Learning in honeybees as a function of amount of reward: Further experiments with color

P. A. COUVILLON and M. E. BITTERMAN
University of Hawaii, Honolulu, Hawaii

Foraging honeybees were trained individually with successively presented targets differing in color, one containing 5 μl and the other 20 μl of 50% sucrose solution, after which preferences were measured in unrewarded choice tests. The targets were conical, designed to control for the possibility of differential delay of reward stemming from the greater detectability of the larger as compared with the smaller drops of sucrose when the drops were presented on the conventional flat targets. The new results for color, like recent results for odor, can be understood on the assumption that the attractiveness of a stimulus increases as a function of the strength of its association with reward and that the effect of amount of reward is on asymptotic strength.

Foraging honeybees trained with two successively presented targets different in color or odor, one always containing a 5-μl drop of sucrose solution and the other a 20-μl drop of the same solution, readily develop a preference for the 20-μl target (Buchanan & Bitterman, 1988, 1989; Couvillon, Lee, & Bitterman, 1991). The results for both color and odor can be understood on the simple nonrepresentational assumption that the attractiveness of a stimulus depends on the strength of its association with reward and that the strength of association depends on the amount as well as the frequency of the reward. In another respect, however, the results are different, those for color suggesting that the effect of the amount of the reward is on the rate of acquisition (the equal-asymptote assumption), and those for odor that the effect is on asymptotic strength.

The information about color comes from a series of three experiments by Buchanan and Bitterman (1989). In the first—a reversal experiment—one of two targets contained 5 μl of sucrose solution on Trials 1-16 and 20 μl on Trials 17-32, while the opposite was true of the other. On the equal-asymptote assumption, the terminal associative strengths of the two targets should have been exactly the same, and, in fact, no differential response to them was found in a subsequent choice test. In the second experiment, 20 differentially rewarded trials with each of two targets—enough to produce a clear preference for the 20-μl target at the outset of training—were given after 10 feedings-to-repletion on each target that were calculated to bring their associative strengths close to the hypothetical common asymptote, and, in fact, no differential response to them was found in a choice test. In the third experiment, there were 12 feedings-to-repletion on one

of two targets, while on the other there were 12 feedings-to-repletion followed by fifteen 5-μl trials that should not, on the equal-asymptote assumption, have made the animals any less likely to approach it, and, in fact, no differential response to the targets was found in a choice test. In two odor experiments by Couvillon, Lee, and Bitterman (1991), the reversal strategy was used again. In one of them, there were sixteen 5-μl trials followed by sixteen 20-μl trials with one of two targets and sixteen 20-μl trials followed by sixteen 5-μl trials with the other. The second experiment was like the first, except that there were twenty-four 5-μl and twenty-four 20-μl trials with each target. In each case, a choice test showed a clear preference for the target that had more recently provided the larger amount, clearly contradicting the equal-asymptote assumption.

How is the discrepancy to be explained? Lee and Bitterman (1990b) suggested that the critical variable in the work with color was the delay, rather than the amount of the reward. Given that inexperienced honeybees locate 20-μl drops more readily than 5-μl drops (Walker, Lee, & Bitterman, 1990), an association between color and sucrose may have developed more rapidly with the larger drops than with the smaller because of closer contiguity between the perception of the target's color and the initial taste of the sucrose, a common asymptote being approached as the animals learned to locate the smaller drops more quickly. In the odor experiments, the possibility of differential delay was ruled out by marking the location of each drop with a readily discriminable white dot, but dots could not be used in the color experiments because they seemed to overshadow the colors (Lee & Bitterman, 1990b). On the assumption that the associative strength acquired by a stimulus paired with sucrose depends on the duration of concurrent exposure, which is longer for larger rewards because the ingestion time is longer (Buchanan & Bitterman, 1989), the amount of reward per se actually may have played no role at all in the color experiments, because stationary honeybees, which readily
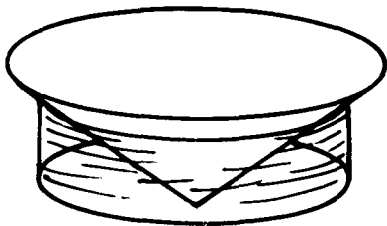
Figure 1. Sketch of a conical target. The widest diameter of the target is 6.2 cm.

process odors, do not seem to process visual stimuli very well under the same conditions (Walker, Baird, & Bitterman, 1989).

In the present experiments, the color–amount relation was studied again, but now with a new technique for equating the delay of the reward. Instead of the flat circles of colored plastic used previously, the targets used here were conical in structure, each set on a base of clear Plexiglas as diagrammed in Figure 1. Because the animal landed at the top of the target and walked down its inside slope to the bottom, where the drop of sucrose was placed, a small drop was encountered as quickly as a large one. If the results of the previous color experiments with flat targets were due only to a difference in the delay of the reward confounded with the difference in the amount, and if the conical targets equate the delay, like experiments with the conical targets should show no effect of the amount.

EXPERIMENT 1

In the first experiment, the animals of a single group were trained as before with successively presented targets, those of one color always containing 5 μl and those of the other containing 20 μl of a 50% sucrose solution. The design called for 16 rewarded trials with each color separately and a terminal unrewarded choice test with the two colors. It may be well to explain again here why the amount of reward was varied within subjects rather than between subjects in these experiments. Since the animals take about 50 μl of sucrose on each visit before leaving of their own accord for the hive, between-subjects variation in the amount of reward would mean only two or three trials per visit for animals trained only with 20 μl but about 10 trials per visit for animals trained only with 5 μl; that is, the amount of reward would necessarily be confounded with either the number of trials or the number of visits. Animals trained with both amounts average two 20-μl trials and two 5-μl trials on each visit.

Method

Subjects. The subjects were 15 honeybees (*Apis mellifera*) from our own hives situated near the laboratory. All were experimentally naive.

Procedure. The training situation was the same as that used in the odor experiments (Couvillon, Lee, & Bitterman, 1991). It consisted of two immediately adjacent windows (each 55 cm wide and

55 cm high) separated by a thin (2-cm) wooden partition around which the animals were required to fly from one window to the other. In the pretraining, experience with both windows was given. A single forager was selected at random from a group of foragers at a feeding station providing 10%–15% sucrose solution, carried to the laboratory, and set down at a large drop (>100 μl) of 50% sucrose solution in a conical pretraining target that was centered on the sill of one of the two windows (the left for half the subjects and the right for the rest). The animal was marked with a spot of colored lacquer as it fed to repletion, after which it was permitted to leave for the hive. Typically, the animal (adapted to 50% sucrose and now finding the lower concentration at the feeder unacceptable) would come back to the laboratory after a few minutes, continuing to shuttle back and forth from the hive as long as sucrose was available there. If the marked animal did not come back after its first placement, it was carried again from the feeding station, where it usually could be found, to the pretraining target. When the animal did return to the first window, the target was picked up after a few seconds and, as the animal continued to feed, centered on the sill of the alternative window, from which the animal left for the hive when replete. On subsequent visits, a pretraining target was presented twice at the second window and once again at the first. The pretraining ended after the animal had returned twice to each window of its own accord.

In all, there were three sets of targets. Those of one set were of orange plastic, those of the second were of blue plastic, and those of the third (used only in the pretraining) were half orange and half blue. The colored plastics were the same as those used in previous color–amount experiments. The targets used on each visit were washed and exchanged for others in their sets after the visit in order to randomize extraneous stimuli.

Arriving from the hive on each training visit, an animal found either an orange or a blue target centered on the sill of one of the windows, the left window on half the visits and the right window on the rest, in quasirandom order. When the animal landed on the target and made contact with the reward, a second target—orange or blue, in quasirandom order—was centered on the sill of the adjoining window, to which the animal would fly after taking the sucrose on the first target. Then the first target was removed, another new target was centered on the sill of the arrival window, and so forth, until the animal was replete and returned of its own accord to the hive. Although the animal itself determined the number of trials on each visit, a running record kept by the experimenter made it possible to control to a considerable extent the number of 20-μl and 5-μl presentations. The actual mean numbers of trials were 16.1 with the 20-μl color and 15.7 with the 5-μl color. The 20-μl color was orange and the 5-μl color was blue for 7 of the 15 subjects, while the opposite was true for the rest.

After its last training visit, each animal returned from the hive to find a pair of fresh targets, one orange and the other blue, set 10 cm apart in a lateral arrangement on the sill of one of the windows, the left window for 7 of the 15 animals and the right window for the rest. For 8 animals, the target with the 20-μl color was to the left of the other, and for the remaining animals to the right, each target now containing a 10-μl drop of tap water (unacceptable, and distinguishable from the sucrose solution only by taste). Upon encountering water on one of the targets, the animal would leave it, then return to it or go to the other, leave again, return again (often only briefly, with no attempt to drink), and so forth, the interval between successive responses increasing as the test continued. All actual contacts with the inside surface of each target, however brief, during a 10-min period were recorded by the experimenter, who pressed one of two hand-held switches that activated counters programmed to print stored frequencies at 30-sec intervals.

Results

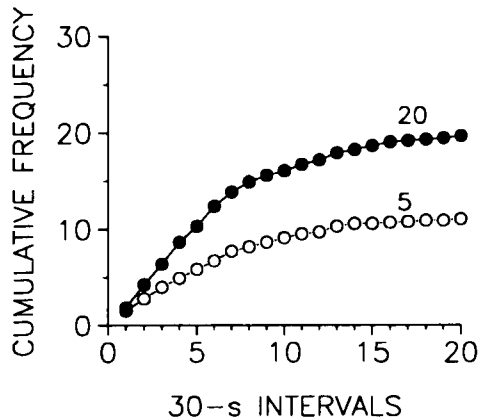In Figure 2, performance in the choice test is plotted in terms of the mean cumulative number of responses to

Figure 2. Mean cumulative number of responses to the 5- and 20-μl colors in the choice test of Experiment 1.

each color in successive 30-sec intervals. As the curves show, there was a clear preference for the 20-μl color. Analysis of variance—based in accordance with our standard practice on uncumulated 30-sec scores in 2.5-min blocks—yielded a significant stimulus (20 vs. 5 μl) effect [$F(1,14) = 14.44$, $p = .0020$], a significant 2.5-min block effect [$F(3,42) = 31.43$, $p < .0001$], and a significant stimulus × block interaction [$F(3,42) = 4.96$, $p < .0050$]. The preference for the 20-μl color points to the effectiveness of amount of reward apart from delay of reward on the plausible assumption that the conical structure of the targets rules out the possibility of differential delay.

EXPERIMENT 2

One possible explanation of the results of Experiment 1 is that the conical structure of the targets serves not only to equate delay but to enhance the processing of color (which now surrounds the animal) during feeding. If so, experiments with conical colored targets and with odor targets dotted to ensure equal delay should give the same results, which is to say that both should show that the asymptote of the acquisition function varies with amount of reward.

Of the various tests of the equal-asymptote assumption employed by Buchanan and Bitterman (1989), the reversal test is perhaps the most powerful. The procedure is to compare responses to two stimuli after one of them has been paired with large reward followed by small reward and the other with small reward followed by large reward. With the growth of associative strength described as

$$\Delta V = \beta(\lambda - V)$$

(where V is the preexisting strength; ΔV, the increment in strength on each trial; β, the growth rate; and λ, the asymptotic strength), it can be shown in general that if two such acquisition functions approach a common

asymptote (for convenience, taken as 1) at different rates (β_1 and β_2), then

$$V_{m+n} = V_{n+m} = 1 - (1 - V_0)(1 - \beta_1)^m(1 - \beta_2)^n$$

(where V_0 is the initial associative strength of each stimulus; V_{m+n} and V_{n+m} are their terminal associative strengths; and m and n are the numbers of trials with the two amounts of reward). The argument, to be sure, is one-sided, because the lack of a preference does not establish that the asymptotes are the same, but a clear preference for either stimulus demonstrates unequivocally that the asymptotes are different. In the present experiment, reversal was studied with the new conical targets instead of the old flat ones.

Method

Subjects. The subjects were 15 experimentally naive foragers from our own hives.

Procedure. The procedure was the same as in Experiment 1, except that 48 trials with each color were planned—twenty-four 5-μl trials followed by twenty-four 20-μl trials with one color (the 5-20 color) and twenty-four 20-μl trials followed by twenty-four 5-μl trials with the other (the 20-5 color). The actual mean numbers of prereversal trials were 24.9 for the 5-20 color and 25.0 for the 20-5 color; the mean numbers of postreversal trials were 24.3 for the 5-20 color and 24.1 for the 20-5 color. The unrewarded choice test given on the final visit was exactly the same as before, with windows and target positions balanced over animals.

Results

In Figure 3, performance in the choice test is plotted in terms of the mean cumulative number of responses to the 5-20 and 20-5 colors in successive 30-sec intervals. The curves show a clear preference for the 5-20 color, which contradicts the equal-asymptote assumption. Analysis of variance yielded a significant stimulus effect [$F(1,14) = 7.77$, $p = .0145$], a significant 2.5-min block effect [$F(3,42) = 8.06$, $p = .0002$], and a significant stimulus × block interaction [$F(3,42) = 3.59$, $p = .0212$].
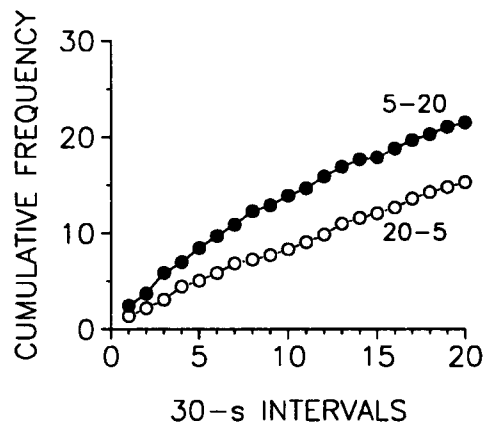

Figure 3. Mean cumulative number of responses to the 5-20 and 20-5 colors in the choice test of Experiment 2.

EXPERIMENT 3

A noteworthy feature of the technique employed in these and preceding experiments in the series is that the effects of training with different amounts of reward are measured in unrewarded choice tests that follow the training. With no index of the changing attractiveness of the stimuli during the training, inferences about the underlying acquisition functions are more indirect than they otherwise might be. In hope of finding such an index, Lee and Bitterman (1990a) measured latency of response in the two-window situation with undotted flat targets differing in color. (Latency of response was defined as the time between leaving a target in one window and landing on a target in the other.) What they found was an orderly decline in latency as training continued, but no tendency for the 20-μl color to be approached more rapidly than the 5-μl color. More recently, however, asymptotic latency has been reported to vary with sucrose concentration (50% vs. 20%) in the same two-window situation (Loo & Bitterman, 1992). Those results made it seem reasonable to look further for an effect of amount on latency by using the conical colored targets, which, unlike the flat targets of Lee and Bitterman (1990a), do show an asymptotic effect of amount on choice.

Method

Subjects. The subjects were 8 foraging honeybees, all experimentally naive, from our own hives.

Procedure. The training procedure was the same as it was in Experiments 1 and 2, except that latency was measured. The experimenter depressed a foot pedal when the animal left the first target presented on each visit, depressed it again when the animal made contact with the second target, again when the animal left the second target, again when it made contact with the third target, and so forth, until the replete animal returned to the hive. A computer measured the times between switch closures, which together with the actual sequence of targets presented on each visit yielded the latency of response on each trial with each color. (The latency of
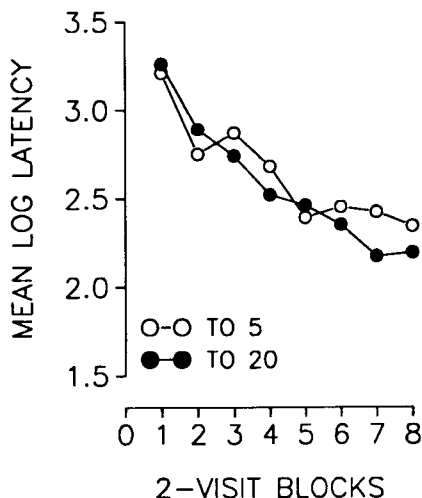

Figure 4. Mean natural log latency of response to the 5- and 20-μl colors in Experiment 3.
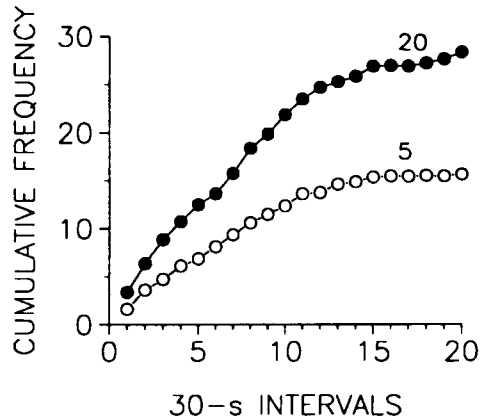

Figure 5. Mean cumulative number of responses to the 5- and 20-μl colors in Experiment 3.

response to the first target presented on each visit could not be measured, because of the lack of a definite starting point on the animal's return from the hive.) There were 16 training visits during which the number of trials (determined by the animals themselves) was somewhat larger than usual; the means were 42.8 for the 20-μl trials and 43.1 for the 5-μl trials. On the final visit, there was an unrewarded choice test that was exactly the same as that in Experiments 1 and 2, with windows and target positions balanced over subjects.

Results

In Figure 4, the performance of the animals in the 16 training visits is plotted in terms of the mean natural log latency of response to the 5- and 20-μl targets in two-visit blocks. The curves show the marked decline previously reported for flat targets differing in color (Lee & Bitterman, 1990a), although the decline is somewhat less precipitous. In the previous work, the latency of response to the 20-μl color was the same as the latency of response to the 5-μl color—the *prospective* latencies did not differ—but there was a difference in the *retrospective* latencies; that is, there was a tendency on each visit to respond a bit less rapidly after the first 20-μl trial than after 5-μl trials (a postingestive effect). A factorial analysis of variance in the data of the present experiment showed a significant effect of two-trial blocks [$F(7,49) = 15.94$, $p < .0001$], but there was no significant prospective effect [$F(1,7) = 1.09$, $p = .3314$] or retrospective effect [$F(1,7) < 1$], nor was there any significant interaction [for the prospective × retrospective interaction, $F(1,7) = 3.31$, $p = .1115$; for the prospective × block interaction, $F(7,49) = 1.57$, $p = .1671$; and for the retrospective × block interaction $F(7,47) < 1$].

In Figure 5, performance in the choice test is plotted in terms of the mean cumulative number of responses to each color in successive 30-sec intervals. As the curves show, there was a clear preference for the 20-μl color, to which the animals made a mean of 28.9 responses as compared with 15.9 responses to the 5-μl color. Analysis of variance yielded a significant stimulus (20 vs. 5 μl) effect [$F(1,7) = 8.91$, $p = .0204$], a significant 2.5-min

block effect [$F(3,21) = 16.96$, $p < .0001$], and a significant stimulus × block interaction [$F(3,21) = 3.15$, $p = .0465$]. Here again, as in the earlier work of Lee and Bitterman (1990a), terminal choice was sensitive to the amount of reward, but the latency of response during acquisition was not.

DISCUSSION

With the delay of reward equated, the role of the amount of reward in the learning of honeybees seems to be the same for color as for odor. The results for both color and odor can be understood on the parsimonious nonrepresentational assumption that the attractiveness of a stimulus depends on the strength of its association with reward and that the effect of the amount of reward is on asymptotic strength rather than on the rate of acquisition. The generalization holds also for the results of color and odor experiments of the same design with concentration of sucrose, in which the course of acquisition could be plotted in terms of the latency of response (Loo & Bitterman, 1992). Although latency has not been found to vary with amount, the confidence with which the equal-asymptote assumption can be rejected is no less for amount than for concentration. The critical evidence in both cases consists of differential performance in reversal experiments, whether the measure is latency, choice, or (as in the work on concentration) both. In runway experiments with rats, which typically show what seem to be asymptotic effects of the amount of reward, differences in performance sometimes disappear with prolonged training (see, e.g., Campbell, Batsche, & Batsche, 1972), but in none of those cases has the possibility of a performance ceiling been ruled out. For rats, as for honeybees, the definitive evidence is provided by reversal experiments that show better performance after large reward following small reward than after small reward following large reward (E. D. Capaldi, 1970; E. J. Capaldi & Lynch, 1967).

Two different interpretations of the effect of reward magnitude on asymptotic performance have been considered in the vertebrate literature and suggest themselves for consideration again here. One way to explain why a color paired with 20 μl of sucrose is preferred to a color paired with 5 μl of sucrose would be to assume that colors and amounts are associated (the two associations differing both in their first terms and in their second terms, or significates) and that choice is based on remembered or "represented" amounts (Hull, 1952; Tolman, 1932). Because preference is a function also of frequency of experience with the alternatives—a more frequently encountered 5-μl color may be preferred to a less frequently encountered 20-μl color (Buchanan & Bitterman, 1988)—choice must be assumed in this view to depend not only on *what* is represented (the *substance* of the representation), but in some sense also on the strength of the representation. As has been noted before (Couvillon, Lee, & Bitterman, 1991), a rigorous representational account

would therefore require some further assumptions about the growth of associative strength with experience and about the interaction of strength and substance in the determination of choice. A more parsimonious, nonrepresentational way of dealing with the results would be to assume that the two associations have a common significate (say, sweet taste) and that choice is based on associative strength, which varies with both the frequency and the amount of reward (Hull, 1943). The data for honeybees as yet provide no compelling reason to reject the simpler view.

To understand fully the role of the amount of the reward in the learning of honeybees, it will be helpful to know how differences in the amount of the reward are detected. There are several possibilities to be considered. One is that drops of different volume are differentially rewarding by virtue, if not of their visual appearance, then of some other physical property that is detected—as the concentration of sucrose is detected—immediately upon contact of the proboscis (the immediate-evaluation hypothesis). A second possibility is that duration of concurrent color–taste or odor–taste stimulation, which is greater for a larger drop than for a smaller because the time required for ingestion is greater, is the critical variable (the concurrent-stimulation hypothesis). Simultaneous odor–taste association has recently been demonstrated in harnessed honeybees (Batson, Hoban, & Bitterman, 1992), although the technique employed gives no evidence even of forward conditioning with visual stimuli. A third possibility is that ingestive aftereffects are associated with target properties that are perceived as the animal moves away from the feeding locus. Recent experiments show that foragers learn the color of nearby landmarks on departure from a feeding place (Couvillon, Leiato, & Bitterman, 1991; Gould, 1988) as well as the color of the feeding place itself (Lehrer, 1991). Conical targets, which surround the animal with color, may be more conducive than flat targets to color learning on departure.

REFERENCES

BATSON, J. D., HOBAN, J. S., & BITTERMAN, M. E. (1992). Simultaneous conditioning in honeybees (*Apis mellifera*). *Journal of Comparative Psychology*, **106**, 114-119.

BUCHANAN, G. M., & BITTERMAN, M. E. (1988). Learning in honeybees as a function of amount and frequency of reward. *Animal Learning & Behavior*, **16**, 247-255.

BUCHANAN, G. M., & BITTERMAN, M. E. (1989). Learning in honeybees as a function of amount of reward: Tests of the equal-asymptote assumption. *Animal Learning & Behavior*, **17**, 475-480.

CAMPBELL, P. E., BATSCHE, C. J., & BATSCHE, G. M. (1972). Spaced-trials reward magnitude effects in the rat: Single versus multiple food pellets. *Journal of Comparative & Physiological Psychology*, **81**, 360-364.

CAPALDI, E. D. (1970). Effect of initial reinforcer magnitude on subsequent resistance to extinction. *Journal of Experimental Psychology*, **86**, 283-287.

CAPALDI, E. J., & LYNCH, D. (1967). Repeated shifts in reward magnitude: Evidence in favor of an associational and absolute (noncontextual) interpretation. *Journal of Experimental Psychology*, **75**, 226-235.

COUVILLON, P. A., LEE, Y., & BITTERMAN, M. E. (1991). Learning in honeybees as a function of amount of reward: Rejection of the equal-asymptote assumption. *Animal Learning & Behavior*, **19**, 381-387.

COUVILLON, P. A., LEIATO, T. G., & BITTERMAN, M. E. (1991). Learning by honeybees (*Apis mellifera*) on arrival at and departure from a feeding place. *Journal of Comparative Psychology*, **105**, 177-184.

GOULD, J. L. (1988). Timing of landmark learning by honey bees. *Journal of Insect Behavior*, **1**, 373-378.

HULL, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.

HULL, C. L. (1952). *A behavior system*. New Haven, CT: Yale University Press.

LEE, Y., & BITTERMAN, M. E. (1990a). Learning in honeybees (*Apis mellifera*) as a function of amount of reward: Acquisition measures. *Journal of Comparative Psychology*, **104**, 152-158.

LEE, Y., & BITTERMAN, M. E. (1990b). Learning in honeybees as a function of amount of reward: Control of delay. *Animal Learning & Behavior*, **18**, 377-386.

LEHRER, M. (1991). Bees which turn back and look. *Naturwissenschaften*, **78**, 274-276.

LOO, S. K., & BITTERMAN, M. E. (1992). Learning in honeybees (*Apis mellifera*) as a function of sucrose concentration. *Journal of Comparative Psychology*, **106**, 29-36.

TOLMAN, E. C. (1932). *Purposive behavior in animals and men*. New York: Century.

WALKER, M. M., BAIRD, D. L., & BITTERMAN, M. E. (1989). Failure of stationary but not of flying honeybees (*Apis mellifera*) to respond to magnetic field stimuli. *Journal of Comparative Psychology*, **103**, 62-69.

WALKER, M. M., LEE, Y., & BITTERMAN, M. E. (1990). Transfer along a continuum in the discriminative learning of honeybees (*Apis mellifera*). *Journal of Comparative Psychology*, **104**, 66-70.