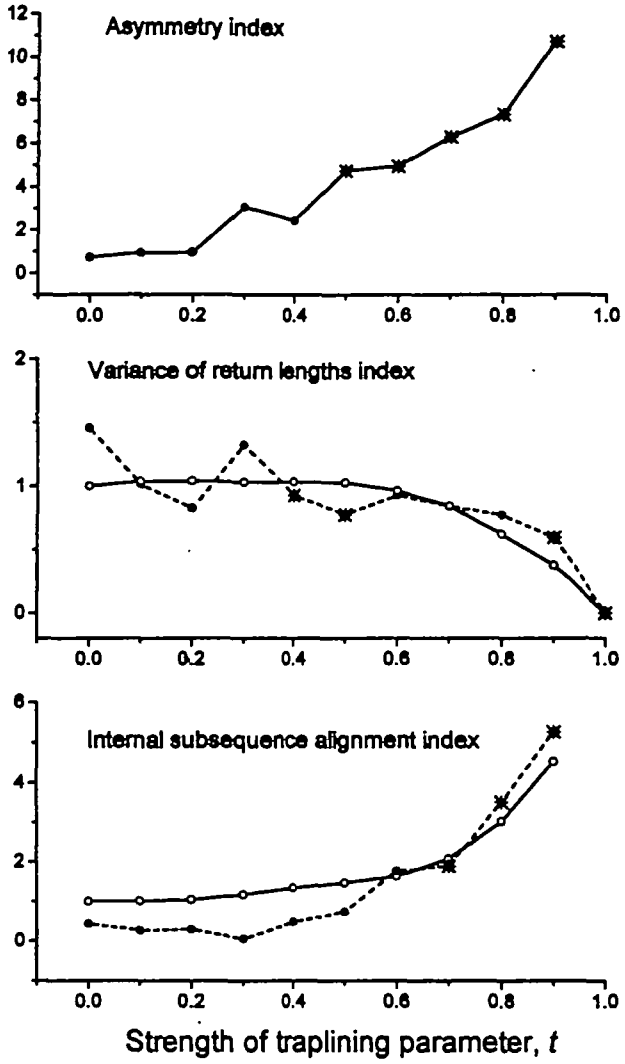

Figure 6
Responses of three indices of traplining to variation in trapline strength in simulated sequences generated by a stochastic algorithm incorporating the trapline-strength parameter, t, as described in the text. Filled symbols represent indices calculated for single sequences of length 401; asterisks replace filled circles for indices that significantly exceed those generated by a null algorithm. Open circles indicate means for 999 stochastic sequences. The three indices are described in the text.

DISCUSSION

Possible levels of traplining

Our study is at the level of the plant. This is a reasonable choice; as Manning (1956) and others have shown, bees do learn the locations of potted plants and return there after the plants have been taken away. However, traplining might also occur at larger or smaller scales (e.g., among flowers, whorls, or inflorescences offered by a single plant, or among clumps of plants or even subpopulations). A full analysis would include these other hierarchical levels.

What the bees are doing

All of the bees examined restricted their foraging spatially to the *Penstemon strictus* array, although all occasionally visited plants of other species as minor components of their foraging (cf. Heinrich, 1979). Our analyses neglect these inconstant visits, not because they are uninteresting but because they were numerically rare and because they were difficult to incorporate in our models. This restriction to a small, local foraging area is interesting in itself. Although it is not a focus of our study of trapline detection, it is a prerequisite for traplining to occur at the scale of our study. Traplining does seem to occur in this system. Of the three focal bees, two were fairly consistent in showing nonrandom, repetitive movement patterns among plants. The third, pink, was fairly consistent in conforming to null hypotheses, although we suspect that we

could have detected significant patterning of her movements with larger data sets.

We do not know whether bumble bees usually trapline. Because many observations are needed, one can detect it only in plant species that receive high visitation; these species probably have unusual reward amounts or schedules. There are also observational biases. In choosing a focal bee for accumulating a large data set, it is natural to choose one that is frequently seen. Many of the bees that we marked in the *Penstemon* array were not seen again. Marking trauma may have driven them off, but they may also have been non-traplining vagabonds. However, other studies (Thomson, 1996; Thomson J and Williams N, in preparation) reinforce our conclusions that *Bombus flavifrons* on *Penstemon strictus* frequently forage similarly to *Bombus ternarius* on *Aralia hispida* (Thomson, 1988; Thomson et al., 1982, 1987). The foraging areas are similar, (about 100 m^2), the circuit times are similar (about 10–15 min) and the sharing of plants by many bees is similar. In both cases, bees make three to four passes through the array on one foraging trip from the nest. Additionally, we know that on *Aralia*, bees gradually move their foraging areas into areas where floral rewards are higher. It appears that bees visit a core set of plants on virtually every pass through the array, but they also occasionally sample other plants. If those plants prove rewarding, they are more likely to be visited on a subsequent pass (Thomson et al., 1982, 1987). We suspect that bees on *Penstemon* do something similar (Thomson, 1996).

If so, there is no reason to expect perfect traplining; indeed, it would be pathological. Rather, traplining probably presents a special case of the common situation where animals need to learn information about their environment and remember it for awhile, but also to forget it when it loses currency (Mangel, 1990; Thomson, 1996). Occasional sampling is probably highly advantageous in a world where plants change in value through time, but it will blur the conservative visitation pattern that must underlie traplining. Therefore, we stress the need for techniques that establish the reality of that underlying pattern.

Comparisons of methods

We have proposed several different tests of traplining. Each of these illuminates a different aspect of potentially repetitious movement patterns among fixed stations. Each can thus be viewed as a different operational definition of traplining. Prospective users should select tests, or devise new ones, that are relevant to the biological situation and compatible with the data. Here, we consider some of the properties of our tests that such users should bear in mind.

Asymmetry test

The asymmetry test, which was developed to test for directional migration (Sokal, 1991) is the only one that compares data to a simple statistical distribution (i.e., binomial expectation). It is readily grasped and shows an important aspect of nonrandomness in interstation moves, and its application allows the construction of the trapline skeleton diagrams, which we find useful. The index derived from this test also responds to a wider range of t values than the other indices (Figure 6). The test is rather far from the essence of traplining, however. It would also be blind to commuting traplining, in which a bee might go back and forth along a particular route. It does not serve to define traplining, but it is a useful adjunct to other measures.

By expressing sequence data as a transition matrix, we draw attention to the role of Markov models as possible descriptors of traplines. C. H. Janson (unpublished manuscript) has extensively analyzed traplining from a Markovian view, so here we will cite a few points. For example, perfect traplining of a subset of plants in an array would result in a nonergodic Markov chain: in such a chain, some destinations are never reached and the effect of the starting position is never dissipated, as happens in the more frequently modeled ergodic condition. Our analysis of individuality of transitions via contingency table could be seen as a special case or subset of a Markov analysis; for example, our analysis could be extended to consider whether bees' departure directions from particular plants depended not only on the plant of departure (first-order Markov process) but on the preceding plant as well (second-order process). Janson (personal communication) shows that the movements of a monkey troop cannot be described a first-order Markov process, which he interprets as evidence for a cognitive spatial map.

Null models

Aside from the asymmetry test, our measures fall into two categories. Some compare animals to animals, others compare animals to null models. Each type has characteristic shortcomings. As mentioned above, random null models are unrealistic unless we temper pure randomness with additional biological constraints. If we do add constraints, our choice of constraints becomes part of the operational definition, and someone who makes different choices will reach different conclusions about the strength of traplining. For example, we built in an aversion to retraces, but we considered only one- and two-step retraces. This decision was based on years of watching bees on many host plants and acquiring the strong impression that such short retraces are so rare—so "unnatural" for bumble bees in general—that we should not permit our null bees to make them freely. We could have gone further, also constraining three-step, four-step, and all higher retraces so that they too would occur as often in our null sequences as in the observed data. Similarly, rather than choosing a destination plant randomly from the eligible members of a chosen neighbor class, we could have made the null bees more likely to go to members with more flowers. At some point, however, we would have so constrained our null sequences that they would have *all* the properties of the observed data. To be useful, a null model has to fall into a window of credibility: if we use unconstrained randomness, we will always reject the null, but the rejection will not constitute traplining. If we constrain too completely, we will never reject the null, and no behavior can constitute traplining. Anyone who chooses the null model approach must also choose the constraints and be able to defend those choices.

With that caveat registered, we suggest that the best practical index or measure of traplining for our study is that based on the variances of returns. It uses data fairly efficiently, is comparatively easy to understand, allows a significance test, and also produces a comprehensible index of the strength of traplining. On the other hand, it is somewhat removed from the ideal definition of traplining in that it does not look directly at sequence similarity. If sequences are similar, returns lengths will be similar, but return lengths could conceivably be similar without much sequence similarity. Furthermore, the weak performance of significance tests on long sequences generated with various values of t is a concern.

The technique closest to the ideal definition of traplining is the comparison of similarity (alignment) of repeated cycles within sequences. Its only real drawback is a serious one, its dependence on long, continuous observations. Given that null sequences must match the observed sequence (they must start at the same plant, be of the same length), it is hard to combine data from different sequences, and only the longest of our sequences have two or more full cycles. Given that few

animals will be easier to observe than bumble bees, this technique will find little application unless a way is developed to use shorter cycles, perhaps by some objective paste-together technique.

Spectral analysis of displacements is a powerful way to search for periodicities in the bee's spatial position, but it is more suited for analyzing a few long sequences in depth than for measuring traplining strength. The product of spectral analysis (the periodogram) is essentially just a transformation of the data from the time domain to the frequency domain. It is therefore as complicated as the original data series itself and does not provide any single parameter that naturally serves as a summary statistic. We recommend spectral analysis of displacement data as an informative way to examine data, but it should always be accompanied by drawing maps of the flight paths, which may be just as informative. It seems to have little to contribute to the search for a general operational definition of traplining, and like all of the null model approaches, it depends on the assumptions built into the null model.

Comparing bees to each other rather than to "randomness" neatly evades the problem of arbitrary assumptions. However, it skips the most fundamental question, Are bees repeating sequences? and jumps to what would logically be a second question, Given that bees show consistent movement patterns, are their patterns individualistic? This question gives information about traplining only if we reject the null hypothesis of homogeneity among bees. If we succeed in showing that different bees make consistently different moves at certain stations, we can consider that as evidence of traplining: consistent differences between bees cannot arise without consistency within bees. However, if bees are basically doing the same thing, we cannot tell whether they are traplining or not. In our case, both contingency tables and the Mantel test indicate that bees do differ, both in their decision-making at particular points in the array and in the similarity of their sequences. A limitation common to both of these methods is the requirement that the bees' foraging areas overlap to some extent.

Of these two approaches, the Mantel test seems closer to the essence of traplining because it examines sequences. On the other hand, it is a less familiar approach than contingency tables. Also, although we have not explored the test's behavior fully, the overall Mantel results seem rather volatile. Eliminating only a few sequences, perhaps 5% of the total, can change the matrix correlation from highly significant to insignificant. Until a more systematic sensitivity analysis has been completed, including an examination of alternative similarity measures, we would be reluctant to recommend this approach. Contingency tables seem adequate for testing individuality, although our use of separate tables and Fisher's combined probabilities test is less elegant than Janson's (1996; unpublished data) approach to the same question using Markov models.

Data requirements and limitations of the array

In the search for statistical signatures of traplining, we repeatedly encounter a requirement for large data sets. The four focal bees were obviously far from being pure, regular trapliners; indeed, several tests support the conclusion that pink was not traplining. The other bees were irregular enough that only by compiling many observations could we find various sorts of repetitive patterns in their movements. It makes sense that we cannot recognize a trapline as a repeated pattern unless we have seen an animal traverse it numerous times during a long slice of time. This leads to a dilemma: statistically, traplines are a moving target.

We know from other studies (Thomson et al., 1982, 1987) and from the comparison of blue early with blue late in this study, that the foraging areas of bumble bees change through time. This is accomplished by adding some new plants to the route while deleting others. Rapid trapline drift will mean that the first sequences in a long data set may resemble the last ones so weakly that the hypothesis of traplining is rejected, even though the animal's route has evolved only through gradual, incremental change. Gradual evolution of this sort is different from sloppy repetition of a basically fixed cycle, but our techniques will not make the distinction: both processes will weaken traplining. If, for example, we had pooled the early and late data sets for bee blue, several indications that were significant in each separate data set may have been insignificant in the whole. This is essentially a stationarity problem, and we have no tested solution to offer. Given enough data, one could gauge the problem by ordinating sequences by the alignment-similarity matrix, then connecting the points for each bee's sequences in temporal order. If the resulting traces form small, nonoverlapping clusters, trapline drift should pose no problem for detecting individuality. If the traces are long, intertwining with others, then stationarity may be insufficient.

It is also evident that observers should strive to get sequences that are as long as possible. Although short fragments of sequences can be useful for analyses based on the transition matrix, other techniques, such as spectral analyses and alignment of cycles within sequences, depend on long records. These methods will not be useful for animals that cannot be followed for long bouts. This brings up another unfortunate trade-off. Our restricted array was designed to ease observation. Working in a much larger natural stand, with irregular plant spacing, would make it harder to maintain continuous records, as bees are more easily lost on longer flights. On the other hand, a larger stand would make it much easier for techniques such as displacement series to distinguish traplining bees from more random foragers. In an infinite stand, a non-traplining bee could potentially drift away from its starting point without limit, more like a true random walk (Kareiva and Shigesada, 1983). In our case, a bee that limited itself to *Penstemon strictus* could drift no farther than a few meters. Such bees necessarily retrace their steps frequently (Figure 4), making it harder to distinguish traplining from constrained null movement.

Using a small array highlights another elusive distinction. We could imagine a bee that uses the entire array by passing through it in a repeatable sequence—an indubitable trapliner. We can also imagine a bee that restricts itself to only half the array, but within that half shows no more repetition than would be expected from a null process. (See the blue late data set, Figure 3.) Both of these bees would show significant traplining according to our tests that compare them to null bees using the entire array. This does not seem desirable, but there is no clear way to fix the problem. To satisfy our curiosity, we reanalyzed the blue late data, using null models in which visits were allowed only to plants 1–9, 11–15, and 19–21. (We also purged the observed data of the single visit to plant 27 that started one sequence.) As expected, the deviation of the observed data from the null models weakened. For the variance of returns test, the Fisher combined significance level declined from .001 to .025; for the alignment of sub-sequences test, it remained insignificant at $.5 < p < .1$.

It seems, then, that significant regularity remains within this data set even when the array is collapsed to include only those plants that the bee actually frequented. This "fix-up" evades the crux, however. Why draw the line at plants that are never visited? Why not further condition the null bee to visit plants in the same proportion as the observed bee? Because this

would eliminate a principal feature of traplines: the freedom of a bee to move preferentially to some plants rather than others. If an observed bee always avoided one particular plant while visiting those all around it, we would consider that avoidance part of its regular route, part of the selectivity that defines traplining. We would not constrain null bees to avoid it also. Although this case may seem different from the type of selectivity shown by blue late, the difference is one of scale and the distinction is gray rather than black and white. There is no way to deconfound avoidance of plants from avoidance of foraging areas as sources of non-randomness. This is another reason why the null model comparisons are arbitrary: they demand assumptions about the available foraging arena as well as about movement rules. For this reason, again, the between-bee comparisons are necessary adjuncts to the null models.

Overall, no single method provides a complete solution to the operational definition of traplining. If the goal is to document that animals are showing statistically significant regularities in their repeated passes through an array of stations, it is probably best to apply as many methods as possible, with each test contributing its unique increment to a general diagnosis. More often and more usefully, model selection will be guided by an adaptive analysis of the problem facing the animal. An investigator might ask what sort of repetitive foraging patterns one might expect from bees using a particular tactic or facing a particular spatial problem. Then, one should select a test that is appropriate to the properties of the data (e.g., sequence length) and the question. We hope that the suggestions here provide a starting point.

We thank Andrea Lowrance, Lisa Rigney, and George Weiblen for field assistance. Useful discussion has been contributed by too many colleagues to name, but Ralph Cartar, Charles Janson, and Neal Williams stand out. Manuscript preparation was funded by the National Science Foundation (IBN 9316792 and a Mid-Career Fellowship to J.D.T.) and eased by the hospitality of of the Hastings Reservation and the Integrative Biology Department at the University of California, Berkeley. This is contribution 971 in Ecology and Evolution, SUNY-Stony Brook.

REFERENCES

Ackerman DJ, Mesler MR, Lu KL, Montalvo AM, 1982. Food-foraging behavior of male euglossinin (Hymenoptera: Apidae): vagabonds or traplinersf Biotropica 14:241–248.

Anderson DJ, 1983. Optimal foraging and the traveling salesman. Theor Popul Biol 24:145–159.

Freeman RB, 1968. Charles Darwin on the routes of male humble bees. Bull Br Mus Nat Hist Hist Ser 3:177–189.

Garber PA, 1988. Foraging decisions during nectar feeding by tamarin monkeys (*Saguinus mystax* and *Saguinus fuscicollis*, Callitrichidae, Primates) in Amazonian Peru. Biotropica 20:100–106.

Gill FB, 1988. Trapline foraging by hermit hummingbirds: competition for an undefended, renewable resource. Ecology 69:1933–1942.

Heinrich B, 1976. The foraging specializations of individual bumblebees. Ecol Monogr 46:105–128.

Heinrich B, 1979. "Majoring" and "minoring" by foraging bumblebees, *Bombus vagans*: an experimental analysis. Ecology 60:245–255.

Janson CH, 1996. Toward an experimental socioecology of primates: examples from Argentine brown capuchin monkeys (*Cebus apella nigritus*) In: Adaptive radiations of neotropical primates (Norconk R, Rosenberger A, Garber P, eds). New York: Plenum Press.

Janzen DH, 1971. Euglossine bees as long-distance pollinators of tropical plants. Science 171:203–205.

Kareiva PM, Shigesada N, 1983. Analyzing insect movement as a correlated random walk. Oecologia 56:234–238.

Mangel M, 1990. Dynamic information in uncertain and changing worlds. J Theor Biol 146:317–332.

Manning A, 1956. Some aspects of the foraging behaviour of bumblebees. Behaviour 9:164–201.

Mantel N, 1967. The detection of disease clustering and a generalized regression approach. Cancer Research 27:209–220.

Milton K, 1981. Distribution patterns of tropical plant foods as an evolutionary stimulus to primate mental development. Am Anthropol 83:535–543.

Morse DH, 1982. Behavior and ecology of bumble bees. In: Social insects, vol. III (Hermann HR, ed). New York: Academic Press.

Proctor M, Yeo, P, Lack A, 1996. The natural history of pollination. London: HarperCollins.

Pyke GH, 1978. Optimal foraging: movement patterns of bumblebees between inflorescences. Theor Popul Biol 13:72–98.

Sokal RR, 1979. Testing significance of geographic variation patterns. Syst Zool 28:227–231.

Sokal RR, 1991. Ancient movement patterns determine modern genetic variances in Europe. Hum Biol 63:589–606.

Sokal RR, Rohlf FJ, 1995. Biometry. San Francisco, California: W. H. Freeman.

Thomson JD, 1988. Effects of variation in inflorescence size and floral rewards on the visitation rates of traplining pollinators of *Aralia hispida*. Evol Ecol 2:65–76.

Thomson JD, 1996. Trapline foraging by bumble bees: I. Persistence of flight-path geometry. Behav Ecol 7:158–164.

Thomson JD, Harder LD, Peterson SC, 1987. Response of traplining bumble bees to competition experiments: shifts in feeding location and efficiency. Oecologia 71:295–300.

Thomson JD, Maddison WP, Plowright RC, 1982. Behavior of bumble bee pollinators of *Aralia hispida* Vent. (Araliaceae). Oecologia 54: 326–336.

Tinbergen N, 1968. Curious naturalists. Garden City, New Jersey: Anchor Books.

Waterman MS, 1989. Mathematical methods for DNA sequences. Boca Raton, Florida: CRC Press.

Wilkinson L, 1990. SYSTAT: the system for statistics. Evanston, Illinois: Systat, Inc.