

Choice in honeybees as a function of the probability of reward

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

Honeybees foraging for sucrose at a laboratory window were trained in a series of ten 100-trial problems to choose between two targets differing in odor, one of them providing 10 μ l of a 50% sucrose solution and the other 10 μ l of water. In 9 of the problems, two odors were used, and the reward ratio was varied systematically over a wide range. In the 10th problem, three odors were used in an ambiguous-cue (A +/B -, B +/C -) design. The results were predicted quantitatively, and with substantial accuracy, from a simple theory of learning and choice developed in previous work on simultaneous discrimination in honeybees.

In a series of recent experiments on discriminative learning in honeybees (see Couvillon & Bitterman, 1991, for a review), individual foragers were pretrained to come from the hive to the sill of an open laboratory window. On each visit to the sill, they chose between two targets different in color, odor, or both, one of the targets containing a large drop of 50% sucrose solution from which feeding to repletion was permitted, and the other target containing a large drop of tap water (unacceptable to the animals and distinguishable from the sucrose only by taste). If an animal went first to the target containing water, it was free at once to correct its choice. Performance in each problem was measured in terms of the proportion of animals choosing correctly on each visit.

The central question in experiments of this sort is how choice is determined by previous experience with the alternatives. To answer it, we must know how the attractiveness of each target changes with experience (i.e., we must discover the rule or rules of learning), and we must know how attractiveness influences performance (i.e., we must discover the rule of choice). Early in the course of the work with honeybees (Couvillon & Bitterman, 1985), it began to appear that the rules both of learning and of choice might prove to be rather simple, but nonetheless capable (when expressed in equational form) of providing precise, quantitative accounts of seemingly complex performance. The impression was confirmed in experiments on phenomena such as reversal learning, dimensional shifting, compound-component discrimination, and conditional discrimination that served to shape the theory and refine the estimates of its parameters (Couvillon & Bitterman, 1986, 1987, 1988).

The first predictions resulting from the theory were tested in a series of overshadowing experiments designed

to challenge its parsimonious (but currently unfashionable) independence rule, according to which the components of a compound stimulus gain and lose associative strength independently with reinforcement and nonreinforcement of the compound (Couvillon & Bitterman, 1989). Honeybees were trained to discriminate between color-odor compounds, either in confounded problems (e.g., green-geraniol positive vs. blue-peppermint negative on all trials) or in control problems with the colors relevant and the odors irrelevant (e.g., green-geraniol positive vs. blue-peppermint negative on some trials and green-peppermint positive vs. blue-geraniol negative on the rest), after which there were transfer tests (both reversal and nonreversal) with unscented green and blue targets differentially rewarded. The prediction from the theory was that more unrewarded responding to the negative color would occur during training on the control than on the confounded problem, and that its associative strength at the end of the training would therefore be less, but that the small difference would not be evident in the transfer tests. In fact, it was not. Subsequent experiments with a more sensitive (extinction) measure did show the associative strength of the negative color to be less after control training than after confounded training, and symmetrical results were obtained for the odors after control training in which the odors were relevant and the colors irrelevant. In none of the overshadowing experiments was there anything incompatible with the independence assumption.

Another unusual feature of the theory is its choice function, which is plotted in Figure 1 along with the only two choice functions—matching and maximizing—that have ever been given serious consideration in work with vertebrates (Mackintosh, 1983). According to the maximizing rule, which was proposed by Hull (1930), the alternative with the greatest associative strength is chosen. According to the matching rule, which was proposed by Estes (1959), responses are distributed between the alternatives in direct proportion to their relative associative strengths. There is no compelling reason, of course, to expect that the choice function will prove to be the same for all ani-

This research was supported by Grant BNS-9010609 from the National Science Foundation and by Research Centers for Minority Institutions Grant RR03061 from the National Institutes of Health. Correspondence concerning this article should be addressed to M. E. Bitterman, Békésy Laboratory of Neurobiology, 1993 East-West Road, Honolulu, HI 96822.

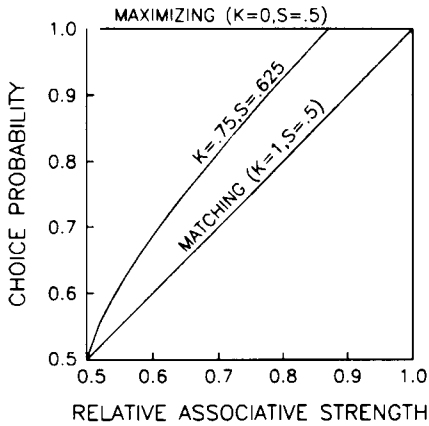


Figure 1. Three functions relating probability of choice to relative associative strength—maximizing ($K = 0, s = .5$), matching ($K = 1, s = .5$), and the function of best fit to the data of all previous choice experiments with honeybees ($K = .75, s = .625$).

mals, and comparative work suggests, in fact, that it may not be (Bitterman, 1975). For honeybees, in any case, the results of previous experiments point neither to matching nor to maximizing, but to a function in some sense intermediate between them. The present experiments, patterned after discrete-trials experiments that have figured prominently in the analysis of choice in vertebrates (Mackintosh, 1974), were designed to inquire further into the nature of the function, which was expected to be strongly constrained by data on probability of reward and change in probability of reward.

THEORY

Each training target is conceived as a compound stimulus—a color-odor compound. The attractiveness of each component of the compound is given by the strength (V) of its association with sucrose. The value of V (scaled from 0 to 1) increases with reward and decreases with nonreward. Change in associative strength on each training trial is described by the linear equation of Bush and Mosteller (1951),

$$\Delta V_A = \alpha_A \beta (\lambda - V_A),$$

in the familiar notation of Rescorla and Wagner (1972): V_A is the associative strength of component A at the start of the trial; ΔV_A is the change in associative strength produced by reward or nonreward; λ is the asymptotic associative strength, which is taken as 1 for reward and as 0 when there is no reward; α_A (scaled from 0 to 1) is the salience of A , which is taken as 1 for each of the strong and highly discriminable colors and odors employed; β (scaled from 0 to 1) is the learning-rate parameter, smaller on rewarded trials than on nonrewarded trials, but relatively large in both cases— $U\beta$ (the incremental, or up, β) = .20, and $D\beta$ (the decremental, or down, β) = .35. The larger $D\beta$ probably reflects the fact that repetitive

errors are permitted and often do occur—that is, an animal may make several unrewarded responses to the target containing water on a given trial before shifting to the target containing sucrose, although only initial errors are recorded and simulated. When the simulated response to a compound of which A is a component is rewarded, $\Delta V_A = .20 (1 - V_A)$; when the response is not rewarded, $\Delta V_A = .35 (0 - V_A) = -.35 V_A$. When the simulated initial choice is incorrect, the V of each component of the nonrewarded compound is decremented, and then the V of each component of the rewarded compound is incremented because the correction method insures that each trial ends with reward. (No inhibitory process is assumed, nor is there any opportunity for the value of V to become negative.)

It should be emphasized that the Rescorla-Wagner notation is used only because of its familiarity. The color and the odor of each target are, as already noted, explicitly assumed to gain and lose associative strength independently in consequence of rewarded and nonrewarded experience with the target. There is, however, some good evidence of the discrimination of color-odor compounds as compounds (Couvillon & Bitterman, 1982, 1988), which is explained in the manner of Rescorla (1972) and of Whitlow and Wagner (1972) on the purely perceptual assumption that new, compound-unique components (CU) are generated by afferent interactions. A target of odor A and color X , for example, is treated as a compound of A, X , and Q , a target of odor B and color Y is treated as a compound of B, Y , and R , and so forth, with Q and R representing the compound-unique components. Like the colors and odors, the compound-unique components gain and lose associative strength independently with reward and nonreward, although their salience is less than that of the colors and odors ($\alpha_{CU} = .7$ for all the compounds used in the experiments). The associative strength of a compound is taken as the simple sum of the associative strengths of all of its components; in the case of AX , for example, $V_{AXQ} = V_A + V_X + V_Q$.

Choice between two targets—say, AX and BY —is based on their relative associative strength, r . The relative strength of AX , r_{AXQ} , is computed as $V_{AXQ} / (V_{AXQ} + V_{BYR})$, and the relative strength of BY , r_{BYR} , as $V_{BYR} / (V_{AXQ} + V_{BYR})$, or $1 - r_{AXQ}$. P_{AXQ} , the probability of choosing AX , is a power function of r_{AXQ} , designated by the parameters $K = .75$ and $s = .625$, and is plotted for $.5 \leq r \leq 1$ in Figure 1. This function was selected on empirical grounds from among the many other possibilities that included matching ($K = 1, s = .5$) and maximizing ($K = 0, s = .5$), all of them generated with the scaling equations,

$$P = .5 + s(2r - 1)^K,$$

for $r \geq .5$, and

$$P = .5 - s(1 - 2r)^K,$$

for $r \leq .5$. These functions should not be confused with empirical plots of the relation between choice ratio and reward ratio that are familiar from the vertebrate litera-

ture. What Figure 1 shows for each of the three plotted functions is the predicted probability of choosing a target as a function of its relative associative strength, which is computed on the basis of previous experience (both rewarded and nonrewarded) with the alternatives.

For information on the selection of parameter values, the papers describing the development of the theory should be consulted (Couvillon & Bitterman, 1985, 1986, 1987, 1988). The values were estimated in a series of systematic (factorial) explorations; at each stage of the work, a search was made for the set of parameter values that would permit the most accurate simulation of *all the data accumulated up to that time*. The simulation procedure was stochastic to begin with and later, for reasons both of accuracy and efficiency, computational (see Couvillon & Bitterman, 1989, pp. 219-220). The familiar stochastic (Monte Carlo) procedure is to determine the choice made by each of many "stat bees" on each visit by consulting a table of random numbers and then to change its individual pattern of associative strengths accordingly. The newer procedure is to compute mean associative strengths on the assumption that the proportion of animals choosing each alternative on each visit is as given by the hypothetical choice function. Goodness-of-fit, measured in terms of the root-mean-square (RMS) deviation of simulated from actual performance, is limited, not only by error in the theory, but also by error in the data and—when the stochastic procedure is used—by random error in the simulation process itself. Simulation error is eliminated when the computational procedure is used.

The best fits to the earliest data, which came from work with targets differing only in odor (Couvillon & Bitterman, 1985, 1986), were provided by a cluster of choice functions similar to $K = .75, s = .625$, and the same functions continued to provide good fits—with $K = .75, s = .625$ marginally the best—to the data of subsequent experiments with targets differing in color as well as in odor (Couvillon & Bitterman, 1987, 1988). The earliest odor experiments suggested that both $U\beta$ and $D\beta$ must be relatively large, but not until the subsequent work on color-odor problems did it become evident that $D\beta$ must be larger than $U\beta$. To begin with, the associative strength of each stimulus at the start of training (designated V_0) was treated as a parameter (and estimated by fitting in the same way as the other free parameters) because of uncertainty as to the effects of the pretraining procedure. Now, $U\beta$ is used to compute the growth of V from 0 on the basis of the number of rewarded responses to each stimulus in pretraining; with two rewarded responses in the standardized pretraining procedure of previous experiments and with $U\beta = .2, V_0 = .36$. The concept of compound uniqueness was introduced when required by the data on conditional discrimination (Couvillon & Bitterman, 1988), but then its application was perfectly general—targets differing only in odor do, of course, have a common color (e.g., AX+/BX-), and targets differing only in color have a common odor. The estimation of α_{CU} was based, therefore, on the new results together with all of the previous ones. The best present estimates of the free

parameters, that is, the values that provide the best fit to what is now an extensive set of data, are those already specified: $U\beta = .2, D\beta = .35, \alpha_{CU} = .7$, and, for the choice function, $K = .75, s = .625$. It was those values that were used to predict the results of the present experiments.

NEW EXPERIMENTS

Separate groups of animals were trained in ten 100-trial choice problems in which probabilities of reward were varied widely. Listed in Table 1 are 9 of the 10 problems studied, each involving two odors (A and B) presented on targets of a common color (X)—AX and BX—and differentially rewarded as indicated by the conventional notation in the table. For example, 100:0 means that one of the odors (A) was rewarded on all trials and the other (B) was never rewarded; 100:0→80:20 means that A was rewarded consistently in the first 40 trials, but only on 8 (quasi-randomly selected) trials of every 10 in the next 60 trials (with B rewarded on the rest); 70:30→30:70 means that A was rewarded (quasi-randomly) on 7 of every 10 trials in the first 40 trials (with B rewarded on the rest) and on 3 of every 10 in the next 60 trials (with B rewarded on the rest); and so forth. (Since the correction method was used, the distribution of rewards was independent of the choices actually made.) The 10th problem—the ambiguous-cue problem, which has produced a variety of different results in vertebrates (e.g., Hall, 1980)—was chosen because it involves three probabilities; with odor A consistently rewarded on 50 A+/B- trials interspersed quasi-randomly with 50 B+/C- trials on which odor B was consistently rewarded, the overall probabilities of reward were 1 for odor A, .5 for odor B, and 0 for odor C. The three odors were presented on targets of a common color, X—AX, BX, and CX. Quantitative predictions of performance in the 10 problems are provided by the theory, all of its parameters having the values estimated previously.

Some preliminary data on probability of reward were reported in the first paper of this series (Couvillon & Bitterman, 1985), but they were of limited value because of the small number of training trials that were given. The standard practice in this work with free-flying foragers

Table 1
Two-Stimulus Problems

Trials			
1-40		41-100	
100:0	—	100:0	
100:0	→	90:10	
100:0	→	80:20	
100:0	→	70:30	
100:0	→	50:50	
100:0	→	20:80	
100:0	→	0:100	
70:30	→	30:70	
50:50	—	50:50	

Note—See text for explanation.

(which is dictated by a variety of practical considerations) is to train each animal in a single experimental session scheduled to last no longer than 3–4 h. During that period, the animal makes repeated visits to the laboratory, on each of which it ingests about 50 μl of sucrose solution and then leaves to deposit the sucrose in the hive, returning to the laboratory for more 4 or 5 min later. In previous experiments of this series, the rewarded alternative always contained a large drop of sucrose solution from which feeding to repletion was permitted, which is to say that there was only a single trial on each visit and the number of trials that could conveniently be scheduled in a training session therefore was severely limited. In the present experiments, the amount of reward given on each trial was only 10 μl of sucrose, and with five trials per visit on average, rather than only a single trial, it was feasible to schedule 100-trial problems.

Method

Subjects. The subjects were 84 honeybees (*Apis mellifera*), all experimentally naive, from hives situated near the laboratory. They were assigned at random to the various problems, 8 to each of the two-stimulus problems and 12 to the ambiguous-cue problem.

Apparatus. The training situation is diagrammed in Figure 2. It consisted of a pair of immediately adjacent windows (each 55 cm wide and 55 cm high) separated by a thin wooden partition around which the animal was required to fly as it shuttled between the windows on the successive trials of each visit. In training with a single trial and feeding to repletion on each visit, a single window can be used because preparations for successive trials can be made in the intervals between visits (while the animal is away at the hive). In training with multiple trials on each visit, successive trials are alternated between two windows, and preparations for Trial $N+1$ in one window are made during Trial N in the other. The two-window situation used here was like that developed by Lee and Bitterman (1990) for discrimination training with single targets dif-

fering in color or odor and containing different amounts of reward, but it was modified for the study of choice by the addition of two hinged panels. The purpose of the panels was to restrict activity to the outer portion of each window; without the panels, substantial position preferences—usually for the target nearer the central partition—sometimes appeared.

The targets were covered petri dishes of gray plastic, 5.5 cm in diameter. Drilled in the cover of each dish, 6 mm from its outer circumference, was a circle of eight equally-spaced holes, .5 mm in diameter. The dishes contained pieces of cotton batting that were impregnated with scents. For the animals trained in the two-stimulus problems, three sets of targets were used, one set scented with peppermint, a second set scented with geraniol, and a third set (used only in pretraining) scented with both. For animals trained in the ambiguous-cue problem, four sets of targets were used—a peppermint set, a geraniol set, a set scented with jasmine, and a pretraining set scented with all three odors. The covers of the targets used on each visit were washed and exchanged for others in their sets after the visit in order to randomize extraneous stimuli.

Procedure. In the pretraining, experience with both windows was given. A single animal was selected at random from a group of foragers at a feeding station providing 10%–12% sucrose solution, picked up in a matchbox, carried to the laboratory, and released at a large drop (> 100 μl) of 50% sucrose solution on a pretraining target. The pretraining target was set on the sill of one of the two windows (the left for half the animals and the right for the rest), as shown in the right window of the diagram in Figure 2. 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, where it deposited the sucrose. Typically, the animal (adapted to 50% sucrose and now finding the lower concentration at the feeder unacceptable) would return to the laboratory after a few minutes, continuing to fly back and forth between the hive and the laboratory as long as sucrose was available there. If the marked animal did not return after its first placement, it was captured again at the feeding station, where it usually could be found, and placed on the pretraining target. When the animal returned to the first window, it was picked up after a few seconds of feeding and carried to the alternative window, where it fed to repletion on another pretraining target. On




Figure 2. The two-window training situation. Shown in the right window is the position of a single target on pretraining trials, and in the left window the placement of a pair of targets on choice trials. The experimenter, in a relatively dark laboratory room, observes the subject through the sliding screens, which are shifted to permit placement and removal of targets. While the subject is feeding at the conclusion of a trial in one window, preparations are made for the next trial in the alternative window, to which the subject flies from the first window after finishing the food available there.

the second return, a pretraining target was to be found again at the alternative window, where feeding to repletion was permitted. (After the first return, there was no further handling of the animal.) In the two-stimulus problems, there were two subsequent returns with feeding to repletion, one to odor A alone in one of the windows and the next to odor B in the other; odors and windows were balanced over the animals in each group. In the ambiguous-cue problem, there were two subsequent returns to pretraining targets, one in each window, with order balanced over the animals. The pretraining of every animal ended after it had returned twice to each window of its own accord.

It may be well to note that the associative strength of each component at the end of pretraining (V_0) was predicted to be larger than in the previous experiments with a single window because the two-window situation required somewhat more pretraining. Specifically, V_0 was computed as .49 for the two-stimulus animals, which had three rewarded returns to each odor in pretraining (two to two-odor

pretraining targets and one return to each single-odor target), and .59 for the ambiguous-cue animals, which had four rewarded returns to each odor in pretraining (four returns to three-odor pretraining targets).

Arriving from the hive on each training visit, an animal found a pair of targets, scented differently and set 10 cm apart on the sill of one of the windows, the left window for half the animals in each group and the right window for the rest. The diagonal placement of the two targets is shown in the left window of the diagram in Figure 2. Of the two targets, one contained a 10- μ l drop of 50% sucrose solution, and the other contained a 10- μ l drop of tap water. Position as well as odor was balanced systematically in each problem. The initial choice made by the animal was recorded, with immediate correction of error permitted, and as soon as the animal found the sucrose, a second pair of targets was placed symmetrically on the sill of the adjoining window. When the animal finished the sucrose, it would fly up from the positive target, whereupon

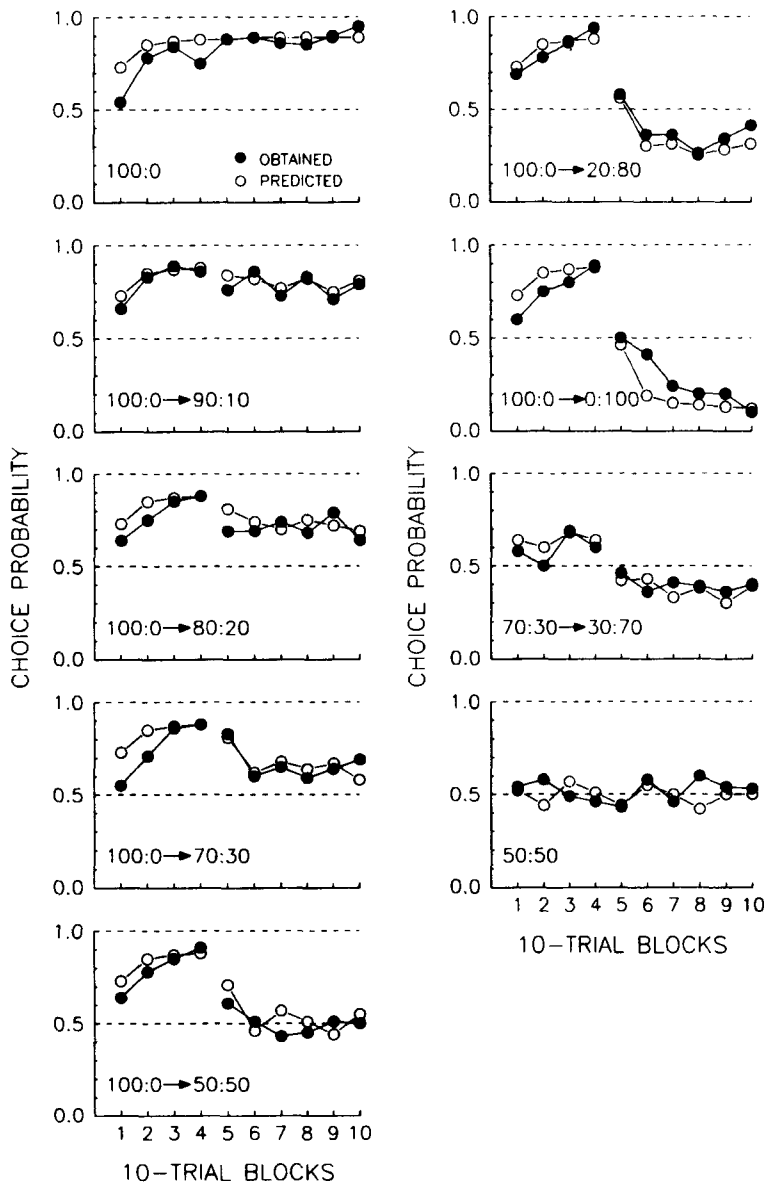


Figure 3. Obtained (filled circles) and predicted (open circles) performance in the nine two-stimulus problems.

both targets in the first window were removed, and the animal went to the adjoining window for a second trial. Then a third trial was given in the arrival window, and so forth, until—typically after five trials—the animal was replete and left of its own accord for the hive. The training was continued until the animal had made 100 choices (which required 20 visits on average) in a single session that lasted about 3 h on average. Each animal was trained in only 1 of the 10 problems.

Results

The performance of the animals in each of the nine two-stimulus problems is shown in Figure 3 (filled circles) along with the performance predicted by the theory with the parameter values specified at the outset (open circles). The curves are plotted in terms of the mean proportion of response (in successive 10-trial blocks) to the original 100% stimulus wherever the schedule was 100:0 at the outset, to the original 70% stimulus in the 70:30→30:70 case, and to one of the two stimuli selected at random in the 50:50 case. Acquisition in the 100:0 training was rapid in each case, and so also, where the reward ratio was changed, was the shift in responding to a level that varied with the new reward ratio. The 70:30 problem produced a small preference for the 70% stimulus and then a corresponding preference for the alternative when the reward ratio was changed to 30:70. In the 50:50 case, the choice ratio remained at about the chance level throughout. It may be well to note that the various levels of performance shown by the curves are not artifacts of averaging. In the 50:50 case, where (as would be expected) variability was highest, the most deviant animal showed a 63:37 preference for one of the alternatives. The reliability of measurement was in general quite high, with standard errors for the separate problems ranging from .038 to .056.

Figure 3 suggests that the performance of the animals was essentially asymptotic in the last four 10-trial blocks of each of the two-stimulus problems; on that assumption, the relation between asymptotic choice ratio and the prevailing reward ratio is shown in Figure 4 (filled circles). There are nine plotted points, one point for each of the problems—the proportion of responses (in Trials 61–100) to the original 100% stimulus for groups with 100:0 training at the beginning, the proportion of responses to the 30% stimulus for the group shifted from 70:30 to 30:70, and the proportion of responses to one of the two stimuli chosen at random for the group with 50:50 training throughout. (There are two 50% points, one for the latter group and another for the group shifted from 100:0 to 50:50.) Plotted also in Figure 4 are the predicted choice ratios (open circles) and the line of best fit to the obtained ratios ($y = .18 + .68x$).

In Figure 5 (filled circles), performance in the ambiguous-cue problem is plotted separately over successive blocks of five trials for the A+/B- and B+/C- pairs. As in previous work in which the reward was feeding to repletion (Couvillon & Bitterman, 1986), each of the 12 ambiguous-cue animals made fewer errors on B+/C- trials than on A+/B- trials ($p < .0003$). Plotted also

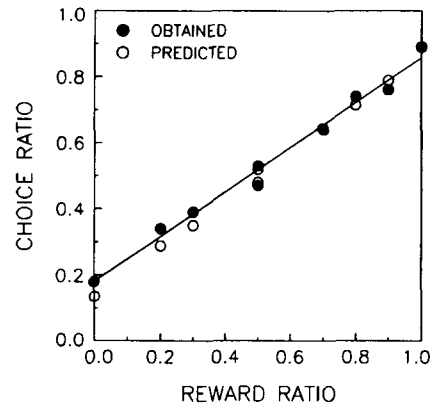


Figure 4. Obtained (filled circles) and predicted (open circles) relation between asymptotic choice ratio and reward ratio in the nine two-stimulus problems. The straight line of best fit to the obtained data is also shown.

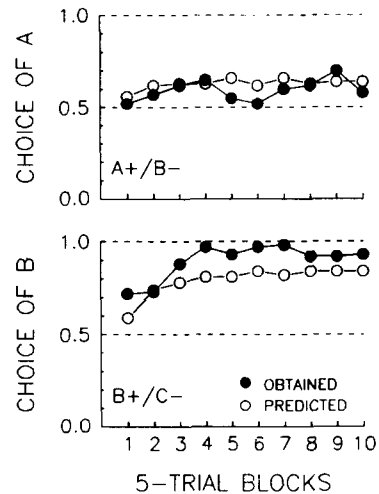


Figure 5. Obtained (filled circles) and predicted (open circles) performance in the ambiguous-cue problem.

in Figure 5 (open circles) is the performance predicted by the theory, according to which the relative associative strength of B on B+/C- trials should be greater throughout than that of A on A+/B- trials.

Inspection of Figures 3 and 5 shows that the principal features of the results are captured nicely by the theory; the RMS deviation of the predicted proportions from the obtained proportions in all 10 problems is .079. There are only two potentially important discrepancies, one of which is that the rate of acquisition in the initial 100:0 training was somewhat overpredicted. During early visits, the animals were still learning to shuttle between the windows—see the latency data of Lee and Bitterman (1990) and Loo and Bitterman (1992), who studied successive rather than simultaneous discrimination in the two-window situation—which may have contributed an element of randomness to the early discrimination performance.

That explanation might be evaluated by beginning the choice training only after a more extended period of shuttle training. The second potentially important discrepancy is that B+/C- performance in the ambiguous-cue problem was underestimated. It suggests that there may be some afferent interaction between the two members of each pair of stimuli, as has been proposed to account for better-than-chance performance of pigeons in "loop" problems (Couvillon & Bitterman, 1992). It might be interesting to study the performance of honeybees in such problems, the simplest instance of which would be concurrent training with A+/B-, B+/C-, and C+/A-. The only other discrepancy worth mentioning is that which occurs for the sixth block of training trials in the 100:0→0:100 problem and which is as likely to reflect sampling error as error in the theory.

DISCUSSION

The purpose of this work, which substantially extends the store of data on learning in honeybees, was to test a wide range of exact rather than merely ordinal predictions from a quantitative theory of discrimination and choice. The independent variable of central interest was probability of reward, which made it necessary to introduce some substantial changes in the training conditions used to provide the data for which the theory was originally developed to account. Specifically, the amount of reward for each correct choice was reduced, the number of choices made on each visit was increased, and the animals were required during training to shuttle back and forth between two sites at which the choices were made. With reasonably successful prediction of the results obtained under the new conditions, confidence in the theory and in its generality is enhanced.

An interesting feature of the linear relation between the choice ratio and the reward ratio shown in Figure 4 is that its slope (.68) is less than 1. It might be argued that performance in the last 40 training trials actually was not asymptotic and that the slope of the relation might have approached 1 with continued training. The postshift performance in most of the problems does not, however, provide much support for that possibility. The theory itself suggests otherwise—simulations of 100 additional training trials on each problem show no significant changes in terminal performance—although why the theory predicts that the slope of the relation remains less than 1 may not be immediately obvious. Consider, for example, the 100:0 problem (AX consistently rewarded and BX never): Given the independence rule, V_A and V_Q (Q is the compound-unique component of AX) approach λ (= 1), while V_B and V_R (R is the compound-unique component of BX) approach zero, but, because V_X —the associative strength of the common color—is substantial (.87 after 200 trials), r_{AXQ} remains relatively low (.77 after 200 trials), and (reading from the choice function $K = .75$, $s = .625$ plotted in Figure 1) P_{AXQ} remains less than 1 (.89 after 200 trials). The dynamics are the same for the

other reward ratios. Nevertheless, because the question is of such importance for the theory, it may be well to make some further asymptotic determinations with larger numbers of trials and subjects.

Given their historical importance in work with other animals, it is interesting to ask how well the matching and maximizing functions fit the honeybee data. The matching function ($K = 1$, $s = .5$) yields a good fit to the data of Figure 4 with $U\beta = .02$, $D\beta = .3$, and $\alpha_{CU} = .8$, the common stimulus (X) again accounting for the fact that the slope is less than 1. On this point, see Estes (1959), who himself published a plot of data on human verbal conditioning remarkably similar to Figure 4 (Estes, 1957). The matching function does not, however, yield a satisfactory fit to the ambiguous-cue data, predicting B+/C- performance (.69) scarcely better than A+/B- performance (.61). Although the matching fit to all of the present data is not much poorer than the fit with $K = .75$, $s = .625$ (RMS = .089 as compared with .079), the matching fit to the ambiguous-cue data alone is considerably poorer (RMS = .222 as compared with .126), as it is to all of the previous data (RMS = .213 as compared with .119).

The best maximizing fit (with $U\beta = .4$, $D\beta = .28$, $\alpha_{CU} = 1$, RMS = .202) predicts consistent choice both of A over B and of B over C in the ambiguous-cue problem (RMS = .378). It also predicts that asymptotic choice ratios will exactly *match* the terminal reward ratios in the two-stimulus problems (plotted as in Figure 1, the predicted ratios lie on a straight line with a slope of 1), the explanation being that the large β s produce what has been called "reward following" in the vertebrate literature (Bitterman, 1975); with choice on each trial heavily determined by the outcome of the immediately preceding trial, choice probabilities tend to mirror reward probabilities. In the 50:50 problem, for example, the maximizing prediction is that the alternative chosen on each trial will *always* be the rewarded alternative of the immediately preceding trial. With $K = .75$, $s = .625$, the predicted probability of reward-following in the 50:50 problem is not 1, but .54; the obtained probability was .53.

These computations emphasize the importance of the distinction, often blurred, between *theoretical* choice functions, which relate predicted choice ratios to ratios of hypothetical associative strengths (as in Figure 1), and *empirical* choice functions, which relate actual choice ratios to actual reward ratios (as in Figure 4). The theoretical matching function need not predict empirical matching, and the theoretical maximizing function need not predict empirical maximizing. Not only may the same theoretical function predict quite different outcomes (depending on the learning assumptions with which it is combined), but different theoretical functions may predict some of the same outcomes. What is required to decide among competing quantitative theories, as among purely verbal theories, is a constantly expanding data base—the broader the base, the smaller the likelihood that it will be compatible with contradictory assumptions. Although the

matching function yields as good a fit to the data of Figure 4 as does $K = .75$, $s = .625$, the fit with the same parameters to the ambiguous-cue data is, as already noted, distinctly inferior, and so also is the fit to all of the data of the wide range of previous experiments on choice in honeybees. What we are looking for, of course, is a theory that provides the best quantitative account—not simply an ordinal account (cf. Daly & Daly, 1982)—of all available data.

It should be obvious that not all features of a broad theory may be tightly constrained, or even constrained at all, by the data of a particular experiment or set of experiments. It would be possible, certainly, to simulate the results of the present experiments without the assumption of compound-unique components that was introduced originally to account for conditional discrimination—as, for example, when honeybees were rewarded for choosing blue rather than green if both targets were scented with peppermint and for choosing green rather than blue when both targets were scented with geraniol (Couvillon & Bitterman, 1988). There is no reason, however, to assume that afferent interaction occurs only under the aegis of conditionality; therefore, compound-unique components must be included in all calculations. By the same token, no assumptions are made that are not required by the data, and without relevant data there would be no way to calculate their potential effect. Before any quantitative statements can be made, for example, about the position of a target, as apart from its color and odor, experiments will be required in which responses to position are differentially reinforced under the standard training conditions employed; there is already some evidence from work under rather different conditions (Klosterhalfen, Fischer, & Bitterman, 1978) that position (which is not a feature of a target but is given only in relation to its surround) will require special treatment. As to the parsimonious independence rule incorporated in the theory, there is no reason to modify or replace it—whatever the vertebrate literature suggests—before new honeybee data are obtained to contradict it.

It is interesting to note that λ is not significantly constrained either by the present experiments with small reward or by the entire set of previous experiments with feeding to repletion as reward; that is, the predictions are essentially the same for all values of λ despite evidence from somewhat different (successive discrimination) experiments that λ increases with amount of reward (Couvillon, Lee, & Bitterman, 1991). In those experiments, amount of reward was varied within subjects, primarily to avoid confounding with number of training trials or number of training visits, but also because work with other animals had shown within-subject designs more likely than between-subject designs to yield amount-of-reward effects (Mackintosh, 1974). The general procedure was to present two different targets individually and in quasi-random sequence, one always containing 5 μ l and the other 20 μ l of a 50% sucrose solution, with a difference in associative

strength inferred from preference for the latter in a subsequent choice test. Efforts to find amount-of-reward effects in between-subject experiments clearly are called for, although a suitable methodology remains to be developed. Negative results in such experiments would call into question the parsimonious assumption that the effect of reward is on associative strength alone and would prompt consideration of a representational alternative, which would, of course, substantially complicate the theory.

REFERENCES

- BITTERMAN, M. E. (1975). The comparative analysis of learning. *Science*, **188**, 699-709.
- BUSH, R. R., & MOSTELLER, F. (1951). A mathematical model for simple learning. *Psychological Review*, **58**, 313-323.
- COUVILLON, P. A., & BITTERMAN, M. E. (1982). Compound conditioning in honeybees. *Journal of Comparative & Physiological Psychology*, **96**, 192-199.
- COUVILLON, P. A., & BITTERMAN, M. E. (1985). Analysis of choice in honeybees. *Animal Learning & Behavior*, **13**, 246-252.
- COUVILLON, P. A., & BITTERMAN, M. E. (1986). Performance of honeybees in reversal and ambiguous-cue problems: Tests of a choice model. *Animal Learning & Behavior*, **14**, 225-231.
- COUVILLON, P. A., & BITTERMAN, M. E. (1987). Discrimination of color-odor compounds by honeybees: Tests of a continuity model. *Animal Learning & Behavior*, **15**, 218-227.
- COUVILLON, P. A., & BITTERMAN, M. E. (1988). Compound-component and conditional discrimination of colors and odors by honeybees: Further tests of a continuity model. *Animal Learning & Behavior*, **16**, 67-74.
- COUVILLON, P. A., & BITTERMAN, M. E. (1989). Reciprocal overshadowing in the discrimination of color-odor compounds by honeybees: Further tests of a continuity model. *Animal Learning & Behavior*, **17**, 213-222.
- COUVILLON, P. A., & BITTERMAN, M. E. (1991). How honeybees make choices. In J. L. Goodman & R. C. Fischer (Eds.), *The behaviour and physiology of bees* (pp. 116-130). Wallingford, U.K.: CAB International.
- COUVILLON, P. A., & BITTERMAN, M. E. (1992). A conventional conditioning analysis of "transitive inference" in pigeons. *Journal of Experimental Psychology: Animal Behavior Processes*, **18**, 308-310.
- 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.
- DALY, H. B., & DALY, J. T. (1982). A mathematical model of reward and aversive nonreward: Its application in over 30 appetitive learning situations. *Journal of Experimental Psychology: General*, **111**, 441-480.
- ESTES, W. K. (1957). Of models and men. *American Psychologist*, **12**, 609-617.
- ESTES, W. K. (1959). The statistical approach to learning theory. In S. Koch (Ed.), *Psychology: A study of a science. Study I: Conceptual and systematic. Volume 2: General systematic formulations, learning, and special processes* (pp. 380-491). New York: McGraw-Hill.
- HALL, G. (1980). An investigation of ambiguous-cue learning in pigeons. *Animal Learning & Behavior*, **8**, 282-286.
- HULL, C. L. (1930). Simple trial-and-error learning: A study in psychological theory. *Psychological Review*, **37**, 241-256.
- KLOSTERHALFEN, S., FISCHER, W., & BITTERMAN, M. E. (1978). Modification of attention in honeybees. *Science*, **201**, 1241-1243.
- LEE, Y., & BITTERMAN, M. E. (1990). Learning in honeybees (*Apis mellifera*) as a function of amount of reward: Acquisition measures. *Journal of Comparative Psychology*, **104**, 152-158.
- 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.

- MACKINTOSH, N. J. (1974). *The psychology of animal learning*. London: Academic Press.
- MACKINTOSH, N. J. (1983). *Conditioning and associative learning*. New York: Oxford University Press.
- RESCORLA, R. A. (1972). "Configural" conditioning in discrete-trial bar pressing. *Journal of Comparative & Physiological Psychology*, **79**, 307-317.
- RESCORLA, R. A., & WAGNER, A. R. (1972). A theory of classical conditioning: Variation in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- WHITLOW, J. W., JR., & WAGNER, A. R. (1972). Negative patterning in classical conditioning: Summation of response tendencies to isolable and configural components. *Psychonomic Science*, **27**, 299-301.

(Manuscript received September 4, 1992;
revision accepted for publication December 15, 1992.)