


HONEYBEE DEMOCRACY

THOMAS D. SEELEY

PRINCETON UNIVERSITY PRESS
PRINCETON AND OXFORD





BUILDING A CONSENSUS



We deprecate division in our Meetings and desire unanimity. It is in the unity of common fellowship, we believe, that we shall most surely learn the will of God.

—*Society of Friends, Book of Discipline, 1934*



A dissent-free decision. This is what normally arises from the democratic decision-making process used by house-hunting honeybees and, quite frankly, I find it amazing. We have seen in the last two chapters how the debate among a swarm's scout bees starts with individuals proposing many potential nesting sites, vigorously advertising the competing proposals, and actively recruiting neutral individuals to the different camps. All this makes the surface of a swarm look at first like a riotous dance party. Yet out of this chaos, order gradually emerges. Ultimately the debate ends with *all* the dancing bees indicating support for just *one* nesting site, usually the best one. Exactly how the scout bees achieve unanimity at the end of a protracted debate is the subject of this chapter.

Consensus building is sometimes the basis of democratic decision making in human groups—such as trial juries, Quaker meetings, and groups of friends—but it is not so common. What is common is for a human group to end a debate, election, or other democratic process with its members strongly divided in their preferences. At this point the group must invoke some formal decision rule, for example, majority rule or a weighted-voting system, to translate its split vote into a single choice. This kind of group decision making has been called “adversary democracy” because it arises from a group of individuals who have conflict-

ing interests and different preferences. In contrast, the group decision making of swarm bees is “unitary democracy” since it involves individuals who have congruent interests (choose the best homesite) and shared preferences (small entrance opening, etc.). Thus, in looking closely at the inner workings of the unitary democracy of a honeybee swarm, we will be examining a democratic process that is intriguingly different from our all-too-familiar adversary democracy. Later in the book (in chapter 10), I will discuss some practical lessons that we humans can learn from the bees for improving human group decision making, especially when the members of a group have common interests, as do the bees in a swarm.

The group solidarity with which a swarm’s scout bees end their debate is critical to the success of the entire swarm. After all, a swarm contains just one queen, so when a swarm takes off to fly to its new home, it needs to do so as a single cohesive entity that travels to a single new homesite. Split decisions are wasteful and can even be fatal. As we have seen with Lindauer’s Balcony swarm (see fig. 4.4), if a swarm takes off with the scouts still strongly advertising multiple homesites, the swarm won’t succeed in moving to any of the sites, hence it wastes time and energy. And if the swarm loses its queen during the aerial tug-of-war between the different parties of scout bees, then it pays the ultimate price of complete failure, for it is doomed without its queen. It seems of paramount importance, therefore, that a swarm’s scouts reach an agreement on just one site among the many that have been found before the swarm launches itself into flight.

A good way to begin to understand how the scout bees achieve unanimity is to reexamine the synoptic records of the scout bees’ debates. Consider the debate of Swarm 3, summarized in figure 4.7. It shows two striking phenomena that must be explained to understand how a swarm’s scouts build a consensus. First, there is the curious way that the support for the winning site—site G in the southwest—grew steadily and ultimately dominated the discussion. Between 1:00 and 3:00 p.m. on July 20, only 4 out of the 30 (13 percent) dancing bees advertised site G. But by 9:00 to 11:00 a.m. on July 21, 32 out of the 52 (62 percent) dancing bees advertised this site. And on the morning of July 22, shortly before swarm departure, 73 out of the 73 (100 percent) dancing bees advertised site G. Presumably, site G was the best of the 11 sites considered by this swarm, since swarms generally choose the best of the candidate homesites under consideration

(chapter 5). So our first critical puzzle about the bee's system of decision making by consensus building is this: What causes the support of the scout bees for the best site to grow and grow throughout a debate?

The second striking phenomenon shown in figure 4.7 is the way that the support for all the poorer sites eventually evaporated. We can see that sometimes the loss of support happened quickly, as in the case of site A in the east. And sometimes it happened gradually, as with site B in the south. But sooner or later, all the bees that performed dances for the poorer sites lost their enthusiasm for them and ceased advertising these sites. The attrition of support for the rejected sites can also be seen at the sites themselves. Figure 5.7 shows, for example, how in the best-of-N experiment on Appledore Island the counts of scout bees at all the nest boxes except the chosen one dropped essentially to zero by the end of each trial. So our second critical puzzle about the bees' method of consensus building is this: What causes the support of the scout bees for the poorer sites to fade over the course of a debate?

Lively versus Lackluster Dances

We know that a swarm contains approximately 10,000 worker bees and that a few hundred of these bees function as nest-site scouts. We also know that a swarm's scouts locate a few dozen candidate nest sites that deserve to be advertised with waggle dances. Each candidate site is originally discovered by a single scout bee, the one who chances to find it while prospecting knotholes, crevices, and other dark places for a good nesting cavity. This means that only a few dozen scouts truly discover the sites that get debated during a swarm's decision making; most scouts learn about and become committed to a particular site by being recruited to it. Each of these recruits follows a dance advertising a site, flies out, locates the advertised site, and makes an independent evaluation. If the proposed residence satisfies her scrutiny, then she too will dance for it when she returns to the swarm.

Given these facts about scout bees, we can view a swarm's democratic choice of its future domicile as a kind of election process in which there are multiple candidates (possible nest sites), competing advertisements (waggle dances) for




Fig. 6.1 The transitions between states that scout bees can undergo, from neutral scout to supporter for a site, and then back to being a neutral scout bee.

the different candidates, individuals who are committed to this or that candidate (scouts supporting a particular site), and a pool of individuals who are still neutral (scouts not yet committed to a site). The scouts supporting a site can produce dances that will convert neutral individuals into additional supporters for their site. Also, the scouts supporting any given site can become apathetic voters and rejoin the pool of neutral scouts. The whole decision-making process can be depicted schematically as a set of positive feedback loops of recruitment of neutral bees into supporters for the different sites, along with “leakage” of some supporters back into the pool of neutral scouts (fig. 6.1).

Looking at the scout bees’ debate in this way, it is clear that in order for the supporters of the highest-quality site to be successful in ultimately dominating the debate, they must do the best job of gaining converts, presumably by showing the greatest zeal in advertising their site. Does this happen? More specifically, when an evangelizing scout bee advertises a potential nest site with a waggle dance, does she adjust the strength of her dance in relation to the absolute goodness of her site? If all the scouts do likewise, then the highest quality site should indeed receive the most compelling advertisements.

The first evidence that this actually happens comes from observations made by Martin Lindauer in the summer of 1953. He set up an artificial swarm in the broad moorlands east of Munich, and there he also set out two empty wooden hives 75 meters (about 250 feet) from the swarm. On the first day of this experiment, scouts from Lindauer’s swarm quickly discovered his two hives sitting exposed in the windswept fields, and they performed rather sluggish dances advertising their two finds. Little by little, there grew a small crowd of inquisitive scouts at each hive. By the end of the first day, Lindauer had labeled 30 dancers total for his two hives. On the second day, Lindauer noticed an exceptionally lively dancer on

the swarm cluster, a scout who turned out to be advertising a snug underground cavity located beneath a tree stump in the corner of a small woodlot. This site was thoroughly protected from the wind by thick bushes, had a 3-centimeter wide (1.2-inch diameter) entrance opening and a 30-liter (27-quart) cavity volume, and was wonderfully dry inside despite heavy rains in recent days. It was a perfect bee home! Lindauer normally killed all bees advertising rogue sites, but on this day he wisely made an exception; this excited bee was allowed to continue announcing her discovery. Within an hour, other boisterous dancers were also indicating the natural nest site, and after another hour, the scouts were dancing unanimously in favor of this site. It was the clear winner in this debate.

The fact that the scout bee that discovered this first-class dwelling place announced her find with an eye-catching dance, even though she had not visited either of Lindauer's test hives, suggested to him that scouts are able to judge the absolute quality of a site through reference to an innate scale of nest-site goodness. Also, the fact that this first dancer and her fellow advocates of the tree stump site danced more strongly than the bees advertising the two hive sites gave Lindauer an indication that a scout's dance provides information not only about a site's *location* but also about its *quality*. He summarized his observations by reporting: "The most lively dances indicate a nest-site of the first quality; second-rate homes are announced by lackluster dances."

Representing Site Quality in Dance Strength

Good decision making by a honeybee swarm depends critically on the scouts adjusting dance strength in relation to site quality, so that scouts advocating higher quality properties are better at attracting additional supporters. Nevertheless, it was not until the summer of 2007 that I looked closely at how nest-site scouts provide information about site quality in their waggle dances. I had recognized for years that Lindauer had made only preliminary observations on this important subject, so I had long known that more convincing evidence was needed, but I had procrastinated.

I left this gaping hole in the analysis open for so long because I had little doubt that what Lindauer had claimed was correct: better sites elicit stronger dances. It





Fig. 6.2 Movement pattern of a bee performing the waggle dance. Each dance consists of a series of dance circuits. Each dance circuit contains a waggle run (W) and a return run (R, alternating right and left). The duration of the waggle run depends on the distance to the target (food source or nest site). The duration of the return run depends on the desirability of the target. As target desirability increases, return run duration decreases, making the dance appear livelier.

was certainly consistent with what I had observed here and there. For example, I had often noticed how some scout bees perform longer and livelier dances than others. Also, from the best-of-5 choice test conducted on Appledore Island where I had seen scouts performing dances, side by side, for either a 40-liter or a 15-liter nest box (see *Window on a Bee's Mind*, in chapter 5), I had seen that the bees reporting on the better homesite performed stronger dances. Furthermore, in previous studies by myself and others on how a honeybee colony wisely deploys its foragers among nectar sources—a group decision-making process that depends on a colony's foragers making graded advertisements of the various nectar sources—we had found that the richer the nectar source that a bee exploits, the greater the number of dance circuits she produces when she returns to the hive and advertises the source. In short, the richer the nectar source, the stronger the waggle dance. We had also figured out how a dancing bee adjusts the number of dance circuits that she produces in relation to nectar-source richness. She does so by adjusting two aspects of her dancing: the *rate* of dance circuit production (R , in dance circuits per second) and the *duration* of dance circuit production (D , in seconds) (see fig. 6.2). The total number of dance circuits produced (C , in dance circuits) in a dancing bee's advertisement is the product of the rate and duration of her dancing ($C = R \times D$). So, richer nectar sources elicit livelier (higher R) and longer-lasting (greater D) dances than do poorer nectar sources. These findings about nectar-source foragers matched perfectly with Lindauer's report that nest-site scouts announced an inferior nesting place with a "faint-hearted dance" while those from a superior nesting place "solicited with a lively and long-lasting dance."

By 2007, however, I had reached the point in my analysis of the bees' house-hunting process where I knew that I really needed to get solid, quantitative information about how the scouts code nest-site quality in their dances. To accomplish this, we would need to work under the controlled conditions provided by Appledore Island. I say *we* because I teamed up with two collaborators on this project, Marielle Newsome, an undergraduate student at Cornell, and Kirk Visscher, a behavioral biologist from the University of California at Riverside. Marielle had done beekeeping with her father and was headed to graduate school at the University of Michigan to study insect behavior, so she was keen to help. Kirk is a longtime collaborator in various bee studies, going back to when we were both students at Harvard, and he has always been the best possible partner: intelligent, skilled, good-natured, and highly enthusiastic.

Our plan called for positioning one artificial swarm in the center of Appledore Island and two nest boxes 250 meters (820 feet) from the swarm but only about 40 meters (130 feet) apart so the swarm's scouts would be likely to find the two boxes more or less simultaneously (fig. 6.3). One box offered a high-quality (40-liter) nesting cavity while the other presented a medium-quality (15-liter) one. For each of the first five to seven scout bees that appeared at each box, we would record with a data logger when she was at "her" nest box and we would

Fig. 6.3 Marielle Newsome recording visits by individual scout bees to a nest box housed in the orange shelter. In the background, 40 meters away, Kirk Visscher is doing the same at the second nest box.






Fig. 6.4 A scout bee that has been labeled with a paint dot on the thorax, applied through the coarse-mesh netting of an insect net.

record with a video camera when she was at the swarm and how strongly she danced to advertise her site. Analysis of the video recordings in the evening would reveal exactly when each scout danced and how many dance circuits she produced. What made the execution of this experiment seem daunting at first was the fact that to examine the behavior patterns of individual scout bees, we would need each scout to be individually identifiable as soon as she was sighted at one or the other nest box. I expected this need would require us to laboriously prepare swarms in which each bee was labeled for individual identification (see fig. 4.5). We certainly had no way to know in advance which of the thousands of bees in a swarm would first appear at our nest boxes, so we couldn't label in advance just the few first pioneering scout bees.

Fortunately, Kirk had an ingenious solution to our scout bee ID problem. In a previous study, he had found that he could apply identifying paint marks to a scout bee during her visit to a nest box without distressing her. To do so, he placed a small insect net over the box's entrance after he saw a scout bee go inside to inspect the box's interior. Then, when the scout came out a minute or so later

and flew into the net, he would gently pin her between folds of the coarse-mesh netting material. Next, he applied paint marks to her thorax through the netting material, taking care to keep the paint from soiling her wings (fig. 6.4). Finally, he would release her near the nest box's entrance opening, hence right where she had been caught. Amazingly, scouts show no sign of being troubled by this bizarre experience—a genuine abduction by aliens—for immediately upon release they resume scrutinizing the nest box.

Working on Appledore Island for most of the month of July, we performed seven trials of our experiment and succeeded in seeing how 41 and 37 scouts advertised the 40-liter and 15-liter nest boxes, respectively. The first thing we noticed in performing this experiment was that the scout bees reported on their sites for at most a few hours, and that an individual scout bee's report was often spread over several trips back to the swarm. These features of scout bee behavior can be seen in figure 6.5, which shows the records of the 11 scouts observed in the trial conducted on July 17, 2007. We see that the first scout bee (labeled with a red dot, hence named Red) showed up at the 40-liter nest box at 9:33 a.m. Red then spent about 10 minutes examining the nest box inside and out, whereupon she flew back to the swarm and excitedly announced her discovery with a waggle dance that lasted for six minutes and comprised 162 dance circuits. She then flew off the swarm, was sighted again at the 40-liter nest box around 10:00 a.m., and then spent another 10 minutes there before returning to the swarm at 10:10 a.m. Now Red spent about six minutes crawling across the surface of the swarm, but this time she did not perform a waggle dance. Indeed, even though she revisited the nest box again from 10:16 to 10:26, the only dance she ever performed was the one extremely enthusiastic and persistent dance containing 162 dance circuits during her first return to the swarm. Note, too, that Red even stopped visiting the nest box after about 10:30 a.m. So, curiously, within about an hour of discovering the 40-liter nest box and announcing her important find with an impassioned waggle dance, the scout, Red, had lost her enthusiasm to perform dances for and make visits to her high-quality nest site. (How and why scout bees stop advertising and visiting a prospective homesite will be discussed later in the chapter.) For the rest of the morning, Red was seen hanging out at the swarm, sometimes crawling slowly about but mostly sitting motionless, utterly indistin-




Fig. 6.5 Activities of 11 scout bees reporting on a 15-liter or a 40-liter nest box. In the horizontal time line for each bee, black diamonds denote sightings of the bee at the nest box, white bars show blocks of time spent at the swarm, and black bars within the white bars indicate periods of waggling. Numbers above the black bars specify the number of dance circuits performed.

guishable from the vast majority of the swarm's quiescent members except for her shining dot of bright red paint.

Looking at the records of the other 10 scout bees shown in figure 6.5, we can see that Red's behavior was typical. Whether a bee visited the 40-liter (high-quality) nest box or the 15-liter (medium-quality) one, the pattern of her visits to her potential homesite was basically the same. Each bee made an initial inspection of the box that lasted 5 to 35 minutes, then she flew back to the swarm and spent 5 to 30 minutes there, often announcing the site with a waggling dance, then she revisited the nest box for another 10 to 30 minutes, and then she returned again to the swarm for another 5 to 40 minutes, perhaps performing another waggling dance. Such shuttling between swarm and nest box usually went on for another

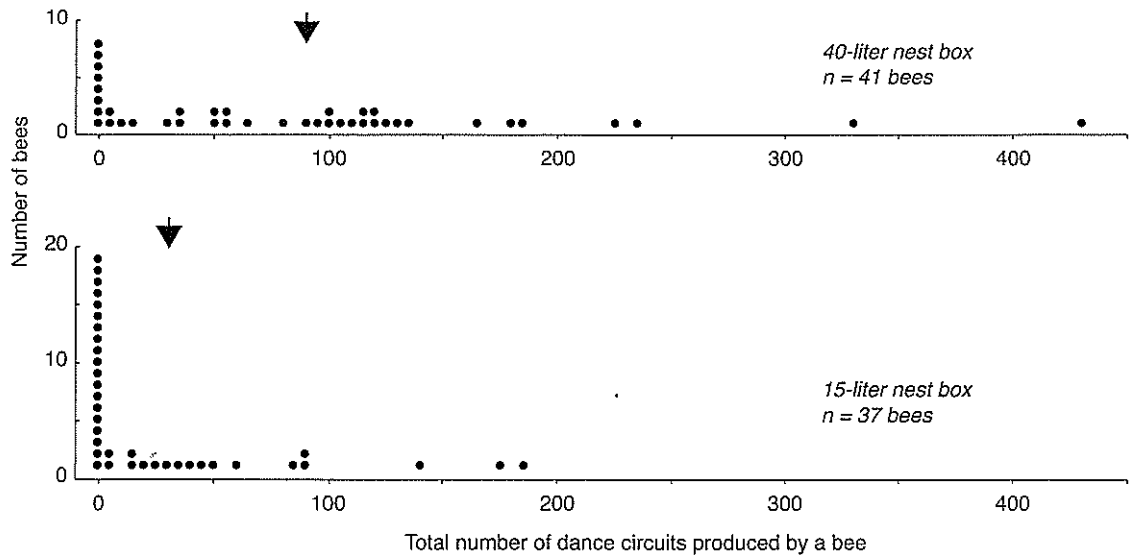


Fig. 6.6 Distributions of the number of dance circuits produced per scout bee, for bees reporting on either a 40-liter or a 15-liter nest box. The number shown for each bee is the sum of all the dance circuits that she produced over her multiple trips back to the swarm. Black arrows indicate the mean values for the two distributions.

hour or so, during which time the scout first lost her motivation to advertise the site and later lost even her desire to visit the site.

Our main finding from this work was that there was a strong difference between the scout bees from the high-quality (40-liter) and the medium-quality (15-liter) nest boxes in how zealously they advertised their sites, that is, in the total number of dance circuits a bee produced before she retired from active scouting duty. As is shown in figure 6.6, there was great variation among the bees of each group, but on average the dance circuit total per bee was higher for scouts from the 40-liter box compared with those from the 15-liter box: 89 versus 29 dance circuits per bee. And it seems clear that these scout bees could tell whether their site was high or medium in quality during the first visit to the site, for upon the first return to the swarm cluster, 76 percent (31 out of 41) of the scouts from the 40-liter nest box advertised it with a dance, but only 43 percent (16 out of 37) of the scouts from the 15-liter box did so.

At first, we were surprised to find so much variation (“noise”) in the strength of the scout bees’ reports on their sites, because the large scatter in the bees’ reports for both the high-quality and the medium-quality sites created much overlap between the distributions of dance strength for these two sites. It is only on average that a higher-quality site elicits a dance with more circuits. On a bit of reflection, however, we realized that for the swarm as a whole the reporting on each candidate site is spread over many bees. So even though there is noisy individual-level reporting of site quality, there is clear swarm-level reporting of site quality. In other words, at the group level there is a sharp difference in strength of advertising between alternatives that differ in quality. The superiority of swarm-level reporting relative to individual-level reporting is demonstrated as follows. If one takes *one* bee’s report at random from each of the two distributions shown in figure 6.6 and compares the number of dance circuits in these two reports, and then one does this over and over, one will find that the advertising of the 40-liter box is stronger than the advertising of the 15-liter box only about 80 percent of the time. This shows that a single scout from the better site does not always make a stronger report than a single scout from the poorer site. But if one takes *six* bees’ reports at random from each of the two distributions, sums the dance circuits in these six reports, and compares the total number of dance circuits for the two groups of six bees, and then one does this over and over, one will find that the total advertising of the 40-liter box is greater than that of the 15-liter box not 80 percent of the time, but 100 percent of the time! This shows that six bees from the better site will always make a stronger collective advertisement than will six bees from the poorer site. So if a swarm finds itself in a choice between two acceptable candidate sites like our 40-liter and 15-liter nest boxes, it is highly likely that the force of persuasion—the total number of dance circuits produced for a site—will be greater for the better site.

The group-level reporting of information about nest-site quality neatly solves the problem of noisy individual-level reporting of this information once there are multiple bees advertising each option. But at the start of the decision-making process, when the scout bees are just starting to discover, inspect, and report back on potential nest sites, there will be only a few bees reporting on each site, so noise in the scout bees’ reporting remains a serious problem at the outset. The poten-

tial for decision-making errors arising from individual-level noise in reporting on sites is especially great when each site is discovered, for if the scout that discovers a site fails to report on it with a waggle dance, the site won't be entered into the scout bees' debate. Indeed, it will be lost from the swarm's attention unless another scout happens to find and report the same site, which is most unlikely. A solution to this problem would be to have each scout bee that discovers a site likely to report on the site and thereby enter it into the debate. Marvelously, the bees appear to do exactly this. In our experiment, Marielle, Kirk, and I found that the two scout bees that first visited the two nest boxes in each trial almost always (with probability of 0.86) performed waggle dances upon return to the swarm, whereas the scouts that visited the same nest boxes subsequently, probably having been recruited to the boxes, were somewhat less likely to perform waggle dances (with probability of 0.55). We do not know what gave the initial scouts their especially strong motivation to dance. Perhaps it was each initial scout's experience of finding the site by herself—not having followed another scout's dance to find it—or of inspecting the site by herself. This “discoverer-should-dance” rule is not foolproof, however. As we have already seen in the best-of-5 choice test, in which swarms were presented with a five-alternative choice (one 40-liter nest box and four 15-liter nest boxes), one swarm failed to choose the best 40-liter option because two scout bees that discovered it independently both failed to report it with dances (see fig. 5.7). Consequently, the swarm “overlooked” the excellent alternative and ended up occupying one of the mediocre ones.

There is one more important feature of scout bee behavior that caught our attention in July 2007: each of the marked scouts visited *just one* of our two nest boxes even though the two boxes were only 40 meters (130 feet) apart, a distance that a flying bee needs only 10 seconds to traverse. Such site fidelity by the scouts is noteworthy because it provides further support for Lindauer's suspicion that when a scout bee evaluates a prospective homesite, she makes an estimate of its absolute quality based on an innate (genetically specified) scale of nest-site goodness. In other words, she does not make an estimate of a site's relative quality by comparing it to other sites that she has visited. Because our swarms were prepared from colonies that had not recently swarmed, we could be sure that our bees had no prior experience as scouts before coming to Appledore Island.

And because we did not see any of them visit more than one nest box on the island, we could be confident that they did not compare one site with the other. Nevertheless, those that visited the high-quality site danced more strongly than those that visited the medium-quality site. Evidently, a worker bee possesses both an innate knowledge of what constitutes an ideal homesite and an innate ability to determine the absolute quality of the site that she has inspected. This is not a far-fetched claim; various studies of worker honeybees have shown that when a flower-naive bee searches for flowers, she spontaneously prefers objects with complex shapes, certain colors (e.g., violet rather than green), and certain odors (floral rather than nonfloral). This innate knowledge of floral cues naturally steers the novice forager's attention toward flowers.

Finally, I should emphasize that almost certainly a scout bee does not consciously think through her evaluation of a site. Instead, she probably does so unconsciously with her nervous system integrating various sensory inputs relating to cavity size, entrance height, and the like, and generating within her a sense of the site's overall goodness. It may be that finding a desirable tree cavity feels to a homeless scout bee as inherently pleasurable as feasting on a delicious meal does to a hungry human being.

The Strong Grow Stronger

One key to understanding why the scout bees' support for the best site grows and grows throughout a debate is that the supporters of the best site advertise it the most strongly. To be precise, the scout bees from the best site produce the greatest number of dance circuits per bee, on average, as we have just seen (fig. 6.6). And this is true in nature, not just in experiments. Consider again the scout bee debate depicted in figure 4.7, in which site G to the southwest prevailed, presumably because it was the best available site. Throughout the debate, the bees advertising site G produced the greatest number of dance circuits per bee. For example, between 3:00 and 5:00 p.m. on July 20, when there was fierce competition among sites A, B, D, and G, the average numbers of dance circuits produced per scout bee for these four sites were 59, 29, 42, and 74. Likewise, the next morning, between 9:00 and 11:00 a.m., when the contest had narrowed to sites




Fig. 6.7 Scout bees tune the strength of their waggles in relation to site quality, which builds a consensus of dancing bees for the best site. Here, two scouts simultaneously discover two potential nest sites, one with a large entrance opening (left) and one with a more desirable small opening (right). Each scout then returns to the swarm and performs a waggles dance for her site, but the scout from the right tree performs three times as many

B and G, the average numbers of dance circuits produced per scout bee for these two sites were 16 and 42. (Note: the bees danced only about half as strongly this morning relative to the previous afternoon because the weather had deteriorated overnight. Indeed, a rainstorm started at the end of the morning. The bees always slow their house-hunting process in cool or stormy weather.)

Because the best site stimulates its supporters to dance the most strongly, its supporters have the highest per capita success in converting neutral scouts into additional supporters. And because these additional supporters will likewise have the greatest per capita success in attracting still more supporters, the differences in number of supporters among sites that differ in quality will grow exponentially. In principle, one group of supporters will eventually overwhelm all the others, which is precisely the pattern that we have seen in the swarm bees' debates (see figs. 4.6 and 4.7).

Figure 6.7 illustrates how this works for the basic situation of two competing sites that differ in quality. The high-quality site on the right, which is more desirable by virtue of its smaller entrance opening, stimulates its supporters to advertise it with 90 dance circuits on average (as did our 40-liter nest box; see fig. 6.6). The medium-quality site on the left, which has a larger entrance opening, elicits 30 dance circuits on average from its supporters (as did our 15-liter nest box). The two sites are discovered simultaneously, each by just one scout, at 10:00 a.m. During the first three hours, the two scouts produce 90 dance circuits and 30 dance circuits, so the relative force of persuasion (total amount of advertising) for the two sites is 3:1. If we assume that 8 neutral scouts are recruited to the two sites, and in proportion to the level of advertising for each site, then by 1:00 p.m. there will be six scouts supporting the high-quality site and two supporting the medium-quality one. (By 1:00 p.m. the two original scouts

waggle dance circuits (blue symbol) as the scout from the left tree (red symbol). The result is that three hours later, the number of bees committed to the right tree has increased sixfold, whereas support for the left tree has increased only twofold, and the majority of dancing bees favor the right tree. After three more hours, the number of scouts at the right tree has ballooned, and the numerous dances in support of this site have nearly excluded the left-tree site from the debate.

will have ceased advertising and visiting the sites.) Now what happens over the next three-hour period? The six supporters of the high-quality site will produce a total of 540 dance circuits (six bees x 90 dance circuits per bee) while the two supporters of the medium-quality site will produce a total of 60 dance circuits (two bees x 30 dance circuits per bee). Thus the relative force of persuasion for the two sites becomes 9:1 during this second three-hour period. If 20 neutral scouts are recruited to these sites (more recruits now than before because there is more advertising), and if they are recruited to the two sites in proportion to the amount of advertising for each, then by 4:00 p.m. there will be 18 scouts supporting the high-quality site but still only two supporting the medium-quality site. So we can see that even though this debate started out with a 1:1 ratio of supporters for the two sites, after three hours the ratio became 3:1, and after three more hours it reached 9:1. We can also see that if the debate continues, it won't be long before the high-quality site achieves complete domination of the debate, just as in nature.

A curious feature of the bee's consensus-building process is that the domination of the debate by one site's supporters can be driven entirely by differences in the per capita strength of advertising of the various sites. One might suppose that building a consensus among the dancing bees would also require the neutral scout bees that are getting converted into supporters to pay attention to the different types of advertisements and ignore the weaker ones representing poorer sites. But in fact, the neutral scouts don't need to follow dances selectively. In the example just given, the neutral bees become supporters for the two sites strictly in proportion to the amount of dancing for the two sites. It is as if a neutral scout simply strolls across the surface of the swarm, follows the first dance that she encounters, gets recruited to the site advertised by this dance, and then becomes a supporter for this site. Although we don't know if this is exactly how a dance-following scout bee behaves, we do have evidence that they do not selectively follow dances for certain sites but instead follow dances at random.

The evidence comes from an experiment conducted by Kirk Visscher and one of our mutual friends, Scott Camazine, a gifted physician, nature photographer, and fellow honeybee fanatic. In December 1995, in the desert east of Indio, California, where large trees are rare and so natural homes for honeybees are scarce,

Kirk and Scott set up artificial swarms (one at a time) and two nest boxes. These boxes attracted the interest of the scout bees from their swarms. Kirk and Scott then labeled for individual identification each scout that performed a dance for one of the nest boxes, and they video recorded all instances of dancing and dance following throughout each swarm's decision-making process. Then they reviewed the recordings to see which of their labeled dancers eventually became dance followers. For those that did, they determined whether each bee selectively followed dances for the nest box that was not the one she had previously visited and advertised, perhaps so she could do some "comparison shopping." Remarkably, they found that the dancers that became dance followers followed dances for the two nest boxes simply in proportion to the amount of dancing for the two boxes. Thus these bees gave no sign of doing anything more sophisticated than following dances chosen at random.

We see, therefore, that the debating scout bees appear to use a simple method to build an agreement: the better the potential homesite, the stronger the dancing of the scout bees supporting it and the greater their effectiveness in recruiting additional supporters for their place. The new supporters of each spot visit and evaluate it for themselves—thereby checking the "claims" of the previous advocates of the site and avoiding untested information being spread like a rumor—and then they likewise announce it with dances, weak or strong according to their evaluations of the place. Bit by bit, because the positive feedback (the recruitment of recruiters) is strongest for the best site, the supporters for this site increasingly dominate the discussion. Complete agreement requires, however, not only that the support for the best site steadily grows, but also that the support for the poorer sites gradually fades. We will turn now to seeing how the support for the losing sites melts away.

The Expiration of Dissent

For an agreement to emerge within a group that is debating multiple options, all of the group's members who start out supporting the losing options must eventually withdraw their support for these options and either switch their support to the winning option or quit the debate altogether. In short, the dissent must

expire. We have seen that this happens in the dance debates among scout bees on honeybee swarms (see figs. 4.6 and 4.7), such that every bee that starts out dancing for a rejected site eventually ceases doing so, but we've not yet seen exactly how this occurs. Back in the early 1950s, Lindauer wrestled with this important puzzle about the bee's consensus-building process but he never quite solved it. He seemed to favor the idea that a scout bee ends her support—her dancing—for one site only when she learns about a superior site and shifts her dancing to it. He expressed this view as follows:

Scout bees that could only find lesser nest sites easily change their votes in favor of a different nest site. Even if they dance for "their" nest site at first, they decrease their dancing bit by bit, become noticeably more interested in the lively dances of the other scout bees and finally take off to seek out the other nest site. On their inspection visits they can now draw a comparison between their own and the new nest site and, if the latter is really more suitable, from now on they also dance for it on the swarm cluster. In this way all the interest of the scout bees is concentrated bit by bit on the best of all the nest sites.

There are two critical elements of this hypothesis for how scouts cease dancing for losing sites: a bee *compares* her old site with a new site (to which she was recruited by lively dances of other bees), and if she finds the new site superior she *converts* to dancing for the new and better site. Thus we can call this the compare-and-convert hypothesis for the expiration of dissent. It is certainly a plausible hypothesis. It is, after all, how we humans usually resolve disagreements in a debate; the group's members propose various courses of action, individuals hear and compare the various proposals, and eventually the individuals who initially favored a losing proposal change their minds and convert to supporting the winning proposal. I suspect that Lindauer reasoned by analogy to consensus building by humans as he worked to understand how the scout bees reach an agreement, for he described the bees as not remaining "stubborn about their first decision" and letting "their minds be changed."

Even though Lindauer stressed the compare-and-convert hypothesis to explain the expiration of dissent among nest-site scouts, he also reported some observa-

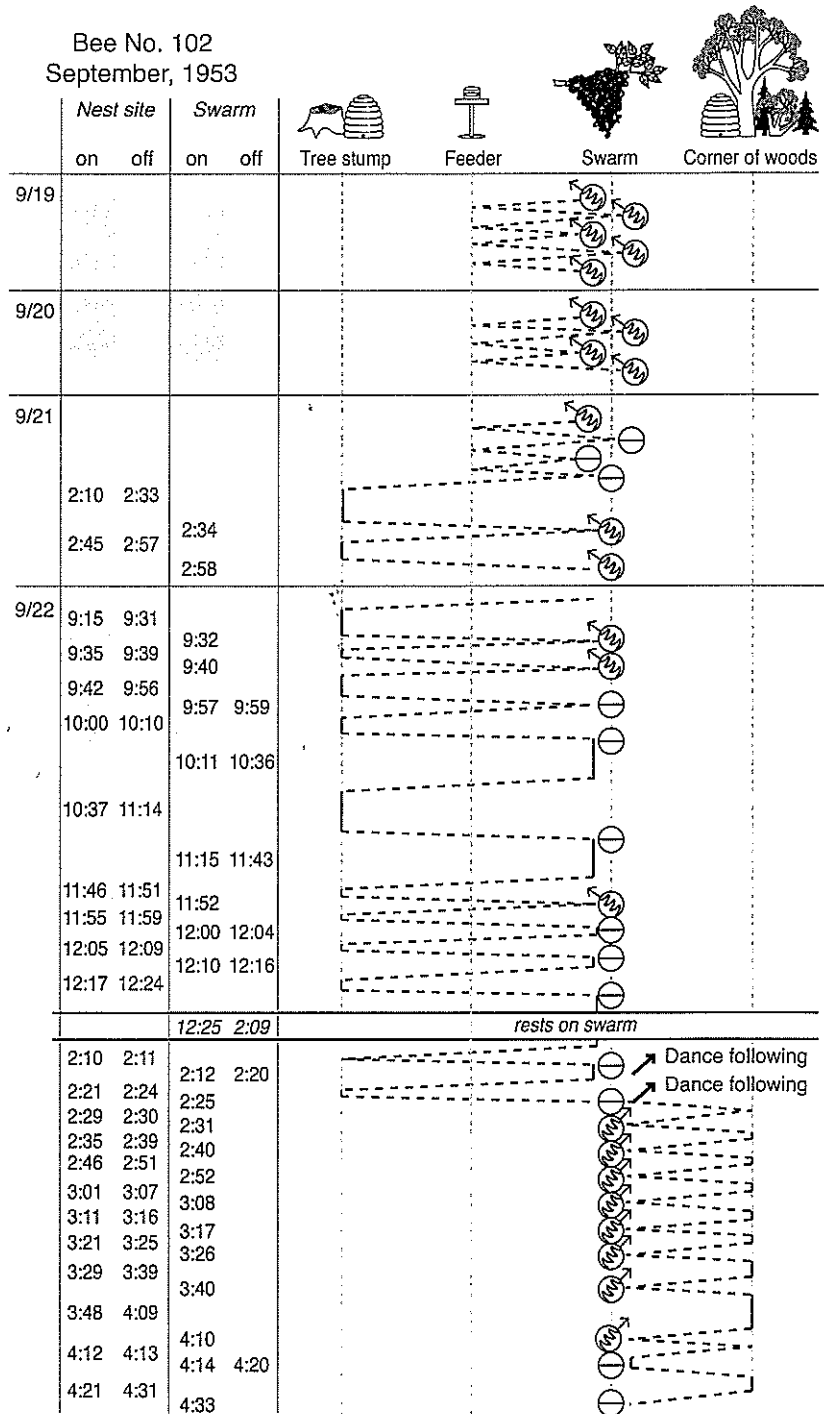
tions that weren't entirely consistent with this hypothesis. For example, he wrote, "It is still not understood why those scout bees that had found an inferior nest site gave up dancing for the site over time, even when nothing changed about their nest site and they had not yet inspected any new housing possibilities." Clearly, he had seen instances in which a scout bee quit dancing for one site even before she knew about another site, hence before she could compare her old site against a new one. Indeed, in his magnum opus of 1955 Lindauer included a beautifully detailed record of one scout who quit dancing for one site and then sat quietly on the swarm for nearly two hours before she began following dances that directed her to a second site (fig. 6.8). This shows clearly that sometimes a scout bee will quit dancing for a site without first comparing it to another site.

The two critical elements of this alternative hypothesis for how scouts cease dancing for losing sites are these: a bee *does not compare* her old site with a new site, and she *does not convert* to dancing for a new and better site. Instead, she simply loses her motivation to dance for one site and then becomes quiescent, not even visiting her site. Thus we can call this the retire-and-rest hypothesis for the expiration of dissent.

Whenever you have two competing and mutually exclusive hypotheses to explain a single mystery, you can determine which one is false by identifying some phenomenon about which the two hypotheses make clearly different predictions. You then go out, observe the critical phenomenon, and see which hypothesis doesn't correctly predict what you have observed. You know immediately that this hypothesis is false. This "strong inference" procedure may sound esoteric, but it is something we all do all the time. For example, if the light doesn't go on in a room when you turn on the light switch, you wonder if the cause is (hypothesis 1) the bulb burned out or (hypothesis 2) the power went out. If the former, then you predict that the lights will work in another room, but if the latter, then you predict that they won't. So you check the lights in another room and when you find that they work, you know immediately that the power-went-out hypothesis is bogus.

To distinguish between the compare-and-convert hypothesis and the retire-and-rest hypothesis for how the dissent among the dancing scout bees expires, I made use of the fact that these two hypotheses make distinctly different predictions about when a scout will cease dancing for a losing site relative to when she follows a dance for another site. A critical prediction of the compare-and-convert

Bee No. 102
September, 1953



hypothesis is that a scout will cease dancing for a losing site *only after* she has followed a dance for another site (and then located this site and compared it to her current site). In contrast, a critical prediction of the retire-and-rest hypothesis is that a scout will cease dancing for a losing site *even before* she has followed a dance for another site. Testing these two predictions was simply a matter of setting up swarms one at a time, labeling with bright paint marks the first few bees that performed dances on each swarm, and then observing these labeled bees steadily whenever they were at the swarm to see when they danced and when they stopped dancing, and when (if ever) they followed the dances of other bees. I focused my attention on the first few dancers to appear on each swarm because I knew from eavesdropping on the scout bees' debates that the early dancers tend to advertise losing sites.

Since I needed to be able to observe all instances of my focal bees producing or following dances, I limited myself to labeling only a few (four to eight) scout bees on each swarm. This, in turn, meant that I needed to repeat the entire observation protocol on several swarms to get data on a sufficient number of bees. The work would be slow going, but this was fine by me. I knew that it would be both pleasurable and valuable to watch steadily my small company of brightly colored scout bees on a swarm—noting for each individual all her comings and goings, and all her dance producings and dance followings—until the swarm finished choosing its new home. Times spent outdoors closely observing bees always include the thrill of discovery.

I watched 37 scout bees in six swarms, which required a total of 66 hours of steady observation. As expected, most (31, or 84 percent) of the scouts first

Fig. 6.8 Records from the life of swarm bee Number 102, who served initially as a forager, then became a scout bee that first advertised a nest site (empty hive) beside a tree stump and later switched to advertising a different nest site (empty hive) at the edge of a woods. Dashed lines show flights to or from the swarm. Solid lines denote times spent at the swarm or one of the nest sites. A circle with a wavy line indicates a dance, and the arrow indicates to which feeder or nest site the dance referred. The swarm was not well fed when it was set out, so some bees (like Number 102) foraged from a feeder at first, then they became sluggish foragers as the swarm became well fed, and finally they began scouting for nest sites.

advertised a site that was eventually rejected and only a few (six, or 16 percent) danced initially for the site that was ultimately chosen by their swarm to be its future home. Of the 31 bees that started out supporting a losing site, 27 ceased advertising their sites before the end of the swarm's decision making and the other four almost did so, for their dancing had become feeble by the time the swarm finished its decision making. The key question, then, is how did the 27 bees that quit supporting a losing site do so? Did they stop dancing only after or even before they had followed dances for other sites? Figure 6.9 shows how three of these bees behaved on one swarm that chose a site in the south for its new home. We see that the first bee, Red, stopped dancing for a losing site in the west on her second trip back to the swarm, and that she did so without first following a dance for another site. Likewise, the second bee, Pink, stopped dancing for a losing site in the southwest on her third trip back to the swarm, and she too did so without first following a dance for another site. It was not until her fourth trip back to the swarm that she followed five circuits of a dance promoting a site to the west, and so possibly learned about an alternative site. Finally, the third bee, Orange, stopped dancing for a losing site in the east on her fifth trip back to the swarm, and just like Red and Pink she did so without first following a dance for another site. So all three of these bees ceased their dancing *before* they followed a dance for another site. Their behavior was typical. Of the 27 bees of interest, 26 (96 percent) stopped dancing for their losing sites *before* they followed dances for other sites and only one (4 percent) stopped her dancing for a losing site *after* she had followed a dance for another site. The finding that only one out of the 27 bees stopped her dancing for a losing site after she had followed a dance for another site indicates that the compare-and-convert hypothesis is incorrect, at least for the vast majority of scout bees. These results also increase our confidence that the retire-and-rest hypothesis is correct.

So what caused the dancers for the losing sites to retire from advertising these sites? Clearly, most were not stimulated to do so by encountering a bee dancing extremely enthusiastically for another site, for most ceased producing dances before they followed any dances. One strong possibility is that the bees were driven to retire from advertising the losing sites by an internal, neurophysiological process that causes every scout to gradually and automatically lose her motivation

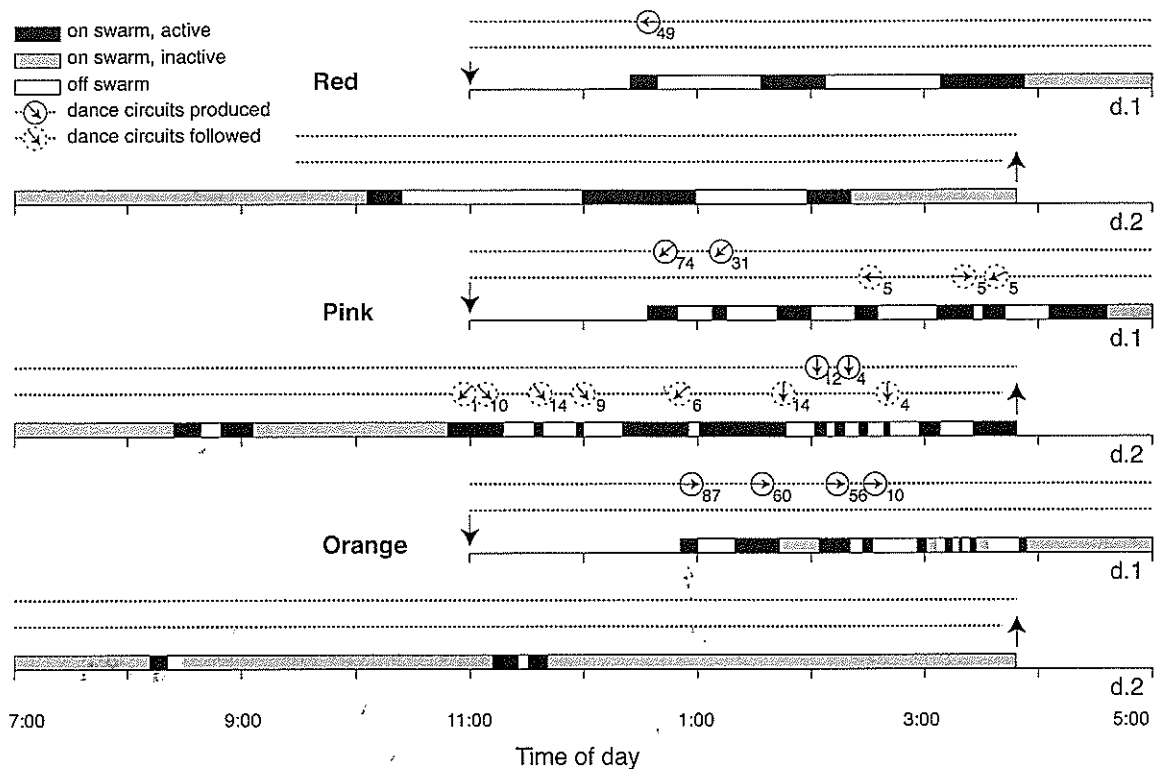


Fig. 6.9 Plots for three scout bees showing when each bee was on or off the swarm and how much dancing she produced or followed each time she was on the swarm. Each bee's history is shown for the two days over which the swarm chose its future home. The large arrows at the start and end of each bee's record denote when the swarm settled and lifted off. Each circle enclosing a small arrow denotes a dance that a bee produced or followed, and the arrow's direction indicates the compass direction of the site (an arrow pointing straight up means north, etc.). The number beside each circle enclosing an arrow shows the number of dance circuits that the bee produced or followed.

to dance for a site, even one that is high in quality. Such a process would foster consensus building among the scouts, for automatic fading of each bee's dancing would prevent the decision making from coming to a standstill with groups of unyielding dancers deadlocked over two or more sites. It might also help the dancers reach unanimity more quickly than they would otherwise, for endowing each bee

with an automatic tendency to lose interest in any given site would make each bee a highly flexible participant in the decision-making process.

One piece of evidence that strongly supports the idea that scout bees have an internally driven tendency to stop dancing for any given site is something I noticed about the 37 scout bees that I watched to test the compare-and-convert and the retire-and-rest hypotheses: each bee reduced the strength of her dancing over consecutive trips back to the swarm. For example, in figure 6.9, we see that for the bee Red the decline in dance strength (number of dance circuits per trip to the swarm) was abrupt: 49 then 0. For the bees Pink and Orange, however, the declines in dance strength were more gradual: 74, 31, then 0; and 87, 60, 56, 10, then 0. (Note: one can also see this consistent drop in dance strength in the dance records of the individual scouts shown in figure 6.5.) When I tabulated all instances, for all 37 scout bees in which a bee made a series of returns to the swarm with dancing for a particular site followed by a return without dancing, I found that the bees had produced 51 such series. They varied in length from one trip back with dancing to six consecutive trips back with dancing. Then I grouped the 51 series into six sets according to series length, and for each set I calculated the mean number of dance circuits in trip 1, in trip 2, and so forth. Finally, I compared the results for the six sets by aligning them with respect to the trip back when the scout bee did not dance, as shown in figure 6.10. This revealed that, regardless of series length, there was a regular pattern of the scouts producing fewer and fewer dance circuits across a series of trips back to the swarm, and that the rate of decline in the number of dance circuits per trip did not differ markedly between bees producing long and short series. On average, there is a remarkably regular decline in the number of dance circuits produced per trip back to the swarm, and the rate of this decline is approximately 15 fewer dance circuits per trip.

It is important to note that the same pattern of steady decay in dance strength is seen with all scout bees, both those advertising a chosen site (high in quality) and those advertising a rejected site (lower in quality). The only difference is that a bee that advertises a high-quality site tends to start her reporting by performing a large number of dance circuits, whereas one that advertises a low-quality

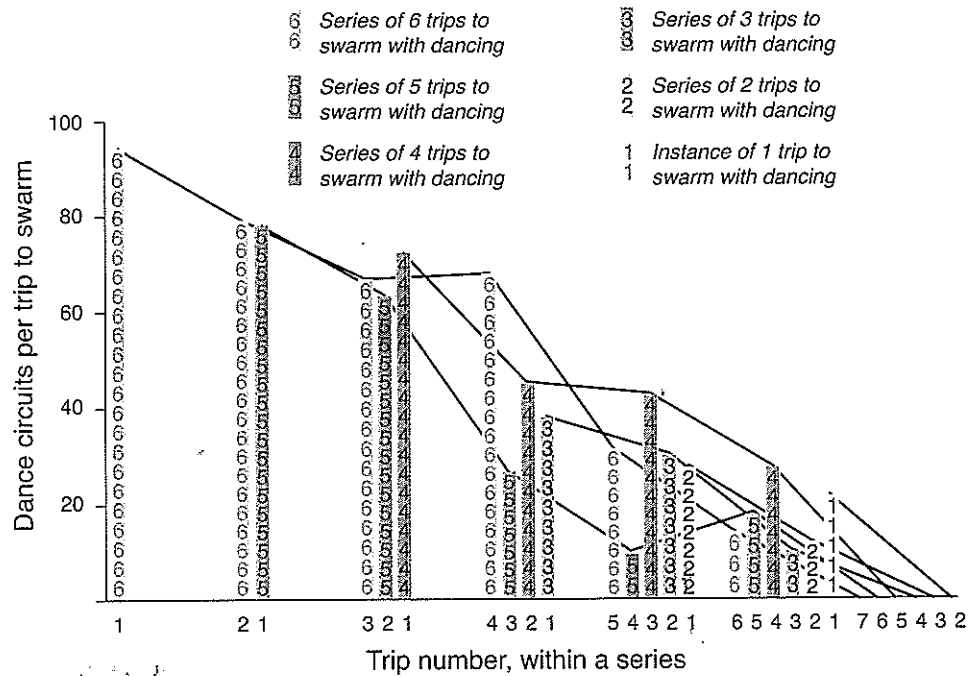
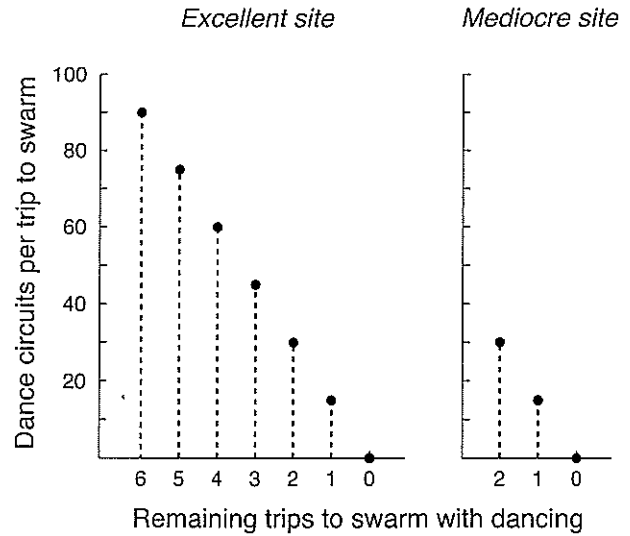


Fig. 6.10 Scout bees decrease the number of dance circuits performed for a site on successive trips back to the swarm. A string of consecutive trips back to the swarm with dancing for a particular site is called a “series”; series vary in length from six trips to one trip. The drop in dance strength per trip (about 15 dance circuits) appears to be a constant, regardless of series length.

site will tend to start her reporting with a smaller number of dance circuits (see fig. 6.5). Because the rate of decay in dance strength per trip back to the swarm is the same for all scouts, a bee from a high-quality site will tend to advertise her site over many consecutive trips back to the swarm (for example, the bee Orange in figure 6.5) and in sum will produce a strong advertisement with many dance circuits, whereas a bee from a medium-quality site will tend to advertise for only a few consecutive trips back to the swarm (for example, the bee Blue-White in figure 6.5) and in sum will produce a weaker advertisement with fewer dance circuits. Consequently, as shown in figure 6.11, a scout bee supporting a superb site, relative to one supporting a poorer site, will be both a longer and “louder”

Fig 6.11 Comparison of the patterns of dancing by scout bees advertising an excellent site or a mediocre site. Both bees reduce their dance strength at the same rate (15 fewer dance circuits per trip back to the swarm), but the bee from the excellent site starts with higher motivation to dance and so dances longer (6 trips vs. 2 trips) and “louder” (90 + 75 + 60 + 45 + 30 + 15 = 315 dance circuits vs. 30 + 15 = 45 dance circuits).



supporter of her site. And as we all know, in any contest for popular support, the side with the most persistent and most zealous supporters is the one most likely to prevail.

It appears, therefore, that a swarm’s scout bees do something sharply different from what humans do to reach a full agreement in a debate. Both bees and humans need a group’s members to avoid stubbornly supporting their first view, but whereas we humans will usually (and sensibly) give up on a position only after we have learned of a better one, the bees will stop supporting a position automatically. As is shown in figure 6.5 and figure 6.9, after a shorter or longer time, each scout bee becomes silent and leaves the rest of the debate to a new set of bees. Figure 6.7 shows how this regular turnover in which scouts are dancing can help a swarm’s scouts quickly reach an agreement, for in this schematic depiction of consensus-building on a swarm all of the bees that were active dancers at 10:00 a.m. have retired by 1:00 p.m., and all those that were active dancers at 1:00 p.m. have retired by 4:00 p.m.

There is, however, one important case in which human group decision making operates in a manner similar to that of honeybee swarm house-hunting. It is how

scientists conduct their social decision making on scientific theories. Many have noted that new and better ideas succeed in scientific debates through attrition, that is, by one generation of scientists retiring from their field and eventually dying off. But before this generation drops from the debate, the next generation of scientists will have listened carefully to the various arguments made by their predecessors, been persuaded by the most compelling claims on the truth, and adopted the new theory. Thus the support for a new and better theory (e.g., the sun-centered theory of Copernicus and Galileo) grows while it fades for an older and poorer one (e.g., the earth-centered theory of Ptolemy). The most often quoted statement describing this social process is by Max Planck: "A new scientific truth does not triumph by convincing its opponents and making them see the light, but rather because its opponents eventually die, and a new generation grows up that is familiar with it." One difference between aged scientists and aged scouts, though, is that the people tend to drop out of the debate reluctantly, sometimes not until death, whereas the bees do so automatically. I cannot help but wonder whether science would progress more rapidly if, in this regard, people behaved a bit more like bees.