Discrimination of color-odor compounds by honeybees: Tests of a continuity model

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

In experiments previously reported, individual honeybees were trained in a variety of problems to choose between two visually identical but differently scented targets, one or the other of which contained sucrose solution. The results could be simulated accurately with simple equations for computing changes in associative strength produced by reinforcement or nonreinforcement and for predicting choice on the basis of relative strength. In the present experiments, the targets used differed in color as well as in odor, and the animals were trained in a variety of problems with color-odor compounds. Contrary to expectation, the new results could be simulated accurately with the same equations as before on the further assumption that the components of a compound gain and lose associative strength independently (independence rule) and that the associative strength of a compound is equal to the sum of the strengths of its components (summation rule).

In recent experiments with honeybees, we have studied the discriminative learning of individual foragers shuttling back and forth between the hive and the sill of an open laboratory window (Couvillon & Bitterman, 1985, 1986). On each visit, two differently scented gray targets were presented, one of them containing a drop of 50% sucrose solution from which there was feeding to repletion; if the target without sucrose was chosen, immediate correction was permitted. The performance of the animals in a variety of problems (including successive reversal and ambiguous-cue problems) was simulated quantitatively with simple equations for computing the strength of the association of each target with sucrose and for predicting choice on the basis of differential strength. Our strategy was first to look for the rule of choice in problems requiring only the most primitive associative principles, and then to extend those principles in whatever way might be necessary to deal with the array of seemingly more complex associative phenomena discovered in previous work on compound conditioning (Couvillon & Bitterman, 1982; Couvillon, Klosterhalfen, & Bitterman, 1983). The strategy was adopted on the perfectly testable assumption that the choice rule identified in the first experiments would continue to apply as the associative features of the theory were elaborated. As explained at the outset (Couvillon & Bitterman, 1985), our primary interest has been in the associative learning of honeybees. We have been interested in choice secondarily, as an index of learning that is more reliable than latency of response to a single target and more efficient than resistance to extinction.

Change in associative strength is expressed by a linear

equation (Bush & Mosteller, 1951) in the now much more familiar notation of Rescorla and Wagner (1972):

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

with V_A representing the associative strength of alternative A , ΔV_A the change in strength produced by reinforcement or nonreinforcement, λ the asymptotic associative strength, and β the learning-rate parameter. Rescorla and Wagner's salience parameter, α , does not appear in the equation on the assumption that the odors used are equally salient; although associative strength may vary with α , the prediction of choice, which (as will be explained in a moment) is based on relative strength, does not. The value of λ is taken as 1 when there is reinforcement (on the assumption that the magnitude of reinforcement is maximal), 0 when there is no reinforcement, and provision is made in the model for the possibility that the value of β is different for reinforcement and nonreinforcement ($U\beta = \text{incremental or up}\beta$, $D\beta = \text{decremental or down}\beta$). When choice of A is reinforced, $\Delta V_A = U\beta(1 - V_A)$; when it is not reinforced, $\Delta V_A = D\beta(0 - V_A) = -D\beta \cdot V_A$. The associative strength of A relative to that of alternative B is designated as r_A and is expressed by the equation

$$r_A = V_A / (V_A + V_B),$$

where V_B is the associative strength of B . A simple computer program evaluates r on the basis of the initial values of V_A and V_B (which are assumed to be greater than 0 in consequence of pretraining with the two stimuli), dictates a choice, and then calculates the resulting change in associative strength. When A is the reinforced alternative on a given visit, V_A alone is changed (incremented) if A is chosen initially; if B is chosen initially, V_B is decremented and then V_A is incremented. Choice is dictated by any one of a family of power functions that may be incorporated in the program, that is, functions relating P_A , the probability of choosing A , to r_A .

This research was supported by Grant BNS 8317051 from the National Science Foundation. Helpful advice from William T. Woodard is gratefully acknowledged, as is the assistance of Hien Lien, Philip Okamoto, Susana Pulawa, Todd Wade, and Mei Yee Wong in collecting the data. Requests for reprints should be addressed to the Békésy Laboratory of Neurobiology, 1993 East-West Road, Honolulu, HI 96822.

To generate the choice functions, we use the scaling equations

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

for $r \geq .5$, and

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

for $r < .5$, which, with $s=.5$, yield a set of sigmoid curves passing through the points (0,0), (.5,.5), and (1,1). When $K > 1$, P changes slowly in the region of $r=.5$ and progressively faster at the extremes; when $K < 1$, the opposite is true; when $K=1$, the relationship between P and r is linear, as shown in Figure 1, where the curve for $K=1$, $s=.5$ is included to provide a frame of reference. (The curves of Figure 1 are plotted only for $r \geq .5$, because all the functions are perfectly symmetrical for $r < .5$). When $s > .5$, with P truncated to $0 \leq P \leq 1$, the functions are steeper, as shown by the curves plotted in Figure 1; the truncation implies, of course, that the choice measure is insensitive to differences in extreme values of r . For each function considered, we set out to find the initial value of $V_A = V_B$ and the values of $U\beta$ and $D\beta$ that would yield the minimal root-mean-square (RMS) deviation of predicted from obtained choices, and we compared the functions in terms of their minimal RMS deviations. Good fits were achieved with the cluster of functions shown in Figure 1, of which $K=.75$, $s=.833$ was marginally the best. In all cases, it was necessary to assume rapid learning—initial values of V_A and V_B (generated by reinforcement in pretraining) on the order of .3 or .4, and β values on the order of .20 or .25.

We were able in our previous experiments to simulate choice rather accurately on the basis of the differentially reinforced odors alone, ignoring such perfectly dis-

crimable properties of the targets as color (which was the same gray for both) or position (which was non-differentially reinforced), and it may be well now to consider why that should be. It is conceivable that color and position were overshadowed by odor, which would require a departure from the simple contiguity-frequency principle of association incorporated in the model. Another possibility is that sensory components other than odor did, in fact, acquire appreciable associative strength, V_C , adding equally to the strengths of the differentially reinforced alternatives, and that by ignoring them we exaggerated r ; it can be shown that, when $V_A > V_B$, $V_A/(V_A + V_B) > (V_A + V_C)/(V_A + V_B + 2V_C)$, and the opposite when $V_A < V_B$. If r was exaggerated (its absolute deviation from .5 unduly large), the error may have been compensated for in the fitting process by a flattening of the choice function, which implies that, contrary to our expectations, complicating the associative features of the model may indeed change our picture of the choice function.

In the present experiments, we set out deliberately to do multicomponent modeling, introducing two colors as discriminable as the odors and differentially reinforcing them along with the odors in ways designed to promote interaction and to facilitate its analysis. We continued to balance position, but otherwise, for the time being, to ignore it on the basis of earlier evidence that suggested that it might require special treatment (Klosterhalfen, Fischer, & Bitterman, 1978); position is not, after all, an intrinsic target property, but is given only in relationship to the structure of the visual field. Our strategy was to begin with the basic "continuity" assumptions (Hull, 1929; Spence, 1936)—that the associative strength of each component of a compound stimulus changes independently with reinforcement and nonreinforcement (*independence rule*), and that the associative strength of a compound is equal to the sum of the strengths of its components (*summation rule*)—to try to identify points in the data at which those assumptions might appear to be inadequate, and then to think about how they might have to be supplemented or changed. The independence assumption is implemented in the model by separately incrementing or decrementing each component of a reinforced or nonreinforced compound. Where A and B are the two odors, X and Y the two colors, and the animal is given a choice of two color-odor compounds, AX and BY , with AX reinforced, V_A and V_X are incremented independently if the animal goes first to AX [$\Delta V_A = U\beta(1 - V_A)$ and $\Delta V_X = U\beta(1 - V_X)$]; if the animal goes first to BY , V_B and V_Y are decremented [$\Delta V_B = -D\beta \cdot V_B$ and $\Delta V_Y = -D\beta \cdot V_Y$] before V_A and V_X are incremented. (It may be well to emphasize that, despite the Rescorla-Wagner notation, shared associative strength is *not* assumed.) By the summation rule, $V_{AX} = V_A + V_X$ and $V_{BY} = V_B + V_Y$. Choice between AX and BY is simulated by computing $r_{AX} = V_{AX}/(V_{AX} + V_{BY})$ and reading P_{AX} from the choice function.

Earlier experiments with resistance to extinction as the measure of associative strength had given abundant evidence of the usefulness of the summation rule (Couvillon

Figure 1. Five choice functions yielding good fits to the data of previous experiments (Couvillon & Bitterman, 1985, 1986) with $K=1$, $s=.5$ included for reference.

& Bitterman, 1980), but evidence also of compound uniqueness (Couvillon & Bitterman, 1982), which has sometimes been held to contradict it, and evidence of overshadowing (Couvillon & Bitterman, 1980, 1982; Couvillon et al., 1983), which may seem incompatible with the independence rule. We had every expectation, in fact, that our new experiments soon would produce results that could not be simulated without departure from these rules. Our only question was about what changes would be required.

METHOD

Subjects

The subjects were honeybees (*Apis mellifera*) that were 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 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 bee's second return to the laboratory of its own accord.

The targets were covered petri dishes of clear plastic, 5.5 cm in diameter. In each cover, to which a disk of green or blue plastic of the same diameter as the cover could be cemented, eight equally spaced holes, .5 cm in diameter, were drilled at the outer circumference. The dishes themselves contained wads of cotton that either were impregnated with the scent of geraniol or the scent of peppermint or were unscented. In all, eight sets of targets were used in the training: green-unscented, blue-unscented, gray-geraniol, gray-peppermint, green-geraniol, blue-geraniol, green-peppermint, and blue-peppermint. The covers 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 were three additional sets of targets, one with gray covers and labeled with both odors, and two with covers half-green and half-blue, one unscented and the other labeled with both odors. The pretraining procedure in every case was to place the animal carried in from the feeding station onto a target that presented all of the stimuli it would encounter in its subsequent training.

On each training visit, the animal found two targets set 10 cm apart on a line parallel to the outer edge of the window sill. One of the targets contained a large drop of 50% sucrose solution, and the other contained a large drop of tap water that was distinguishable from the sucrose only by taste. The target first chosen by the subject 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. The various training problems will be described as the results are presented. Each subject was trained only in a single ex-

periment, with no fewer than 8 subjects per problem and sometimes considerably more.

RESULTS

Color Versus Odor

In modeling our previous results for odors, we were able to ignore their salience, which was assumed to be equal. Fortunately, we can continue to ignore salience in our work with two dimensions, because the results of three different experiments give no reason to assume a difference in the salience of the odors and colors used, as indexed either by rate of learning or by competition for control of choice.

One-dimensional training. Shown in Figure 2 is the performance of two groups of 8 subjects each that were trained for 10 visits either with unscented green versus blue targets or with geraniol- versus peppermint-scented gray targets. For half the color animals, green was reinforced, and for the rest, blue was reinforced; for half the odor animals, geraniol was reinforced, and for the rest, peppermint was reinforced; for both groups, the reinforced alternative was presented equally often on the right and on the left in quasi-random sequence. (These balancing procedures were employed routinely in the various training problems and will not in general be referred to again.) It may be seen in Figure 2, which is plotted in terms of the proportion of animals in each group choosing correctly on each visit, that the two colors were discriminated as readily as the two odors. The mean number of errors was 1.50 in the color problem and 1.75 in the odor problem (for the median test, Fisher's exact $p = .24$).

Two-dimensional training with one relevant dimension. Plotted in Figure 3 is the performance of two groups of 8 animals each that were trained for 10 visits to choose between targets that differed both in color (green vs. blue) and in odor (geraniol vs. peppermint). For the color-relevant animals, color was differentially reinforced and odor was nondifferentially reinforced (e.g., green-peppermint positive vs. blue-geraniol negative on half the visits and green-geraniol positive vs. blue-peppermint

Figure 2. The performance of two groups of animals in a one-dimensional problem, one trained with color and the other with odor.

Figure 3. The performance of two groups of animals in a two-dimensional problem, one trained with color relevant and the other with odor relevant.

negative on the rest). For the odor-relevant animals, odor was differentially reinforced and color was nondifferentially reinforced (e.g., green-peppermint positive vs. blue-geraniol negative and blue-peppermint positive vs. green-geraniol negative). The curves show that learning to respond in terms of the colors rather than the odors was as rapid as learning to respond in terms of the odors rather than the colors. The mean number of errors was 2.38 in the color-relevant problem and 2.25 in the odor-relevant problem (median test, exact $p = .34$).

Transfer after confounded training. In another experiment, 64 animals were trained for either 6 or 20 visits on a confounded problem, with one color-odor compound (e.g., green-geraniol) reinforced and the other (blue-peppermint) unreinforced. Then half the 6-visit animals and half the 20-visit animals were trained for 10 visits with color relevant and odor irrelevant, and the rest were trained with odor relevant and color irrelevant (a 2×2 design with 16 animals in each group). The idea here was that performance in the second problem might indicate that the animals had relied primarily on color rather than odor, or on odor rather than color, in the first problem. Consider an animal that is trained in the first problem with green-geraniol reinforced and blue-peppermint unreinforced and must choose between green-peppermint and blue-geraniol on the first visit of the second problem: If it has learned primarily about odor in the first problem and if odor is relevant in the second (blue-geraniol reinforced), it will choose correctly and have little difficulty on subsequent visits, but if it has learned primarily about color rather than odor in the first problem, or if color rather than odor is relevant in the second problem, it will tend to choose incorrectly. The experiment was done with two different amounts of training in the first problem on the chance that the animals might learn first about one of the dimensions and only later about the other, or even that they might learn first about both dimensions and come only later to attend primarily to one of them.

In Figure 4, the results for all animals trained in the color-relevant problem are compared with those for all animals trained in the odor-relevant problem. The results

for animals with 6 training visits in the first problem are pooled with those for animals with 20 visits, because amount of training in the first problem made no difference (median test, exact $p = .50$). The two curves show much the same pattern of choice in the color-relevant and odor-relevant problems. On the first, third, and fourth visits, in which the animals encountered compounds different from those experienced in the prior confounded training, the color-relevant performance was somewhat better than the odor-relevant performance, suggesting more control by color, although the difference is not significant. The mean number of errors was 1.44 in the color-relevant problem and 1.63 in the odor-relevant problem (median test, exact $p = .40$). It seems reasonable, then, to proceed with our simulations on the assumption of equal salience for all of the stimuli employed.

The ease with which we have been able to find a pair of colors as discriminable as the odors may be surprising in view of general assertions in the literature that bees learn about odor more quickly than they learn about color (e.g., Gould, 1982, p. 255; Lindauer, 1971, p. 44). It should be obvious, however, that the outcome of such experiments will depend on the particular colors and odors employed, their intensity and similarity, and the manner of their presentation. Fewer errors will be made in training with a pair of dissimilar colors or odors than with a pair of similar ones, which means that the similarity of the alternatives must be equated in any comparison of the rate of color and odor learning. The outcome of training in two-dimensional problems with one dimension relevant and the other irrelevant must depend as well on the relative intensity of the stimuli quite apart from modality. In our first comparison of color-relevant and odor-relevant problems (Klosterhalfen et al., 1978), which actually involved (among others) the same colored plastics used here, the odor discrimination was somewhat more rapid than the color discrimination, primarily, we suspect, because the odors were more intense; not only were the odor wells larger, but the animals were trained in an enclosure through which there was less air movement than there was

Figure 4. The performance of two groups of animals in a two-dimensional problem, one trained with color relevant and the other with odor relevant, after both had been trained in a confounded problem.

at an open laboratory window. Experiments that would justify generalizations such as Gould's and Lindauer's simply have not been done, and it is not clear at all, in fact, just how one would do them properly.

One- Versus Two-Dimensional Problems

To begin to get some idea of how well our continuity model is able to deal with multicomponent data, we compared the performance of separate groups of bees in one-dimensional problems and in four different two-dimensional problems—a confounded problem, two color-relevant problems with odors that varied either within or between pairs, and a color-relevant problem with a constant irrelevant odor.

Confounded compared with one-dimensional training. Plotted in Figure 5 is the performance on visits 1–10 of the 32 animals, previously referred to, that were trained for 20 visits in a confounded problem with the colors green and blue and the odors geraniol and peppermint. That performance is compared with the performance of the 16 animals, also previously referred to, that were trained in one-dimensional problems, either with unscented green and blue targets or with geraniol- and peppermint-scented gray targets (the two curves of Figure 2 pooled). There is no clear evidence of summation—the mean number of errors was 1.34 in the confounded problem and 1.63 in the one-dimensional problem (median test, exact $p = .17$). This outcome is predicted by the continuity model on the assumption of equal salience. In a discrimination between $AX+$ and $BY-$, reinforcing AX equally increments V_A and V_X , while nonreinforced experience with BY equally decrements V_B and V_Y . With $V_X = V_A$ and $V_Y = V_B$, $r_{AX} = V_{AX}/(V_{AX} + V_{BY}) = (V_A + V_X)/(V_A + V_X + V_B + V_Y) = 2V_A/2(V_A + V_B) = r_A$. X and Y contribute nothing, then, to the ease of discrimination, although the animals do learn about them. Where X and Y are less salient than A and B , performance in the confounded problem will be better than in $X+Y-$ (according to the continuity model), although no better than in $A+B-$.

That the animals trained in the confounded problem learned both about color and about odor is suggested by

Figure 5. The performance of two groups of animals, one trained in a confounded problem and the other in a one-dimensional problem with color or odor.

the fact that they shifted as readily to the color-relevant problem as to the odor-relevant problem (Figure 4). Nevertheless, we analyzed the data in one other way in order to examine the possibility that individual animals were attending primarily to one or the other dimension in the confounded training. Of 64 subjects trained in the confounded problem, 30 chose correctly on the first visit of the second problem, and their performance in the next nine visits was compared with that of the 34 subjects that chose incorrectly. If correct choice on the first visit reflected a set for the relevant dimension and incorrect choice reflected a set for the irrelevant dimension, better subsequent performance would be expected of the animals that had chosen correctly, but no significant difference was found. The mean number of errors in visits 2–10 was .93 for the animals whose first choice was correct and 1.03 for the animals whose first choice was incorrect (median test, exact $p = .47$).

Modes of presenting irrelevant stimuli. We turn now to a comparison of the performance of three groups of bees trained for 10 visits to discriminate between green and blue in the presence of nondifferentially reinforced peppermint and geraniol, but with the odors presented in different ways. For a group of 8 irrelevant-within animals, the alternative compounds always differed in odor (e.g., green-peppermint positive vs. blue-geraniol negative and green-geraniol positive vs. blue-peppermint negative). For 8 irrelevant-between animals, both alternatives were scented with peppermint on half the visits and with geraniol on the rest (green-peppermint vs. blue-peppermint and green-geraniol vs. blue-geraniol). For 8 constant-irrelevant animals, only one or the other of the two odors was used throughout (green-peppermint vs. blue-peppermint for half the animals and green-geraniol vs. blue-geraniol for the rest). One of our interests here was in the possibility that the irrelevant-within problem might be more difficult than the irrelevant-between problem because of more successful competition for attention by odor in the within case; there is some evidence of such an effect in rats (Wortz & Bitterman, 1953), although it may have quite another explanation. We were interested in the constant-irrelevant problem because in some earlier work (Couvillon et al., 1983) it had seemed to be more difficult than the irrelevant-between problem.

Performance in the three problems is compared in Figure 6, which gives no indication of differential difficulty. The mean number of errors was 2.37 for within, 2.69 for between, and 2.06 for constant (Kruskal-Wallis $H = 2.02$, $p > .30$). Equal performance in the three problems is, in fact, predicted by the continuity model. With A and B (the odors) nondifferentially reinforced in the irrelevant-within and irrelevant-between problems, $V_A = V_B$, and r_{AX} is the same in all problems— $V_{AX}/(V_{AX} + V_{BY}) = V_{AX}/(V_{AX} + V_{AY})$. Why the constant-irrelevant problem should have been relatively difficult under the quite different conditions of our previous experiment, we are unable to say.

Figure 6. The performance of three groups of animals trained in two-dimensional problems with color relevant and odor either irrelevant-within, irrelevant-between, or constant.

Confounded training versus two-dimensional training with one relevant dimension. The continuity model clearly suggests that performance will be better in the confounded color-odor problem than in problems with one dimension relevant and the other irrelevant. We have already shown that r_{AX} in the confounded problem ($AX+BY-$) is equal to r_A in the one-dimensional problem ($A+B-$). For the two-dimensional problem with one relevant and one irrelevant dimension (e.g., $AX+BY-$ and $AY+BX-$), $r_{AX} = (V_A + V_X)/(V_A + V_X + V_B + V_Y)$. With X and Y nondifferentially reinforced ($V_X = V_Y$), $r_{AX} = (V_A + V_X)/(V_A + V_B + 2V_X)$, and, since $V_B < V_A$, $r_{AX} < r_A$. This deduction from the continuity model is confirmed by Figure 7, in which the confounded curve of Figure 5 is compared with a color-relevant curve based on the pooled data of Figure 6. The mean number of errors is 1.34 for the confounded problem and 2.38 for the problems with only one of the two dimensions relevant, a significant difference by the median test (exact $p = .02$). In our previous work, too (Couvillon et al., 1983), performance was better with color and odor confounded than with color relevant and odor irrelevant or with odor relevant and color irrelevant.

Other Two-Dimensional Problems

Here we consider performance in several different two-dimensional problems chosen in the hope of defining the limitations of the continuity model. One new problem was a pseudodiscrimination (neither the colors nor the odors differentially reinforced), and another was a dimensional shift (color relevant and odor irrelevant after training with odor relevant and color irrelevant). We also used three problems in each of which the colors were reinforced on a 75:25 schedule (one color on three of four visits and the other color on the rest) with the odors reinforced on a 50:50, 75:25, or 100:0 schedule.

Transfer after true discrimination and pseudodiscrimination. Two groups of 16 animals each were trained for 10 visits with color relevant and odor irrelevant. For one group, this training followed 16 visits with odor relevant and color irrelevant (dimensional shift). For

the second group, it followed 16 pseudodiscrimination visits on which both color and odor were nondifferentially reinforced (green-peppermint vs. blue-geraniol on half the visits and green-geraniol vs. blue-peppermint on the rest, with each compound equally often positive and negative). We were interested here in trying to model pseudodiscrimination as we had earlier modeled nondifferential reinforcement of odor alone (Couvillon & Bitterman, 1985), as well as in the possibility that the pseudodiscrimination experience might have a deleterious effect, either general or specific to the stimuli employed, on subsequent performance (learned irrelevance). We were interested also in how well we might be able to model dimensional shifting without the assumption of attentional modification (Sutherland & Mackintosh, 1971).

The performance of the two groups in the second stage of training (with color relevant and odor irrelevant) is plotted in Figure 8. The curve for the pseudodiscrimination animals looks very much like that for naive animals trained from the outset on the same problem (Figure 3), which is to say that the pseudodiscrimination experience had no effect (median test, exact $p = .25$). If the nondifferentially reinforced odors and colors acquired substantial associative strength in the first stage of training, it was not sufficient, given the rate of learning and the shape of the choice function, to produce measurable impairment of performance in the second stage. The marked early fluctuation in the curve for the dimensional-shift animals is readily understandable in terms of the training schedule. On the first and third visits of the second stage, the now-positive color was presented with the formerly negative odor and the now-negative color was presented with the formerly positive odor; on the second visit, the now-positive color was presented with the formerly positive odor and the now-negative color was presented with the formerly negative odor. Results of this sort, reported for rats (Bitterman & Coate, 1950), painted turtles (Graf & Tighe, 1971), and goldfish (Tennant & Bitterman, 1973), are easily understood in terms of continuity theory.

75:25 color problems with odor on various schedules. Three groups of 12 animals each were trained for

Figure 7. The performance of two groups of animals in two-dimensional problems, one confounded and the other with color relevant (one-relevant).

Figure 8. The performance of two groups of animals in a two-dimensional problem with color relevant, one after pseudodiscrimination (PD) and the other after training with odor relevant (dimensional shift).

16 visits with the colors reinforced on a 75:25 schedule—one color positive on 12 visits and the other positive on the remaining visits, in quasi-random sequence. For one group, the odors were nondifferentially reinforced (a 50:50 schedule); for the second, the odors, like the colors, were reinforced on a 75:25 schedule; and for the third, one of the two odors was consistently reinforced (100:0). We had no qualitative prediction of what the performance in each case would be like, but were interested, to begin with, only in whether it could be simulated with reasonable accuracy on the basis of the continuity assumptions. The obtained and simulated results are presented below following a description of the simulation procedure.

Simulations

Visit-by-visit performance in each of the following problems was simulated:

1. The one-dimensional problem: $A+B-$.
2. Three two-dimensional problems with one relevant dimension: $X+Y-$ with A and B irrelevant within pairs ($AX+BY-$, $BX+AY-$), irrelevant between pairs ($AX+AY-$, $BX+BY-$), or constant ($AX+AY-$).
3. The confounded problem: $AX+BY-$.
4. Transfer from the confounded problem to the two-dimensional problem with one dimension relevant and the other varying within pairs.
5. The two-dimensional pseudodiscrimination problem: $AX+BY-$, $BY+AX-$, $AY+BX-$, $BX+AY-$.
6. Transfer from the pseudodiscrimination problem to the two-dimensional problem with one dimension relevant and the other varying within pairs.
7. Transfer from the two-dimensional problem with one dimension relevant and the other varying within pairs to the same problem with relevant and irrelevant dimensions interchanged (dimensional shift).
8. The two-dimensional problem with one dimension on a 75:25 reinforcement schedule and the other on a 50:50 schedule.
9. The two-dimensional problem with one dimension on a 75:25 schedule and the other on a 75:25 schedule.

10. The two-dimensional problem with one dimension on a 75:25 schedule and the other on a 100:0 schedule.

The simulations were done with each of the five choice functions that yielded the best fits to the combined data of our previous one-dimensional experiments (Figure 1). Simulations were done also with $K=1$, $s=.5$, which did not yield a good fit to the earlier data, simply to be sure that the shape of the choice function continued to matter. For each function, there were 100 simulations with four values of initial $V_A=V_B=V_X=V_Y$ (.1, .2, .3, .4), five values of $U\beta$ (.15, .20, .25, .30, .35), and five values of $D\beta$ (the same as for $U\beta$) combined factorially. Each simulation was done with 100 stat bees trained exactly as the real bees had been trained in each of the 10 problems. A simple computer program evaluated r on the basis of the initial V values, dictated response on the basis of the choice function, and then calculated the resulting changes in associative strength. Goodness of fit was measured in terms of the *RMS* deviation of the simulated from the obtained proportions of choices, one computation on the basis of untransformed proportions and another after arcsin transformation, each proportion, P , being replaced by an angle, ϕ , measured in degrees, according to the equation, $\phi=(180/\pi)\sin^{-1}P^{.5}$. The *RMS* scores based on untransformed proportions have more intuitive meaning but—because the variance of a proportion decreases with distance from .5—may not give sufficient weight to small deviations at the extremes of the scale.

Each computation was based on 108 deviations selected to provide a conservative indication of goodness of fit. Since the theory was developed originally to deal with the one-dimensional problem, and since (as has already been noted) predictions for the one-dimensional problem and for the confounded problem are mathematically identical with equally salient components, the fit to the one-dimensional problem (1) was not considered in the computations. Furthermore, since the predictions for the three different treatments of the irrelevant dimension in problems with one relevant and one irrelevant dimension (2) are (again as already noted) mathematically identical, the computations were based on deviations of the combined simulations for the three treatments from the combined performance of the groups trained with the three treatments (Figure 6); to do otherwise might be to exaggerate the goodness of fit. The 108 choices used in the computations consisted of the following: (a) visits 2–10 (9 choices) with combined treatments in (2)—the first visit was excluded because any deviation from .5 either simulated or in the actual performance must be sampling error; (b) visits 2–10 (9 choices) in (3); (c) visits 1–10 (10 choices) in (4); (d) visits 2–16 (15 choices) in (5); (e) visits 1–10 (10 choices) in (6); (f) visits 1–10 (10 choices) in (7); (g) visits 2–16 (15 visits) in (8); (h) visits 2–16 (15 choices) in (9); and (i) visits 2–16 (15 choices) in (10).

The five choice functions that had yielded good fits in our previous work with one-dimensional problems con-

tinued to yield good fits here. The best were for $K=.75$, $s=.833$ (RMS for untransformed proportions = .114) and for $K=.5$, $s=.625$ (RMS = .115), both with the parameters $V_A=V_B=V_x=V_y=.4$, $U\beta=.15$, $D\beta=.25$ —very much the same as those that yielded the best fits to the one-dimensional data. For the other three functions, the RMS values were .127 or .128. For $K=1$, $s=.5$, by contrast, the minimal RMS was .163. The arcsin transformation made little difference: Again, the best fits were for $K=.75$, $s=.833$ and for $K=.5$, $s=.625$ (arcsin RMS = 8.1° in both cases), while the arcsin RMS values for the other three choice functions ranged from 9.2° to 9.5° ; for $K=1$, $s=.5$, it was 12.0° . All the simulations shown in the figures that follow are for $K=.75$, $s=.833$; initial V values = .4, $U\beta=.15$, $D\beta=.25$.

In Figure 9, simulated performance in a one-dimensional problem is compared with simulated performance in a confounded problem. No summation is predicted, and, as shown in Figure 5, no summation was obtained.

Figure 10 shows simulated performance in two-dimensional problems with one relevant dimension and the irrelevant stimuli presented in three different ways (within, between, or constant). No differences are predicted, and as may be seen in Figure 6, no differences were obtained.

In Figure 11, simulated performance in a confounded problem is compared with simulated performance in a problem with one relevant and one irrelevant dimension.

Figure 9. Simulated performance in confounded and one-dimensional problems.

Figure 10. Simulated performance in two-dimensional problems with one dimension relevant and the other either irrelevant-within, irrelevant-between, or constant.

Figure 11. Simulated performance in a confounded problem and in a two-dimensional problem with one relevant dimension.

Figure 12. Obtained (large circles) and simulated (small circles) performance in a two-dimensional problem with one relevant dimension after training in a confounded problem.

The prediction that the confounded performance will be better is confirmed by the empirical results shown in Figure 7, although the obtained difference seems a bit smaller than the predicted difference.

Figure 12 compares obtained and simulated performance with one relevant and one irrelevant dimension after confounded training. The obtained performance is represented by pooling the results for the color-relevant and odor-relevant animals (Figure 4), which are not significantly different and should not be, according to the theory, if color and odor are equally salient.

Obtained and simulated pseudodiscrimination performance are plotted in Figure 13. The fluctuations in the obtained performance are not simulated very closely, but the discrepancies do not seem to be due to systematic error; it should be noted that we are working here with proportions in the region of .5. Figure 14 shows the subsequent performance, obtained and simulated, in training with one dimension relevant and the other irrelevant. No detrimental effect of the pseudodiscrimination experience on the subsequent performance is either predicted (compare with the one-relevant curve of Figure 11) or obtained (compare with the one-relevant curve of Figure 7).

The obtained and simulated dimensional shifts (from odor-relevant to color-relevant) are shown in Figure 15. If there is any meaningful discrepancy here, it is that the real bees accomplish the shift somewhat more rapidly than the stat bees.

Figures 16–18 show obtained and simulated performance in the three problems with the colors reinforced on a 75:25 schedule and the odors on a 50:50 schedule (Figure 16), a 75:25 schedule (Figure 17), or a 100:0 schedule (Figure 18). The wide fluctuations in obtained proportions of correct choice are tracked satisfactorily by the model.

It should be noted that goodness of fit is limited by random error in the simulation process (σ_s) and in the data (σ_d), as well as by model error (σ_m). To estimate σ_s , we did a second simulation for 100 stat bees with $K=.75$, $s=.833$, $V_A=V_B=.4$, $U\beta=.15$, $D\beta=.25$. On the basis of the *RMS* deviation of the two sets of simulated values from each other, which is .056, σ_s may be estimated from the equation

$$\sigma_s^2 + \sigma_s^2 = (.056)^2$$

as .040. On the assumption of independence, it is possible also to estimate σ_{dm} (σ_d and σ_m combined) from the equation

$$\sigma_s^2 + \sigma_{dm}^2 = (.114)^2$$

as .106. As we have pointed out before (Couvillon & Bitterman, 1986), the reliability of the data is a serious limitation in this work. The number of stat bees can be increased easily (to minimize σ_s), but the number of real

Figure 13. Obtained (large circles) and simulated (small circles) performance in pseudodiscrimination training.

Figure 14. Obtained (large circles) and simulated (small circles) performance in a two-dimensional problem with one relevant dimension after training in a pseudodiscrimination problem.

Figure 15. Obtained (large circles) and simulated (small circles) performance in a dimensional shift problem.

Figure 16. Obtained (large circles) and simulated (small circles) performance in a two-dimensional problem with one dimension reinforced on a 75:25 schedule and the other on 50:50.

bees can be increased substantially (to minimize σ_d) only at considerable expense.

DISCUSSION

The performance of our honeybees in a variety of problems requiring the discrimination of color-odor compounds gives no qualitative indication of noncontinuity and can, in fact, be simulated quantitatively with considerable accuracy on the basis of a simple component model incorporating the independence and summation rules. The problems studied thus far do not, of course, maximally constrain the model, and it might even be anticipated that honeybees can solve problems (such as the conditional discrimination, $AX+BX-$ and $BY+AY-$) which cannot in principle be handled by the model in its present form, but our strategy has been to approach the potential difficulties gradually. In considering new directions, it is helpful to inquire into the strengths as well as the weaknesses of the model that is taken as the point of departure, and we know now that the strengths are substantial.

We certainly can say that the hypothesis about the general shape of the choice function based on our previous work with one-dimensional problems is supported by the results for two-dimensional problems. Furthermore, the variety of those problems and the good fits achieved

Figure 17. Obtained (large circles) and simulated (small circles) performance in a two-dimensional problem with both dimensions reinforced on 75:25 schedules.

Figure 18. Obtained (large circles) and simulated (small circles) performance in a two-dimensional problem with one dimension reinforced on a 75:25 schedule and the other on 100:0.

suggest that the associative features of the model may not be intrinsically wrong, but in need, at most perhaps, of some supplementation and refinement. It is conceivable, for example, that conditional discrimination and other configural phenomena can be modeled successfully on the principle of "afferent neural interaction," which was proposed by Hull (1943), either, as he did, with a generalization function that ultimately will be necessary in any case or with compound-unique components (Rescorla, 1972; Whitlow & Wagner, 1972). Nor would change in the associative features of the model inevitably be required by the discovery of some sort of attentional or salience-modulating mechanism, which might easily be grafted onto the existing structure. It should be noted, however, that no indication of such a mechanism appeared in our early work on dimensional transfer in color-odor problems (Klosterhalfen et al., 1978). The most serious deficiency of the present model may lie in its failure to provide for the development of inhibition, which has already been found in extinction tests (Couvillon & Bitterman, 1980), but to which the present experiments were not designed to be sensitive. It is when we come to address the problem of inhibition more directly that we may be called upon to make some fundamental revisions of the model, as, for example, to substitute increment in inhibition for reduction in V , with all of the attendant difficulties (Woodard & Bitterman, 1976), or, as did Rescorla and Wagner

(1972), to permit V to take on negative values, itself by no means a trouble-free alternative. Given the relative simplicity of the animal with which we are dealing and the remarkable efficiency of our training technique, it may not be too unrealistic, however, to hope for more rapid progress than has been possible in work with vertebrates.

REFERENCES

- BITTERMAN, M. E., & COATE, W. B. (1950). Some new experiments on the nature of discrimination learning in animals. *Journal of Comparative & Physiological Psychology*, **43**, 198-210.
- 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*, **96**, 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. (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., KLOSTERHALFEN, S., & BITTERMAN, M. E. (1983). Analysis of overshadowing in honeybees. *Journal of Comparative Psychology*, **97**, 154-166.
- GOULD, J. L. (1982). *Ethology: The mechanisms and evolution of behavior*. New York: Norton.
- GRAF, V., & TIGHE, T. (1971). Subproblem analysis of discrimination-shift learning in the turtle (*Chrysemys picta picta*). *Psychonomic Science*, **25**, 257-259.
- HULL, C. L. (1929). A functional interpretation of the conditioned reflex. *Psychological Review*, **36**, 498-511.
- HULL, C. L. (1943). *Principles of behavior*. New York: Appleton-Century-Crofts.
- KLOSTERHALFEN, S., FISCHER, W., & BITTERMAN, M. E. (1978). Modification of attention in honeybees. *Science*, **201**, 1241-1243.
- LINDAUER, M. (1971). *Communication among social bees*. Cambridge: Harvard 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: 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.
- SPENCE, K. W. (1936). The nature of discrimination learning in animals. *Psychological Review*, **43**, 427-429.
- SUTHERLAND, N. S., & MACKINTOSH, N. J. (1971). *Mechanisms of animal discrimination learning*. New York: Academic Press.
- TENNANT, W. A., & BITTERMAN, M. E. (1973). Some comparisons of intra- and extradimensional transfer in discriminative learning of goldfish. *Journal of Comparative & Physiological Psychology*, **83**, 134-139.
- 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.
- WOODARD, W. T., & BITTERMAN, M. E. (1976). Asymptotic reversal learning in pigeons: Mechanisms for reducing inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*, **2**, 57-66.
- WORTZ, E. C., & BITTERMAN, M. E. (1953). On the effect of an irrelevant relation. *American Journal of Psychology*, **66**, 491-493.