

Reciprocal overshadowing in the discrimination of color-odor compounds by honeybees: Further tests of a continuity model

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

Honeybees were trained to discriminate between simultaneously presented color-odor compounds, one group with color and odor confounded and a control group with color relevant and odor irrelevant; in subsequent differentially reinforced training with the colors in the absence of the odors, the performance of the two groups was the same (Experiment 1). When, however, response to the colors was measured in a 10-min extinction test, discrimination was found to be poorer after confounded training (Experiment 2), and like results were obtained in an extinction test with the odors after control animals had been trained with odor relevant and color irrelevant, the confounded animals showing poorer discrimination of the odors than the controls (Experiment 3). The results of the first two experiments, in which overshadowing of color by odor was found only with an extinction test, require us to take seriously the possibility that our previous modeling experiments (with probability of correct choice in differentially reinforced training as the measure of performance) may have been insufficiently sensitive to noncontinuity effects. Our first efforts to model extinction suggest, however, that all the results of the present experiments can be understood without sacrifice of the parsimonious independence principle.

In recent experiments with honeybees (Couvillon & Bitterman, 1985, 1986, 1987, 1988), we studied the discriminative learning of individual foragers shuttling back and forth between their hives and the sill of an open laboratory window, where they chose between two targets presented simultaneously, one containing sucrose solution and the other containing water. Our first results (Couvillon & Bitterman, 1985, 1986), for targets that differed only in odor, could be simulated accurately with simple equations for computing growth or decline in associative strength produced by reinforcement or nonreinforcement and for predicting choice on the basis of relative strength. Later (Couvillon & Bitterman, 1987), using targets that differed in color as well as in odor, we were able with the same equations to simulate performance in a variety of compound problems on the assumption that the associative strength of each component of a compound changes independently with reinforcement or nonreinforcement of the compound (*independence rule*) and that the associative strength of a compound equals the summed strengths of its components (*summation rule*). Surprised by this success on the basis of prior experiments with other stimuli which had seemed to call both rules into question (Couvillon & Bitterman, 1982; Couvillon, Klosterhalfen,

& Bitterman, 1983), we suspected that the problems chosen for the 1987 analysis did not sufficiently constrain the model, and we began, therefore, to look for more demanding ones.

The summation rule was challenged by evidence of compound uniqueness which had appeared in one of the 1982 experiments. In the first stage of that experiment, only a single target (from which the animal fed to repletion) was presented on each visit. On some visits, the target was labeled with an orange disk and the scent of jasmine, and, on others, with a yellow disk and the scent of lemon. The second stage consisted of a 10-min extinction test with four targets presented simultaneously, two labeled with the training compounds and two with compounds produced by recombination of the training components (orange-lemon and yellow-jasmine), in which the animals showed a clear preference for the training compounds. These results suggested that it might be possible to demonstrate conditional discrimination with our standard colors and odors in simultaneous discrimination experiments of the kind we were modeling, which in fact it was (Couvillon & Bitterman, 1988). We found that honeybees could be trained to choose one of two colors on the basis of a common odor (e.g., a green rather than a blue target when both were scented with peppermint, but blue rather than green when both were scented with geraniol), or to choose one of two odors on the basis of a common color. Clearly, then, the associative strength of a compound such as green-geraniol was different from the sum of the strengths of green and geraniol, although a fundamental revision of the 1987 model did not seem to be required. It was sufficient merely to add the purely

This research was supported by Grant BNS-8709785 from the National Science Foundation. Helpful advice from William T. Woodard is gratefully acknowledged, as is the assistance of Brendan Clark, Pharis Mohideen, Nina Morissette, Myong Song, Rebecca Steen, and Bich Ngoc Tran in collecting the data. Requests for reprints should be addressed to the Békésy Laboratory of Neurobiology, 1993 East-West Road, Honolulu, HI 96822.

perceptual assumption—borrowed from the vertebrate literature (Rescorla, 1972; Whitlow & Wagner, 1972)—that interaction among the components of a compound generates a relatively salient, compound-unique component which gains and loses associative strength in the same way as any other component and contributes in the same way to the strength of the compound. The independence and summation rules continued to apply.

We turn now to a challenge to the independence rule presented by evidence of the overshadowing of color by odor that appeared in two experiments, one reported in our 1982 paper and the other in our 1983 paper. They will be referred to as the 1982 and 1983 overshadowing experiments. In the former, an experimental group of honeybees was fed repeatedly on a gray target labeled with an orange disk and the scent of jasmine, while a control group was fed on a gray target labeled only with an orange disk. Then there was a 10-min extinction test with two gray targets, one labeled with an orange disk and the other unlabeled, in which the control group showed a clear preference for the target labeled with the orange disk but the experimental group did not. In the 1983 experiment, honeybees were trained to discriminate between two simultaneously presented targets that differed both in color and odor for a confounded group, but in color alone (the odor was the same) for a control group; in the second stage of the experiment, the two colors were differentially reinforced as before, with the targets now unscented for both groups. The performance of the confounded group was better in the first stage and poorer in the second stage than that of the control group.

Although results of this sort often have been taken to mean that the components of a compound stimulus compete with each other for attention (Sutherland & Mackintosh, 1971) or for associative strength (Rescorla & Wagner, 1972), they are not, in fact, entirely out of reach of the independence principle. Work on compound-component discrimination, which was reported along with the work on conditional discrimination in our 1988 paper, suggests a different description of the 1982 overshadowing experiment than the one given in the preceding paragraph: The control group now may be thought of as having been trained, not on an unscented target, but on a target with an odor different from jasmine—we refer to it simply for convenience as the “batting” odor (the odor of the unimpregnated cotton batting inside the target, which is replaced by the odor of any substance with which we choose to impregnate the batting). Furthermore, what we described before as a gray target labeled with an orange disk now may be thought of as a target different in color from the plain gray target. We say, then, that the experimental group was trained with orange-jasmine, the control group was trained with orange-batting (orange disk with the odor of the unimpregnated cotton batting), and then both were tested with orange-batting versus gray-batting. In the light of this analysis—especially when the compound-unique orange-batting component is considered

—the stronger preference for orange-batting displayed by the control group is just what continuity theory would lead us to expect, although it remains to be understood why the experimental group should have shown no preference whatever.

The 1983 overshadowing experiment is of particular interest to us because our model was constructed to deal with performance in simultaneous discrimination training. In the 1983 experiment, the confounded group made fewer errors than the control group—which is to say it had substantially less unreinforced experience with the negative color—and on that ground alone might have been expected to do somewhat more poorly in the test, with no reflection whatever on the validity of the independence rule (Wagner, 1969). It should not be supposed, however, that such an experiment therefore has no value, because it is possible with a quantitative theory to predict, not only *how much* less experience with the negative color the confounded group will have in the first stage of training, but *how much* more poorly it will perform in the second stage. We made such predictions for the standard stimuli used to provide the data on which our model is based and tested them in the first experiment to be reported.

EXPERIMENT 1

The design used here (slightly different from that used in 1983) was suggested by cue-redundancy experiments with rats which have yielded some evidence of overshadowing and—although the term was not yet invented at that time—some puzzling indication of potentiation as well (Lovejoy & Russell, 1967; Sutherland & Andelman, 1967). To look for overshadowing of color by odor, we trained one group of honeybees in a confounded color-odor problem and a second group with color (green vs. blue) relevant and odor (geraniol vs. peppermint) irrelevant, after which both groups were trained with the “unscented” colors. For example, a subset of confounded animals was trained with green-geraniol positive and blue-peppermint negative on all visits in the first stage, while a subset of controls was trained in the same way on half the visits and, on the rest, with green-peppermint positive and blue-geraniol negative. In the second stage, both subsets were trained on all visits with unscented-green (green-batting) positive and unscented-blue (blue-batting) negative. One of the variables studied in this experiment was the amount of training given in the first stage, and another was the nature of the test (reversal vs. nonreversal) in the second stage. One confounded group and one control (color-relevant) group—the 8NR groups—were trained in 8 visits with the colors and odors, and then trained in 10 more visits with the “unscented” colors differentially reinforced as before (nonreversal test). For a second pair of groups—the 16NR groups—the training was the same as for the first pair, except that there were 16 visits in the first stage. For a third pair of groups—the 16R groups—there were 16 visits in the first stage and

a reversal test in the second stage (the positive color of the first stage was negative in the second, and the negative color was positive); the reversal procedure was used on the basis of general vertebrate experience which suggested that it might be more sensitive. In all, there were six groups of animals, whose modeled and actual performances were compared.

The performance of six groups of 100 stat-bees, trained as the real bees were to be trained, was simulated with the model described in our 1988 paper. The simulation was done with the same choice function ($K = .75, s = .625$)—the function relating probability of choice to relative associative strength—and with the same parameters governing the growth and decline in associative strength that provided the best overall fit to the combined compound data of 1987 and 1988. The salience (α) of each of the colors (green and blue) and each of the odors (geraniol and peppermint), which were shown in our 1987 work to be equally discriminable, was 1, and the α of the batting odor of each “unscented” test target, not previously estimated, also was taken as 1, although with some misgivings because the batting odor might be thought to be less salient than the others. The α of each compound-unique component was .7. The incremental β ($U\beta$) was .2 and the decremental β ($D\beta$) was .35 (reflecting the rapid changes with reinforcement and nonreinforcement observed in the behavior of the animals). The associative strengths (V) of the colors and odors established in the pretraining was .3, and that of the batting odor was 0, no generalization from the training odors being assumed. The model predicts that the performance of the confounded group of each pair will be better than that of the control group in the first stage of training. It also predicts that the differential associative strength of the colors will be less in the confounded group than in the control group at the end of the first stage of training, because the control group will have had more unrewarded experience with the negative color. After the eighth visit, the (mean) associative strength of the positive color in both groups (equal at the outset and remaining equal because the correction method is used) is about .88, with the associative strength of the negative color about .19 in the confounded group and about .12 in the control group. After the 16th visit, the associative strength of the positive color in both groups is about .98, with the associative strength of the negative color about .18 in the confounded group and only about .07 in the control group. Simulated performance in the second stage is plotted in Figure 1. All the nonreversal groups do well from the outset, with a slight decline in performance as training proceeds that may be attributed to the increasing associative strength of the new (batting) odor, which is common to both compounds, and the consequent decline in the relative strength, r , of the positive compound. In the 16R case, both groups choose the previously positive color at the outset and reverse at about the same rate. The model predicts, then, that there will be no evidence of overshadowing.




Figure 1. Predicted (simulated) proportions of correct choice on each visit in the second stage of Experiment 1. 8NR, nonreversal performance after eight training visits; 16NR, nonreversal performance after 16 training visits; 16R, reversal performance after 16 training visits.

Method

Subjects. The subjects were 48 honeybees (*Apis mellifera*) foraging for nectar. They came from our own hives, situated in the vicinity of the laboratory, and all were experimentally naive.

Procedure. Individual bees were pretrained to fly from their hives to the laboratory and to drink to repletion from a large drop (about 100 μ l) of 50% sucrose solution on a target that was set on the sill of an open window. An animal was selected at random from a group of foragers at a feeding platform equipped with a large jar of 10%–12% sucrose solution, carried in a small matchbox to the laboratory, and placed on the target. There it was permitted to drink its fill of 50% solution (during which time it was marked with a spot of colored lacquer) and then to fly to the hive. Typically, the animal returned to the laboratory of its own accord in 3–4 min, continuing thereafter to shuttle back and forth between the hive and the laboratory as long as food was available. If the marked bee did not return to the laboratory after the first placement, it was picked up again at the feeding platform (where it usually could be found), carried back to the laboratory, and set down on the target once more. The pretraining ended with the subject's second return to the laboratory of its own accord. It should be noted that the work could be done with individual foragers because recruitment is not a problem

under our conditions. Intruders (bees which have not themselves been carried to the laboratory) are rare and can quickly be captured, as is each subject at the conclusion of work with it.

The targets were covered petri dishes of clear plastic, 5.5 cm in diameter. In each cover, eight equally spaced holes, .5 cm in diameter, were drilled at the outer circumference. Some of the covers were green, and others were blue. The dishes themselves contained cotton batting that either was impregnated with the odor of peppermint or with the odor of geraniol, or was unimpregnated. In all, six sets of targets were used in the training: green-peppermint, blue-peppermint, green-geraniol, blue-geraniol, green-batting, and blue-batting. The covers of the targets used on each visit were washed and exchanged for others in the same set after the visit in order to balance extraneous stimuli. For purposes of pretraining only, there was an additional set of targets, each of which was half-green, half-blue, and scented both with peppermint and geraniol.

On each training visit, the animal found two targets placed 10 cm apart on a line parallel to the outer edge of the window sill. One of the targets (half the time on the right and half the time on the left, in balanced quasi-random orders) contained a drop of 50% sucrose solution, and the other contained a drop of tap water (unacceptable and distinguishable from the sucrose only by taste). The target first chosen on each visit was recorded, and the trial ended when the animal had found the correct target, fed to repletion, and gone back to the hive. Color was balanced over groups, with green positive for half the animals in each group and blue positive for the rest. Odor also was balanced: For half the confounded animals, the positive color was paired with geraniol and the negative color with peppermint in the first stage of training, and the opposite was true for the rest; for each control animal, of course, each color appeared equally often with each odor. In the second stage of training, the unscented green and blue targets were used, again with position varied over visits in balanced quasi-random orders. As already noted, there were eight visits in the first stage for the 8NR confounded and control groups, but 16 visits for all the others. In the 10 visits of the second stage, the positive and negative colors were reversed for the 16R groups, but differentially reinforced as before for all the others. There were 8 subjects in each of the six groups.

Results

As expected from the simulation and from our 1987 data, the confounded animals made significantly fewer errors than did the control animals, both in eight trials (median test, Fisher's exact $p < .0001$) and in all 16 trials ($p = .0055$). Performance in the second stage of the experiment, which is plotted in Figure 2, bears out the prediction (cf. Figure 1) that there will be marked positive transfer in the nonreversal groups and marked negative transfer in the reversal groups, but little difference in the number of errors made by confounded and control groups (median test, $p > .05$ in each case). A feature of the simulated curves that is not evident in the empirical curves is the slight decline over visits in the performance of the nonreversal groups. Whether the discrepancy reflects model error (as, e.g., in the shape of the choice function) or a limitation on the reliability of the choice measure, it is difficult to say. The simulated curves are smoother, of course, than the obtained curves, which are based on the performance of only 8 animals per group; it will be remembered that the number of stat-bees (which are far cheaper to run than real bees) was 100 per group.

Quite apart from the fact that the main results of this experiment are predicted from the theory, we must consider




Figure 2. Proportions of correct choice on each visit in the second stage of Experiment 1. 8NR, nonreversal performance after eight training visits; 16NR, nonreversal performance after 16 training visits; 16R, reversal performance after 16 training visits.

the question of why they are so at variance with those of the 1983 experiment, which showed marked overshadowing of color by odor. One explanation may be that the pairs of colors and odors used in the 1983 experiment were different in salience or discriminability, whereas the standard colors and odors of our modeling experiments are not; with the 1983 stimuli, for example, the odor-irrelevant (color-irrelevant) problem was mastered somewhat more quickly than the color-relevant (odor-irrelevant) problem, but in the 1987 experiments with the standard stimuli the two problems proved to be equal in difficulty. Our next step, therefore, might have been to repeat Experiment 1 with more intense odors in the hope of finding overshadowing so that we could try to model it, but we chose instead to consider the possibility that we could find overshadowing even with our standard stimuli if a different testing method was employed.

The model itself suggests that performance in a test with differential reinforcement may not be optimally sensitive to differences in associative strength prevailing at the outset, especially perhaps (given the steepness of the choice

function) for extreme values of *r* (relative associative strength), but in general because learning is so rapid. If testing with differential reinforcement is not sensitive to the effects of the substantial differences in experience with the negative color recorded in Experiment 1, it may not be adequate either for exploring what may be more subtle effects of competition. We decided, therefore, to use the extinction procedure of the 1982 experiment, which we suspected—on the basis of a variety of informal, cross-experiment comparisons—to be more powerful.

EXPERIMENT 2

The stimuli used in this experiment were the same as in Experiment 1, and the general logic also was the same—learning about colors was tested after training with confounded or irrelevant odors—but the testing procedure was different. The test (like that used in the 1982 overshadowing experiment) was a period of unreinforced exposure to the two unscented colors as a pair, with the number of responses to each recorded.

Method

Subjects. The subjects were 32 foragers from our own hives. All were experimentally naive.

Procedure. There were two groups of animals in the experiment, one confounded and one control group, with 16 animals in each group. The stimuli used, the pretraining procedure, and the training procedure were exactly the same as they were for the 16-vision groups of Experiment 1. On the single test visit, which followed the 16th training visit, unscented green and blue targets were presented for a period of 10 min, with position balanced over subjects within groups. Both targets contained water, and the number of responses made to each of them was recorded.

Results

As expected, the training results of this experiment closely resemble those of Experiment 1, the control group making more errors and therefore having more unreinforced experience with the negative color than the confounded group (median test, exact *p* < .0006). The acquisition curves are plotted in Figure 3.

In Figure 4, the test results are plotted in terms of the mean cumulative number of responses to the previously positive (S+) and negative (S-) colors over successive




Figure 3. Proportions of correct choice on each visit in the first stage of Experiment 2.

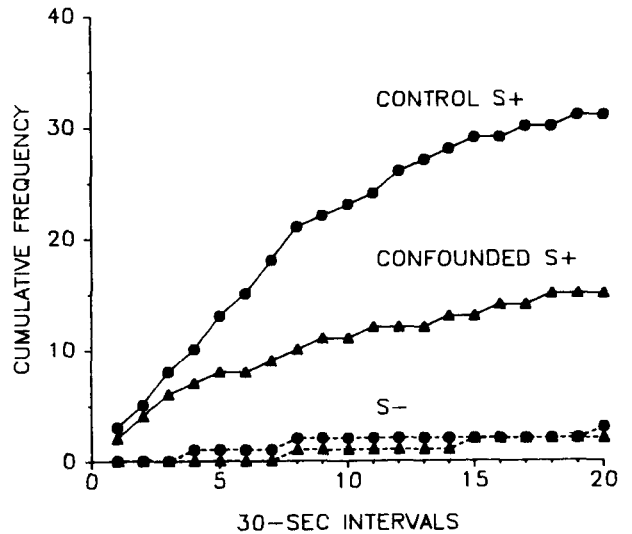


Figure 4. The course of extinction in Experiment 2. The curves show mean cumulative number of responses to the formerly positive (S+) and negative (S-) colors over successive 30-sec intervals of the 10-min extinction period.

30-sec intervals of the 10-min extinction period. The animals of both groups responded almost exclusively to S+. The critical finding is that the confounded animals showed poorer discrimination (in the sense that they responded less to S+) than did the control animals, which may be taken as a demonstration of overshadowing. Analysis of variance yields a significant effect of type of training (confounded vs. control) [*F*(1,30) = 9.20, *p* = .0049], a significant stimulus effect (S+ vs. S-) [*F*(1,30) = 70.76, *p* < .0001], and a significant interaction of stimulus × type of training [*F*(1,30) = 9.25, *p* = .0049]. Comparison of these results with those of Experiment 1 suggests that the absence of an overshadowing effect in Experiment 1 can be attributed to the insensitivity of the testing procedure employed. The question now to be faced is whether the substantial difference in the extinction performance of the confounded and control groups of Experiment 2 can be accounted for in terms of the relatively small difference in the associative strengths of the negative colors which we have computed on the basis of our continuity model, or whether the independence principle will have to be abandoned. To answer this question, we must consider how the model can be extended to deal with extinction, but it may be well before doing so to test another prediction from the model, which is that the color-odor overshadowing is reciprocal.

EXPERIMENT 3

In the 1982 experiment already described that showed overshadowing of color by odor, potentiation of odor by color also was demonstrated with the same stimuli and the same procedure. An experimental group was trained with orange-jasmine (a gray target labeled with an orange disk and the scent of jasmine) and a control group was

trained with gray-jasmine (a gray target scented with jasmine). Then there was a 10-min extinction test with two gray targets, one scented with jasmine (gray-jasmine) and the other "unscented" (gray-batting), in which the Experimental group showed a *greater* preference for jasmine than did the control group. Whatever the reason for those results, which cannot, of course, be explained in the same manner as the overshadowing results, the model predicts that in an experiment like Experiment 2, with our standard stimuli, there will be overshadowing rather than potentiation. After 16 training visits with color and odor confounded for one group of animals and with odor relevant (color irrelevant) for a second, the associative strengths of the odors will be exactly those computed for the colors in Experiments 1 and 2. The difficulty of the odor-relevant problem will be equal to that of the color-relevant problem (which has been demonstrated already in our 1986 paper), and the relative associative strength of the positive odor will be greater after color-relevant training than after confounded training, again because of greater unrewarded experience with the negative odor.

Method

Subjects. The subjects were 32 foragers from our own hives, all experimentally naive.

Procedure. The procedure was exactly the same as in Experiment 2, except that the control group was trained with odor relevant (color irrelevant) and both groups were tested with a pair of gray targets labeled with the odors. Again there were 16 animals in each group.

Results

As expected from our previous work (1987), the training results of this experiment were much like those of Experiment 2, the control group making more errors than the confounded group and therefore having more unreinforced experience with the negative odor (median test, exact $p < .0004$). The acquisition curves are plotted in Figure 5, and, in Figure 6, the extinction curves, which show the predicted overshadowing, not potentiation. Analysis of variance yielded a significant effect of type of training (confounded vs. control) [$F(1,30) = 9.44$, $p = .0045$], a significant stimulus effect (S+ vs. S-) [$F(1,30) = 196.91$, $p < .0001$], and a significant inter-




Figure 5. Proportions of correct choice on each visit in the first stage of Experiment 3.

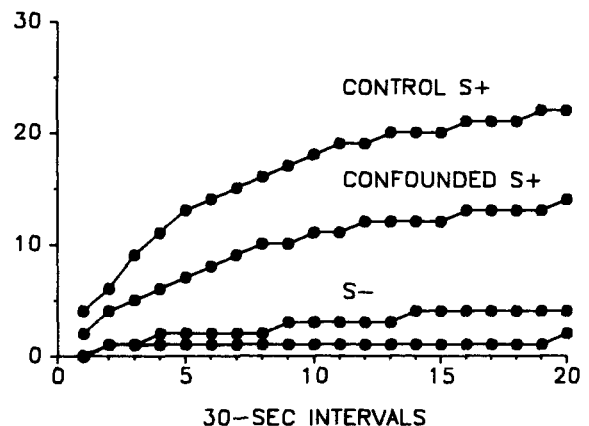


Figure 6. The course of extinction in Experiment 3. The curves show mean cumulative number of responses to the formerly positive (S+) and negative (S-) odors over successive 30-sec intervals of the 10-min extinction period.

action of stimulus \times type of training [$F(1,30) = 7.94$, $p = .0085$].

Comparison of Figure 6 with Figure 4 suggests that the test targets used in the second experiment were more discriminable than those used in the present experiment, and it may be well to note that we were prepared for such a difference. In Experiment 2, where the animals were extinguished with the "unscented" training colors (green-batting and blue-batting), it was safe, on the basis of our experience, to assume that there would be essentially no generalization from the training odors to the odor of the unimpregnated batting, but here the animals were extinguished with the training odors on targets of a common new color (gray) expected to have some not-inconsiderable generalized strength, which would decrease the relative strength of the S+ compound. In our 1982 overshadowing-potentiation experiments, animals extinguished after training with orange-jasmine showed a strong preference for gray-jasmine over orange-batting, which may have been due in part at least to greater generalization from orange to gray than from jasmine to batting. Much the same interpretation can be offered of a comparable early result reported by von Frisch (1914).

DISCUSSION

In Experiments 1 and 2, we looked for overshadowing of color by odor in simultaneous discrimination training with color-odor compounds. Tests with differential reinforcement (both reversal and nonreversal) in Experiment 1 proved insensitive to differences in training experience, but extinction tests showed overshadowing in Experiment 2 and reciprocal overshadowing (of odor by color) in Experiment 3. These results suggest the possibility—not entirely unforeseen (Couvillon & Bitterman, 1988, p. 74)—that our simple model has succeeded so well in the past because probability of correct choice under conditions of differential reinforcement is too crude a mea-

sure to capture what may prove to be relatively subtle non-continuity effects. It may be well nevertheless to emphasize that the term "overshadowing" is used here in a purely empirical sense, with no reflection whatever on the validity of the independence principle. It remains to be seen whether the continuity model, extended to deal with extinction, can account for the results of Experiments 2 and 3 in terms of the less frequent experience with S- of confounded as compared to control animals.

Our approach to the problem has been to treat the continuous extinction procedure as affording opportunity for a series of discrete choices between the two alternatives presented, with the probability of choice determined as before by their relative associative strength (r). The choice function $K = .75$, $s = .625$, used to simulate acquisition, is used again for extinction. To map the choices in time, the latency of each choice (or interresponse time) is computed from the equation $L = U \cdot (1/r^M)$, where L is the latency (in seconds), r is relative strength of the stronger alternative, and U and M are scaling factors. The latency curve for $U = 6$ and $M = 5$, the values that turned out to provide the best fit to the pooled data of Experiments 2 and 3, is plotted in Figure 7. Assuming at the outset that latency must depend both on the relative strength of the alternatives and on their absolute strengths, we were surprised in the course of our explorations to come upon a reasonable fit to the data on the basis of relative strength alone, and we decided to see how far it would take us. Not much guidance on this point is available in the vertebrate literature, although there is some supporting evi-

dence from a T-maze experiment with rats by Spear and Pavlik (1966), who measured speed of response in extinction after 12-1 (12 pellets in one arm and a single pellet in the other), 12-12, and 1-1 training: Irrespective of the alternative actually chosen on any given extinction trial, "turning" and "committed" speeds were greater in the 12-1 group than in the other two, which responded with equal speed.

An assumption which turned out to be correct was that we would not be able to fit the extinction data with the large decremental β (on the order of .3) required in our simulations of simultaneous discrimination training, and the reason seems perfectly clear: In simulating simultaneous discrimination training, we have been restricting our consideration to initial errors. The associative strength of each component of the unrewarded alternative is decremented only once when a stat animal chooses it on any visit, despite the fact that a real animal may respond to it repeatedly before shifting to the rewarded alternative, and the decremental β therefore reflects the effect of a series of unreinforced responses rather than of one alone. In extinction, by comparison, all responses to each alternative are considered, which may be responsible at least in part for the greater sensitivity of the extinction test to differences in associative strength developed in training. We have, in fact, been thinking for some time that it may be useful to try to model repetitive errors in simultaneous discrimination training. For the moment, however, we were content merely to add $D\beta$ to the list of parameters—with U and M , now three in all—to be estimated in our attempt to deal with the extinction data. The best fit was for $D\beta = .08$. The associative strengths of the colors and odors at the start of extinction were assumed in the simulations to be the same as those used to predict the results of Experiment 1. For simplicity, the possibility of generalization, both from the training odors to the batting odor in Experiment 1 and from the training colors to gray in Experiment 2, was ignored.

In beginning to extend the model to deal with extinction, we turned from the stochastic simulation procedure used previously to a computational one, which not only brings a substantial gain in efficiency, but eliminates a component of simulation error contributed by the stochastic strategy. The old procedure was to determine the choice made by each stat-bee on each visit by consulting the choice function and a table of random numbers, and then to change its individual pattern of associative strengths accordingly. The new procedure is to compute average associative strengths on the assumption that the proportion of animals choosing each alternative on each visit is as given by the choice function. Although there is no simple way to demonstrate the mathematical equivalence of the two procedures, extensive empirical testing shows close agreement between their results in simulations both of simultaneous discrimination training and of extinction. Consider, for example, the simulated proportions of correct choice in the 16-visit training of Experiments 1, 2, and 3, which were arrived at both stochastically (in this




Figure 7. Simulated latency of response in extinction as a function of the relative associative strength of the stronger alternative. The scaling parameters are $U = 6$ and $M = 5$.

case, with 1,000 stat-bees) and computationally: For the confounded problem, the stochastic result was .929 and the computational result was .931; for the control problem, the stochastic result was .778 and the computational result was .779. A variety of other comparisons of this sort also have shown close agreement.

The fit of the model with $U = 6$, the embarrassingly unbiological $M = 5$, and $D\beta = .08$ to the pooled data of Experiments 2 and 3 is shown by Figure 8, in which the obtained and computationally simulated extinction curves are compared (root-mean-square deviation = 1.46 responses). The principal discrepancy is that response by the confounded animals to S+ is somewhat overpredicted. It should be noted that the model succeeds very well in translating the greater training experience of the control animals with S- into a stronger preference for S+ in extinction; since the *relative* associative strength of S+ is greater for the control animals, the number of responses to it is greater to begin with and the latency of response is shorter.

It may be well to emphasize that, according to continuity theory, the associative strength of S+ should be the same in the control and confounded animals, because they have had the same number of rewarded experiences with S+. To test the assumption that the differential response of the control and confounded animals to S+ was due only to the presence of S-, we did another experiment that




Figure 8. Simulation of the extinction performance of the confounded and control groups of Experiments 2 and 3 with $U = 6$, $M = 5$, and $D\beta = .08$. The large points show the pooled performance of the animals of the two experiments, and the small points show the simulated performance.

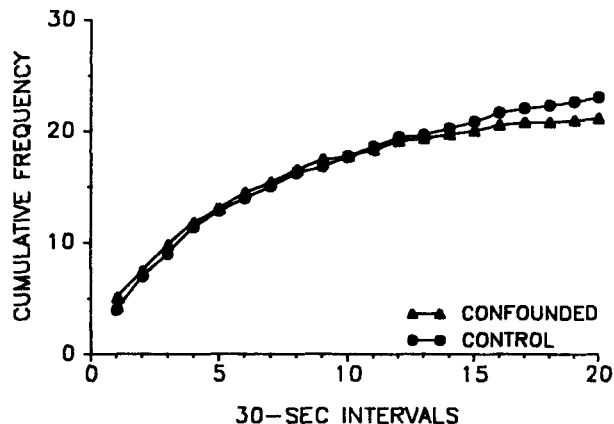


Figure 9. The course of extinction in a supplementary experiment like Experiment 2 in all respects except that only the formerly positive stimulus (S+) was present in extinction. The curves show mean cumulative number of responses over successive 30-sec intervals of the 10-min extinction period.

was like Experiment 2 in every respect, except that S+ alone was present in extinction and responses to it alone were recorded. Performance in the first stage closely resembled that shown in Figure 3, the control group making more errors and therefore having more experience with the negative color than the confounded group (median test, exact $p < .0008$). In the second stage, as Figure 9 shows, responding to S+ was very much the same in the two groups. Analysis of variance yields neither a significant groups effect ($F < 1$), nor a significant groups \times time interaction ($F < 1$). These results justify the emphasis in our extinction model on relative strength.

Despite the reasonably good fit demonstrated in Figure 8, there are some commanding reasons for dissatisfaction with the extinction feature of the model, which is, of course, only in a very early stage of development. For one thing, the assumption that latency is based entirely on relative associative strength must be wrong, if only because it provides no way of simulating extinction of response to a single target (as illustrated, for example, by Figure 9) or even to a pair of targets equal in associative strength. As a matter of fact, we can find almost (although not quite) as good a fit on the assumption that latency is determined entirely by the absolute strength of the *chosen* alternative (with r still determining the choice), the explanation of which may be that relative and absolute strength are highly correlated in the data. The latency equation will have to be reconsidered in relation to the results of a wider range of experiments. Another reason for dissatisfaction with the extinction feature of the model is that it makes no contact with the substantial overlearning effects which have been found both in tests with singly presented targets and in choice tests (Buchanan & Bitterman, 1988; Couvillon & Bitterman, 1984; Shinoda & Bitterman, 1987) and which suggest at the very least that a frustration component will have to be added to the latency equation when the conditions under which

those effects occur are more sharply defined. Clearly, we have a long way to go in the development of the extinction feature of the model, which may reasonably be assigned a high priority in our future work with honeybees.

Although the results of the present experiments seem to offer no serious challenge to the independence rule, there remain to be considered the results of some other experiments that do. We take up in turn: (1) the potentiation of odor by color reported in our 1982 paper, which has already been mentioned, (2) the reciprocal overshadowing of color and *position* reported in our 1983 paper, and (3) the overshadowing of a common stimulus in "true" discrimination (TD) as compared with pseudodiscrimination (PD) training reported in our 1983 paper.

1. The model in its present form cannot account for the potentiation, although potentiation need not contradict the independence rule if, as Durlach and Rescorla (1980) have proposed, it is a product of within-compound association. In our 1982 paper, we reported that—with the stimuli there employed—our bees seemed to form color-odor and odor-color associations at different rates, which, we suggested, might account for the fact that we found potentiation of odor by color but not of color by odor. The model must sooner or later be extended to take account of within-compound association, whether or not potentiation can be simulated in the process.

2. One difficulty presented by the 1983 color-position experiments is that the performance in training of the confounded groups was no better than that of the controls. A second difficulty (related to the first) is that the confounded animals, which showed poorer discrimination both of color and position (reciprocal overshadowing), had no less experience with the negative color than did the color-relevant controls and no less experience with the negative position than did the position-relevant controls. Evidence that color-position processing is somehow different from color-odor processing comes also from earlier work by Klosterhalfen, Fischer, and Bitterman (1978), who found better intradimensional than extradimensional transfer in color-position problems but not in color-odor problems. Position is not, of course, an intrinsic target property, but given only in relation to the structure of the visual field, about which honeybees are sometimes assumed to learn differently. The color and odor of a target are learned on the approach to it, according to Gould (1986), but landmark learning "takes place during departure . . . thereby violating the normal CS/US relationship" (p. 174). It was the opinion of Opfinger (1931), who is cited in support of this dramatic claim, that honeybees *do* learn about the environment of a target on the approach to it, and about the immediate environment of the target (which certainly would include our window frame) *only* on the approach to it. What Gould must have been referring to is Opfinger's opinion that there is at least *some* learning about characteristics of the more remote environment on the initial departure from a newly discovered feeding place. We find no hint in this work as to how the difference between our color-

odor and color-position results is to be understood. The only lead we have is that, as reported both in our 1978 and 1983 papers, the approach behavior itself is different (more direct) in position-relevant as compared with color- and odor-relevant problems, which sets the stage for a mediating-response interpretation.

3. Our 1983 TD-PD experiments were patterned after some well-known experiments with rats and rabbits by Wagner, Logan, Haberlandt, and Price (1968): Honeybees trained with orange-jasmine and orange-violet compounds (two odors and a common color) that were differentially reinforced (TD training) responded less to the unscented color in subsequent extinction tests than did honeybees trained with the same compounds nondifferentially reinforced (PD training), despite the fact that there was much more unreinforced experience with the color in the course of the PD training. These results (which resemble those for vertebrates) certainly would not be predicted from our model in its present form, although performance in a test with differential reinforcement following TD and PD training has already been simulated accurately (Couvillon & Bitterman, 1987). Here again, therefore, we may have an expression of the greater sensitivity of extinction tests to differences in associative strength, which in the TD-PD case would not be expected from the independence rule. Another possibility, of course, is that we are dealing with effects of inconsistent reinforcement per se that are specific to extinction.

Although substitution of the principle of shared associative strength (Rescorla & Wagner, 1972) for the independence principle might seem to be warranted by the TD-PD extinction results, that change alone would not make it possible to accommodate both the color-odor and the color-position overshadowing data: If our standard colors and odors are equally salient (as evidenced by the equal difficulty of the color-relevant and odor-relevant problems), and if the 1983 colors and positions are equally salient (as evidenced by the equal difficulty of the color-relevant and position-relevant problems), then the color-odor and color-position results should be the same. Nor would substitution of the shared-strength principle alone, of course, make it possible to account for the potentiation results. It is likely, in fact, that several different mechanisms not yet represented in the model are responsible for the troublesome data, and it is incumbent upon us to try to find out what they are. With precise quantitative simulation of all the results of such experiments as the criterion, we cannot be satisfied simply to explain away failures in terms of the narrowness of the model, but must seek constantly to expand it as necessary to make simulation possible.

REFERENCES

- BUCHANAN, G. MCC., & BITTERMAN, M. E. (1988). Learning in honeybees as a function of amount and frequency of reward. *Animal Learning & Behavior*, *16*, 247-255.
- BUSH, R. R., & MOSTELLER, F. (1951). A mathematical model for simple learning. *Psychological Review*, *58*, 313-323.

- COUVILLON, P. A., & BITTERMAN, M. E. (1980). Some phenomena of associative learning in honeybees. *Journal of Comparative & Physiological Psychology*, **94**, 878-885.
- 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. (1984). The overlearning-extinction effect and successive negative contrast in honeybees (*Apis mellifera*). *Journal of Comparative Psychology*, **98**, 100-109.
- 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., KLOSTERHALFEN, S., & BITTERMAN, M. E. (1983). Analysis of overshadowing in honeybees. *Journal of Comparative Psychology*, **97**, 154-166.
- DURLACH, P. J., & RESCORLA, R. A. (1980). Potentiation rather than overshadowing in flavor-aversion learning: An analysis in terms of within-compound associations. *Journal of Experimental Psychology: Animal Behavior Processes*, **94**, 878-885.
- FRISCH, K. VON. (1914). Der Farbensinn und Formensinn der Biene [The sense of color and form in bees]. *Zoologische Jahrbücher (Zoologie und Physiologie der Tiere)*, **35**, 1-182.
- GOULD, J. L. (1986). The biology of learning. *Annual Review of Psychology*, **37**, 163-192.
- KLOSTERHALFEN, S., FISCHER, W., & BITTERMAN, M. E. (1978). Modification of attention in honeybees. *Science*, **201**, 1241-1243.
- LOVEJOY, E., & RUSSELL, D. G. (1967). Suppression of learning about a hard cue by the presence of an easy cue. *Psychonomic Science*, **8**, 365-366.
- OPFINGER, E. (1931). Über die Orientierung der Biene an der Futterquelle [On the orientation of bees to a food source]. *Zeitschrift für Vergleichende Physiologie*, **15**, 431-487.
- 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: Variations in the effectiveness of reinforcement and non-reinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64-99). New York: Appleton-Century-Crofts.
- SHINODA, A., & BITTERMAN, M. E. (1987). Analysis of the overlearning-extinction effect in honeybees. *Animal Learning & Behavior*, **15**, 93-96.
- SPEAR, N. E., & PAVLIK, W. B. (1966). Percentage of reinforcement and reward magnitude effects in a T maze within and between subjects. *Journal of Experimental Psychology*, **71**, 521-528.
- SUTHERLAND, N. S., & ANDELMAN, L. (1967). Learning with one and two cues. *Psychonomic Science*, **7**, 107-108.
- SUTHERLAND, N. S., & MACKINTOSH, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- WAGNER, A. R. (1969). Incidental stimuli and discrimination learning. In R. M. Gilbert & M. S. Halliday (Eds.), *Animal discrimination learning* (pp. 83-111). New York: Academic Press.
- WAGNER, A. R., LOGAN, F. A., HABERLANDT, K., & PRICE, T. (1968). Stimulus selection in animal discrimination learning. *Journal of Experimental Psychology*, **76**, 171-180.
- 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 April 20, 1988;
revision accepted for publication October 31, 1988.)