

*COMPUTATIONAL BEHAVIOR DYNAMICS:
AN ALTERNATIVE DESCRIPTION OF NEVIN (1969)*

CHARLES P. SHIMP

UNIVERSITY OF UTAH

A computational processing behavior-dynamic model was instantiated in the form of a computer program that "behaved" on the task developed by Nevin (1969). In this classic discrete-trials experiment, the relative frequency of choosing a response alternative matched the relative frequency of reinforcement for that alternative, the local structure of responding was opposite that predicted by momentary maximizing (i.e., the probability of a changeover decreased with run length), and absolute and relative response rates varied independently. The behavior-dynamic model developed here qualitatively reproduced these three results (but not in quantitative and specific detail) and also generated some interesting, as-yet-untested predictions about performance in Nevin's task. The model was discussed as an example of a stochastic behavior-dynamic alternative to algebraic behavior theory.

Key words: discrete-trials choice, computer model, matching, maximizing

A classic experiment by Nevin (1969) continues to play a role in theoretical analyses of molecular and molar accounts of behavior processes (Nevin, 1982; Silberberg, Hamilton, Zirix, & Casey, 1978; Williams, 1990). The underlying problem is how to describe the behavioral adaptation that occurs as a result of an organism's experience with the environment. Interestingly, there seems at present no single account of any type that handles several of the results Nevin obtained. An integrated account is clearly needed.

There is increasing interest in stochastic, behavior-dynamic approaches to describing behavior processes (Shimp, 1989; Staddon & Bueno, 1991). The purpose of the present paper is therefore to describe an example of the kind of alternative account a behavior-dynamic approach can provide for Nevin's experiment. The computational model described here be-

longs to a family of stochastic, real-time models that have been applied previously to a variety of other behavioral phenomena: Very closely related models have simulated timing performances (Shimp, 1984b), interresponse-time performances (Shimp, 1984b; Shimp, Childers, & Hightower, 1990), and concurrent ratio performances (Shimp, 1984a).

It will be useful to review Nevin's procedure and three chief results he obtained. In the first of his two experiments, he arranged a concurrent variable-interval (VI) 1-min VI 3-min schedule of food reinforcement for pecks on two keys in a discrete-trials procedure.¹ Both keys were lit for 2 s, or until a response occurred to either one. A response either was reinforced with 4-s access to food or initiated a 6-s intertrial interval (ITI), during which the chamber was dark. The two VI schedules ran independently. Sessions lasted 30 min, and sufficient training was provided so that daily performance was not obviously changing. In the second experiment, reinforcement became available beginning with the seventh postreinforcement trial: On the seventh trial, a re-

This research was supported in part by NIMH Grant R01 MH 42270, the Biomedical Sciences Support Grant Committee of the University of Utah, and the University Research Committee of the University of Utah. I thank Frances Friedrich for her comments on the model, William Harris of PsyComp Services for his help in the development of the computer programs, and Shahla Mohammad-Hosseini for her help in the analysis of the simulation results. I also thank Jack Marr, Bill Palya, and the administration of Jacksonville State University for hosting the Jacksonville Conference on Behavior Dynamics, at which preliminary results with the present model were presented. Numerous participants and guests at the conference, as well as the hosts, provided helpful comments on the research described here.

¹ Nevin's procedure involved left and right keys, the colors of which varied over trials in a somewhat unusual way. Schedules were associated with key colors (red and green) rather than positions (left and right). There was no important effect reported due to the key color procedure. The present discussion will refer simply to Nevin's two response alternatives. It seems likely that it is safe to think of these as left and right keys.

inforcer was arranged for one of the keys, with a probability of .75 for one response and .25 for the other. The reinforcer was held until collected.

Three results for which it would be helpful to have a single account are the following. First, in both experiments, the relative frequency of pecks on one alternative approximately equaled the relative frequency of reinforcements contingent on responding on that alternative. This result is, of course, the familiar matching result. Second, in Experiment 1, the longer a subject continued to respond over successive trials to the key with the overall higher rate of reinforcement, the less likely it was to switch to the "worse" schedule: Response perseveration on the "better" key increased as a function of the number of previous responses on that key. The fact that the likelihood of a changeover from the better to the worse schedule decreased as the number of trials since the last changeover increased is in sharp contrast to the outcome expected by momentary maximizing (Shimp, 1966). In terms of that account, the local structure of responding maximizes local reinforcement probability, and, in a concurrent interval schedule, the probability of reinforcement for changing over from the better to the worse schedule increases as time since the last changeover increases. Third, in Experiment 2, the relative frequency of responding to a key continued to approximate the matching value over postreinforcement trials for which the absolute probability of responding varied, ranging from an initially low value after reinforcement when no reinforcements were arranged, to a relatively high value after a number of trials closer to seven, when reinforcement became available. Thus, absolute and relative response rates varied independently.

In an ideal world, a model might handle every feature of behavior anyone might have an interest in, and might do so to any desired degree of quantitative precision. The present model is not intended to achieve such breadth and accuracy, even for Nevin's data. The present goal is to begin to provide a behavior-dynamic alternative account for Nevin's data. As noted above, an alternative seems needed because after 20 years, it is not clear that any model currently handles his chief results. For instance, neither momentary maximizing nor the matching law describes the local organization of choices he obtained. The former pre-

dicts the opposite of Nevin's results, and the latter's predictions have not been articulated sufficiently clearly to permit them to be tested. It therefore seems acceptable in a first attempt to develop a behavior-dynamic model for Nevin's results to emphasize scope at some expense in quantitative precision, in the manner of several contemporary computational processing models (Kehoe, 1989; Wearden & Clark, 1988).

THE MODEL

The model's separate behavior processes can be summarized fairly simply. The discriminative stimulus confronting a subject on a moment-to-moment basis is described as a collection of dynamic, rapidly shifting features. This stimulus is assumed simultaneously to reflect both current environmental stimuli and aftereffects of recent stimuli. A subject's reinforcement history with each possible configuration of this shifting stimulus is assumed to control behavior on a moment-to-moment basis. The goal is to develop a model that is sufficiently well described on a moment-to-moment basis so that it actually "behaves" (i.e., generates a behavior stream; Shimp, 1989).

A comment is in order regarding the language in terms of which the separate behavior processes are described. Natural language, including many of its applications in cognitive psychology, encourages appeals to hidden agents, homunculi, and assorted logical infinite regressions (Skinner, 1957; Wittgenstein, 1953). In the following exposition, natural language expressions involving remembering, knowing, and so on are occasionally used metaphorically to simplify and make more intuitive what otherwise would consist of nothing but sections of a computer program; it should not be forgotten, however, that translation into purely mechanical terms is possible. A principal design objective here is to develop a dynamic model that behaves and learns automatically in response to whatever are the model's environmental experiences, without control or guidance by any hidden agent.

Behavioral Processing of a Stimulus

The model "sees" a stimulus as a collection of component stimuli (Neimark & Estes, 1967; Skinner, 1932). The model's visual response is perhaps most conveniently likened to a dig-

itized image on a computer monitor, with any given pixel illuminated or not, depending on the presence or absence of a corresponding stimulus of sufficient intensity. "Features" of the model's visual response correspond to pixels, according to this analogy, and like pixels, features could correspond to spatial locations and represent other dimensions as well (e.g., color). To apply the model to Nevin's (1969) procedure, one needs to attribute stimulus properties to the keylights, to each of the two responses, and to the food hopper, because these are the events in terms of which Nevin's contingency is defined. The occurrence of any of these events is assumed to "activate" all the component features of the model's visual response to these events. There is no spontaneous activation of component features in the absence of the corresponding stimulus. In the simulations described here, there were five features each in a visual response to a keylight, to a response on a key, and to the food hopper. There were, therefore, 25 features in all, because there were two keys, two responses, and the food hopper. This number of features was selected to ensure computational tractability (simulation time increased geometrically with number of features). A small amount of exploration with other numbers suggested that there is nothing peculiar in terms of the simulation results about five features per stimulus.

Aftereffects of Recent Events

Each feature in the model's visual response is, at any given moment, either active or inactive. Short-term aftereffects of a stimulus derive from deactivation of features in the visual response. The deactivation of a feature is assumed to depend purely on time: In every short interval of time, Δt , there is a probability, P_f , that any activated feature will become deactivated. Throughout all the simulations described here, the value of Δt was 0.5 s. Two simplifying assumptions are made about this deactivation: The value of P_f for one feature is the same as that for any other feature, and P_f does not depend on time since a feature was activated or on the activation status of any other features. These assumptions are discussed in greater detail in Shimp et al. (1990).

Learning

The behavioral processes described so far imply that what the model "sees" dynamically changes as a function of environmental stimuli

and time. At any moment, the configuration of 25 active and inactive features can be represented as a 25-bit binary number. Every possible momentary configuration corresponding to a subject's visual response during the experiment is assumed to correspond to one of these 2^{25} numbers. In practice, in the context of any particular experimental task, only a relatively small fraction of these configurations actually appears.

These momentary configurations have associated reinforcement histories. The model associates configurations of features, not individual features, with reinforcement histories. Accordingly, which specific features are still active is critical, because it is the entire specific pattern or configuration that defines a particular 25-bit binary number.

Such a number corresponding to a particular pattern is added to a list defining the model's "memory" for reinforcement history provided the following conditions are met. If a response occurs (see below) when the model's momentary stimulus configuration has a particular numerical value, call it x , and a reinforcer is delivered, then that configuration x will be added to the model's memory if it is not already there. A configuration y is removed from memory with probability P_r , provided that a response occurs when the model's momentary stimulus configuration has a value equal to y , a response occurs, and it is not reinforced. Any pattern is associated at any moment with at most one response: The model cannot, for instance, make a left response if the current pattern is associated with the right response. The only way for such a pattern to become associated with a left response is for the pattern first to be probabilistically removed (according to parameter P_r) from "memory" by occurring and producing a right response, which is then unreinforced, and then subsequently reoccurring and producing a left response by chance (according to parameter P_0). If this randomly produced left response were reinforced, the pattern would be reinstated in memory and a left response would subsequently be made when that pattern reoccurred.

Responding

The model asks at each moment if its current visual response (recall this is a 25-bit binary number) has an associated reinforcement history: It merely asks if the current number is in the list of numbers defining its memory of

its reinforcement history. If it is, the model "pecks" the "key" appropriate to the reinforcement history (it emits the response it "remembers" having been reinforced). If it is not, the model responds with only a low base-rate probability, P_0 , and chooses randomly between the two keys. When a key peck occurs, the five features corresponding to a peck on that key are activated. Thus, the model subsequently "knows" what key it last pecked because more of that key's stimulus features, rather than of the other key's features, will be more likely to be activated. In other words, this is the model's way of handling the stimulus aftereffects of a response.

In summary, the model's current stimulus configuration serves at each moment as a stimulus, either an S^D or an S^A , corresponding to the model's environmental history. If the configuration is an S^D , the model emits the response it recalls having been reinforced in the presence of that S^D . If the configuration is an S^A , the model responds randomly between keys and with only a low base-rate probability.

Miscellaneous Details of Simulations

These ideas about real-time behavior processes were instantiated in a computer program to simulate pigeons' performances. Simulated subjects were given the equivalent of either 20 or 100 days of training on any given condition. The VI schedules used in the simulations were constant reinforcement-per-opportunity schedules with a minimum inter-reinforcement interval of 1 s. The two schedules ran independently, as in Nevin's (1969) experiment. Unless otherwise specified, the numerical parameters of the reinforcement contingencies in the simulations corresponded to those in Nevin's experiment. Performance during the last 5 days of training were analyzed. Each simulation started with no reinforcement history; performances were initially entirely due to baseline responding and came under the control of a reinforcement contingency only after a simulated subject encountered it and acquired a reinforcement history.

Performances of several subjects were simulated for each condition, with each subject in a group having identical numerical values of the three theoretical parameters, P_f , P_r , and P_0 . The number of subjects in a condition depended primarily on the experimenter's availability, which varied greatly over days, and on

which computers (IBM® model 80s or an IBM® model 95) were used. Enough subjects were run in any given condition, in any case, so that additional subjects were judged to be unlikely to change the qualitative picture. Each simulation used a different random seed for the random number generator.

Parameter estimation was informal: Due to the long required computing time, best fitting values of theoretical parameters are not available. To obtain results for 6 simulated subjects for a single experimental condition for a single set of numerical values for the theoretical parameters, 3 to 20 hours (depending on the contingency, amount of training, and parameter values) were required. A total of well over 1,000 hours was required for the various simulations, of which those presented here are a small sample. Several dozen different parameter combinations were tried, with results from earlier simulations informally guiding experimentation with subsequent parameters.

RESULTS AND DISCUSSION

The model's interacting stochastic behavior processes and the probabilistic nature of the task automatically produced behavior variability. Variability across subjects within a condition arose without different theoretical parameters or different initial reinforcement histories. In a sense, then, all the results displayed below for a given schedule condition and for a fixed parameter set are for a single subject. The results simply show how the same subject behaved differently on different exposures to the same contingency.

No effort is made here to capture the specific details of each of Nevin's (1969) 3 subjects' performances. The handling of such details will probably have to await the future development by the computer industry of higher speed computational resources.

Nevin (1969) focused exclusively on a particular procedure, a concurrent VI 60-s VI 180-s schedule arranged with 6-s intertrial intervals, so that functional relations involving relative frequency of reinforcement, absolute frequency of reinforcement, or trial spacing were not generated. It was of interest here, however, to explore the model's performances in terms of functional relations. Therefore, absolute and relative frequency of reinforcement and the intertrial interval were varied in com-

binations not yet empirically investigated, but for which interesting new predictions are shown below.

Figure 1 shows the relative frequency of occurrence of a response alternative averaged over the last 5 of 20 (upper right and middle left) or of 100 sessions as a function of the obtained relative frequency of reinforcement for that alternative. The reinforcement schedule was the following for each panel: top left, concurrent VI 2 s VI 6 s, concurrent VI 6 s VI 6 s, and concurrent VI 6 s VI 2 s; top right, concurrent VI 10 s VI 30 s and concurrent VI 30 s VI 10 s; middle left, concurrent VI 60 s VI 180 s, concurrent VI 60 s VI 60 s, and concurrent VI 180 s VI 60 s; middle right, concurrent I 360 s VI 1,080 s and concurrent VI 1,080 s VI 360 s; bottom left and right, concurrent VI 60 s VI 180 s and concurrent VI 180 s VI 60 s. The corresponding best fitting straight lines and values of r^2 were the following: top left, $y = 0.975x + 0.010$ and $r^2 = 0.990$; top right, $y = 0.810x + 0.085$ and $r^2 = 0.870$; middle left, $y = 0.673x + 0.126$ and $r^2 = 0.805$; middle right, $y = 0.759x + 0.133$ and $r^2 = 0.816$; bottom left, $y = 0.909x + 0.060$ and $r^2 = 0.818$; bottom right, $y = 0.808x + 0.092$ and $r^2 = 0.934$. Each point represents one simulation of a subject's performance. The intertrial interval was 6 s, trial duration was 2 s, and food hopper duration was 4 s (the same values as in Nevin, 1969). The parameter values used in the top four panels were $P_r = .025$, $P_f = .150$, and $P_0 = .025$. In the bottom left and right panels they were .025, .0005, and .015, and .025, .001, and .025, respectively. The diagonal represents the matching line.

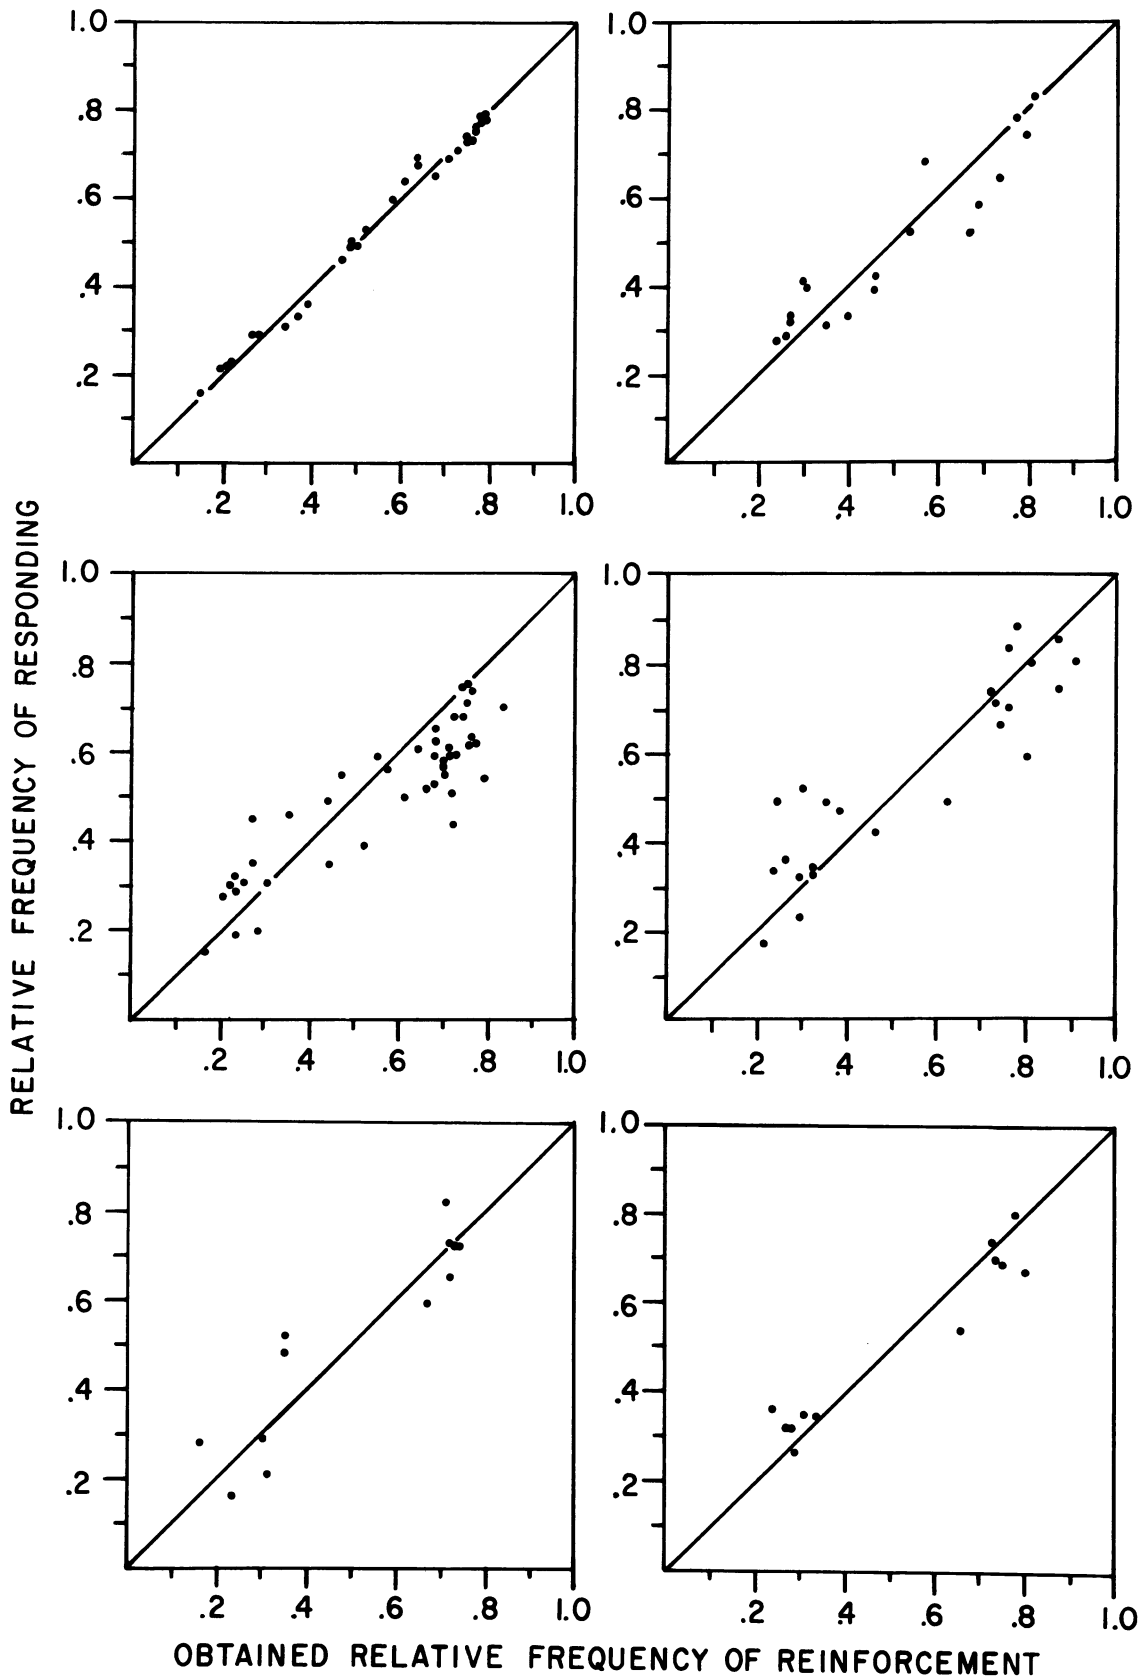
The upper left panel shows a case extreme in several ways. First, the absolute frequency of reinforcement was far greater than in Nevin's (1969) experiment. Second, the degree of obtained matching was extraordinary. This matching needs to be interpreted with some delicacy, however, because the high reinforcement rate guaranteed that virtually every response was reinforced, so matching to obtained reinforcement was forced by the procedure. What was not forced was the particular type of matching: All nine points in the bottom left quadrant were produced by the concurrent VI 6-s VI 2-s schedule, all in the middle were produced by the concurrent VI 6-s VI 6-s schedule, and all 15 in the upper right quad-

rant were produced by the concurrent VI 2-s VI 6-s schedule. Thus, relative frequency of responding, on average, closely matched programmed as well as obtained relative reinforcement rate.

Interestingly, the very precision of conformity to the algebraic matching function in this special use reminds us of the unrealistically low (i.e., zero) variability predicted by that function. The algebraic matching function typically underestimates real-world variability.

The remaining top three panels of Figure 1 show that the model produces more nearly realistic levels of variability when absolute frequency of reinforcement is reduced. Undermatching is also evident to varying degrees in these panels. (Recall that when a momentary stimulus configuration does not have a "remembered" reinforcement history, the model chooses randomly, thus driving preference toward indifference.) The obtained undermatching may be in the range obtained with related free-operant concurrent procedures, as reviewed by Wearden and Burgess (1982) and Davison and McCarthy (1988, p. 85). These reviews suggest that slopes of the best fitting straight line in the range shown here should occur frequently. The bottom two panels of Figure 1 show results obtained with two other sets of parameters, for absolute and relative reinforcement frequencies that were the same as Nevin (1969) used. Again the degree of undermatching does not seem discriminably different from that reported in the literature review by Wearden and Burgess (1982). In general, it is not clear that a researcher would be able to tell the difference between these simulation results and corresponding results in the empirical literature by the Turing test criterion (Turing, 1950).

Figure 2 illustrates how a behavior-dynamic model can generate theoretical predictions for as-yet-unexplored empirical conditions. The left and right panels in Figure 2 show preference as a function of the intertrial interval when absolute reinforcement rates are unusually high or unusually low, respectively. Each point represents the average over the last 5 of 100 sessions of one simulation of 1 subject's performance in the procedure of Experiment 1 in Nevin (1969). Trial duration was 2 s and food hopper duration was 4 s, as in Nevin's experiment. The parameter values used were $P_r = .025$, $P_f = .150$, and $P_0 = .025$. (These






Fig. 2. The relative frequency of one response alternative as a function of the intertrial interval, for a concurrent VI 2-s VI 6-s schedule (left panel) or a concurrent VI 360-s VI 1,080-s schedule (right panel). The horizontal dashed line represents programmed relative frequency of reinforcement. See text for details of the model that produced these results.

theoretical values are the same as in the top four panels of Figure 1.) In both panels of Figure 2, undermatching was obtained at an intertrial interval of 1 s, and approximate matching was obtained with Nevin's intertrial interval of 6 s. With a long intertrial interval of 24 s, a very high reinforcement rate produced extreme overmatching, whereas a low reinforcement rate produced approximate matching but with very great variability.

For the past 20 years, an implicit assumption seems to have been that Nevin's (1969) matching result for a particular combination of intertrial interval and absolute and relative reinforcement frequency was suggestive of a general matching outcome. The present results show how a behavior-dynamic model can help identify previously unexamined assumptions and hence possibly important experimental manipulations.

Figure 3 shows the relative frequency of switching as a function of successive choices of the better alternative. The relative frequency of switching was calculated on the basis of the number of opportunities for switching at each run length, as in Nevin (1969). The critical feature of Nevin's results was that the function went down, not up, as predicted by momentary maximizing.

Left and right panels in Figure 3 correspond to individual and average relative frequencies of changeovers as produced by the model. The intertrial interval was 6 s, trial duration was 2 s, and food hopper duration was 4 s, all as in Nevin (1969). The values of P_r , P_b , and P_0 are equal to .025, .150, and .025 in the top panels; .025, .001, and .025 in the middle panels; and .025, .0005, and .015 in the bottom panels. The functions go down, not up, and in this qualitative sense resemble Nevin's results. The parameter values in Figure 3 are the same as those in the three panels of Figure 1 corresponding to concurrent VI 1-min VI 3-min schedules. The same parameter values therefore cause the model's performance to be qualitatively correct for both preference and switching behavior.

Figure 4 shows results of letting the model learn to respond on the procedure of Experiment 2 in Nevin (1969). The main question for a first approximation to answer seems to be whether the model can demonstrate independence of absolute and relative responding. An additional and more specific question would be whether the model can display matching while absolute responding increases from a low to a high level as a function of postreinforcement trial number.

Fig. 1. The relative frequency of one response alternative as a function of the obtained relative frequency of reinforcement for that alternative. The solid line is the matching function. See text for details of the model that produced these results.

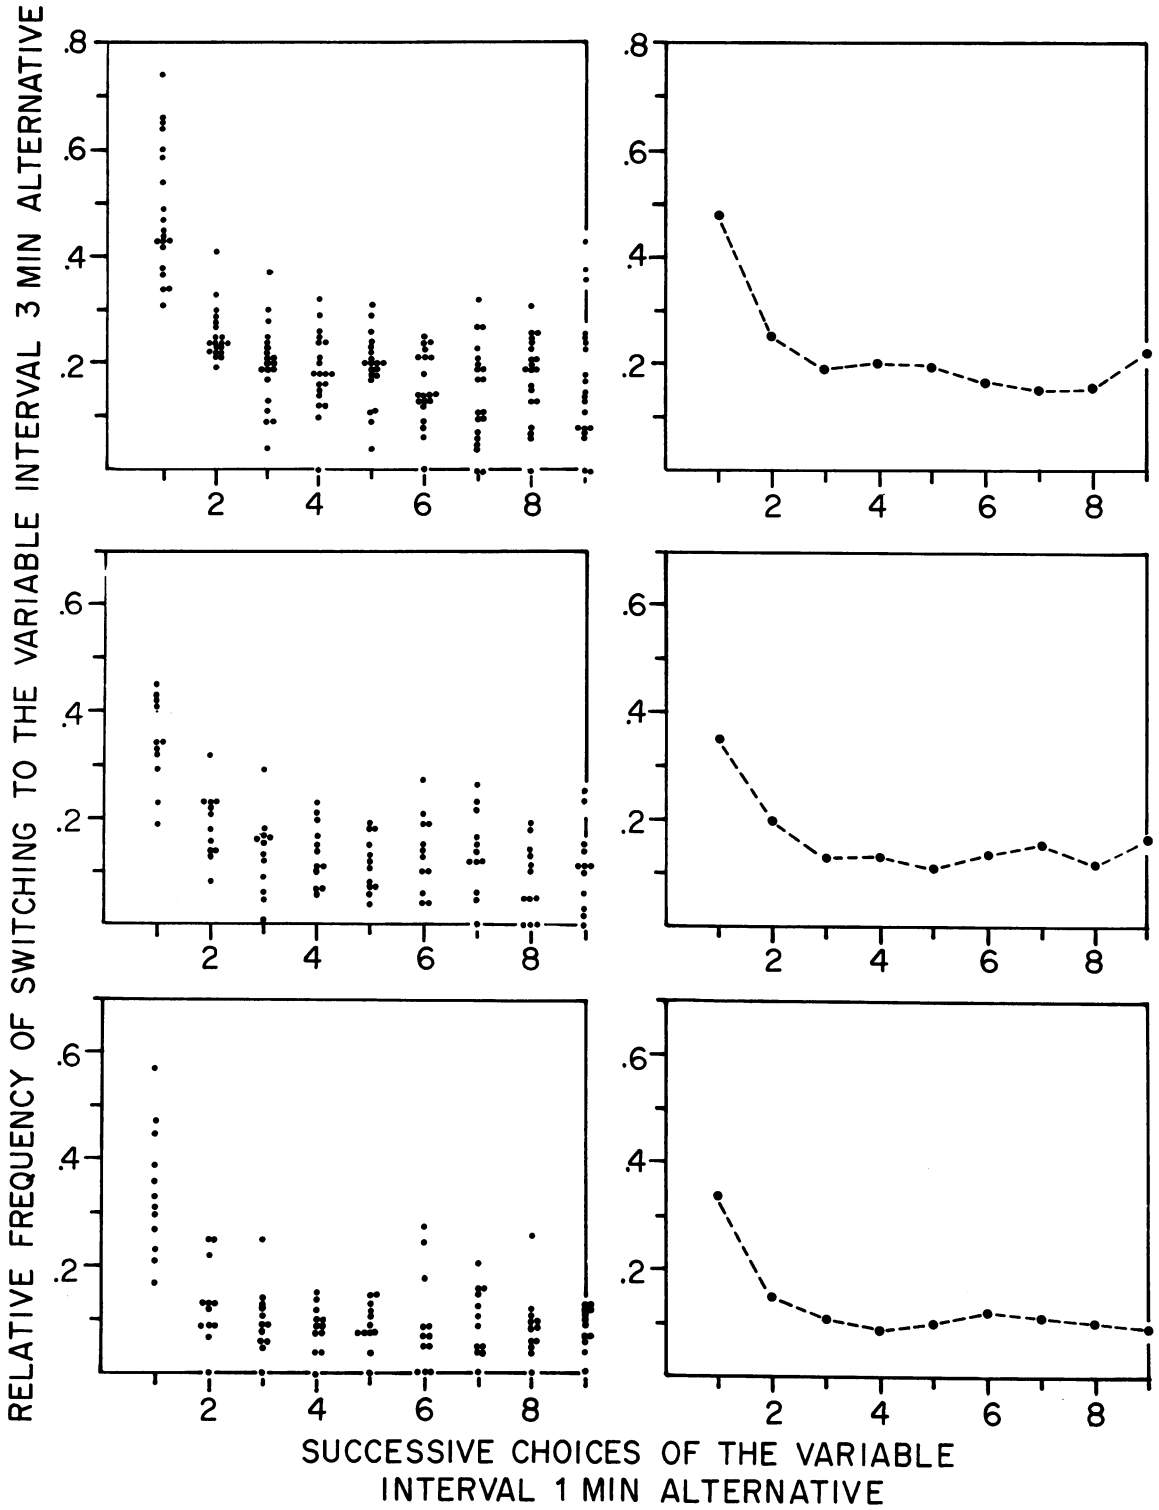


Fig. 3. The relative frequency of switching to the alternative associated with a VI 180-s schedule as a function of the number of preceding successive choices of the alternative associated with a VI 60-s schedule, averaged over the last 5 of 20 (top panel) or 100 (middle and bottom panels) sessions. See text for details of the model that produced these results.

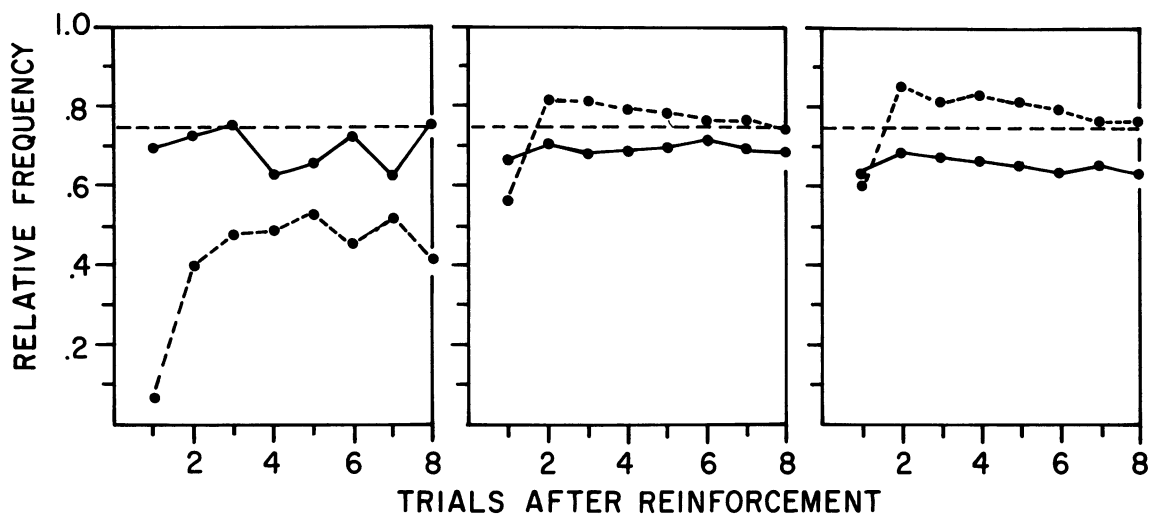


Fig. 4. The relative frequency of responding given a trial (dashed line), and the relative frequency of choosing the alternative associated with a VI 60-s schedule (solid line), as a function of the ordinal trial number after reinforcement. Reinforcement was scheduled only on the seventh trial and was held until collected, as in Experiment 2 of Nevin (1969). The horizontal dashed line at .75 indicates the scheduled relative frequency of reinforcement. See text for details of the model that produced these results.

Figure 4 shows averages across 6 subjects to facilitate discriminating the two separate functions in each panel. The functions are averages over the last 5 of 100 (left and right) or 20 (middle) sessions. The three panels in Figure 4 correspond to three different parameter sets. Parameter values for P_r , P_i , and P_0 are .025, .075, and .0025 in the left panel; .010, .150, and .010 in the middle panel; and .025, .150, and .025 in the right panel. Those for the right panel are the same as those for the top four panels in Figure 1. The absolute response rate was lowest in the left panel and highest by a slight margin in the right panel. There was no systematic corresponding variation in relative response rate, which in all three panels is below the matching value. The degree of undermatching was about the same as in several of the panels in Figure 1. In Nevin's Experiment 2 (1969), in contrast, all 3 birds' preferences were, on the average, quite close to the matching value.

The critical result is that the degree of undermatching was not related to the absolute rate of responding. That is, absolute and relative responding did not covary. However, no parameters were found that produced any other than a very crude qualitative description of how absolute responding varied as a function of postreinforcement trial number. When, as in the left panel, response rate began very low,

it never reached a sufficiently high value to conform to Nevin's results, where absolute response likelihood was above .9 after five trials. When, as in the middle and right panels, the absolute response rate was quite high after several trials, it was too high on the first postreinforcement trial, peaked too early on the second or third postreinforcement trial, and then slightly decreased. In short, several specific details of the predicted shapes of the absolute response-rate function are wrong.

Thus, the model can learn to behave on the task of Nevin's Experiment 2 (1969) so that there is the correct qualitative dissociation between absolute and relative responding, but the model does not capture several specific quantitative features of his data.

GENERAL COMMENTS

Neither Nevin's (1969, 1979, 1982) account (based on the matching law) nor Shimp's (1966) account (based on momentary maximizing) of discrete-trials choice behavior has provided integrated accounts of various features of the results reported by Nevin (1969). The present behavior-dynamic model provides a stochastic alternative to these previous deterministic accounts. The present model learns from scratch, demonstrates a degree of undermatching that is reasonably consistent with the

most nearly relevant available literature, shows an appropriately increasing response perseveration as a function of the number of successive choices of the preferred response alternative, and shows a degree of independence of absolute and relative response rate in Experiment 2. Thus, although the model most definitely does not perfectly describe Nevin's results, it appears to be a viable alternative to other existing accounts. In addition, it makes some new predictions about experimental conditions not yet investigated.

As with any current theory of behavior, the present model leaves many unanswered questions. For instance, no set of parameters was found that simultaneously produced matching and the appropriate magnitude of change in absolute response frequency in Nevin's (1969) Experiment 2. Also, some of the model's more successful descriptions of the results of Nevin's Experiment 1 (e.g., the bottom panels of Figure 1) used parameters different from those in Figure 4 for his Experiment 2. The parameters for the bottom panels of Figure 1 actually produced *increasing* absolute response probabilities in Experiment 2 (not shown here). The procedures of Experiments 1 and 2 were so different that different parameters are perhaps not too troubling. Nevertheless, an integrated account would surely be simpler if the same or similar parameters applied in both cases. It should be noted that the minimum requirement involving parameter invariance was obtained: The same parameters produced both undermatching and the correct changeover function in Experiment 1. This within-experiment parameter invariance is essential and was obtained.

Other unanswered questions pertain to the model's novel predictions. Only substantial empirical work will answer the question of whether the predictions in Figure 2 concerning effects of intertrial interval and absolute reinforcement rate are at all descriptively accurate. At present, these predictions serve primarily to remind us that sometimes having theoretical predictions in advance of data can help identify potentially interesting experiments.

Still other questions focus on why the model behaves as it does. The statistical complexities of a stochastic interactive model naturally far exceed those of a static algebraic model. One knows at the outset that an algebraic model is

only a stylized, oversimplified description of behavior. But this oversimplification is accepted because of the elegant simplicity and tractability of an algebraic model, and the ease with which it can be understood. The study of the rich complexity of an actual behavior stream is sacrificed in favor of an elegantly simple account. With interactive stochastic real-time models, the situation is in a sense reversed. One may hope in principle for a more realistic description of complex behavior dynamics but at the expense of a readily understandable model; the behavior a model emits is an emergent property of a collection of interacting, stochastic, component behavior processes. In the present case, each process can be separately described quite simply. But the result of these processes interacting on a moment-to-moment basis is much less readily apparent. Indeed, to find out what the result is, one has to resort to computer simulation. A challenge for future work will be to provide, perhaps in part through the field in computer science of "scientific visualization," an intuitively satisfying understanding of complex stochastic models.

Presumably it will be some time before an experimenter will be unable to tell the difference between a behaving stochastic model and a behaving pigeon in Nevin's (1969) procedure according to the conventional Turing test for computational models (Turing, 1950). In the meantime, it should prove quite challenging to contribute to progress toward this specific goal as well as toward the more general goal of developing behavior theory through what might, by analogy to "computational vision" (Marr, 1982), be called "computational behavior."

REFERENCES

- Davison, M., & McCarthy, D. (1988). *The matching law*. Hillsdale, NJ: Erlbaum.
- Kehoe, E. J. (1989). Connectionist models of conditioning: A tutorial. *Journal of the Experimental Analysis of Behavior*, *52*, 427-440.
- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. San Francisco: W. H. Freeman.
- Neimark, E. D., & Estes, W. K. (Eds.). (1967). *Stimulus sampling theory*. San Francisco: Holden-Day.
- Nevin, J. A. (1969). Interval reinforcement of choice behavior in discrete trials. *Journal of the Experimental Analysis of Behavior*, *12*, 875-885.

- Nevin, J. A. (1979). Overall matching versus momentary maximizing: Nevin (1969) revisited. *Journal of Experimental Psychology: Animal Behavior Processes*, *5*, 300-306.
- Nevin, J. A. (1982). Some persistent issues in the study of matching and maximizing. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol 2. Matching and maximizing accounts* (pp. 153-165). Cambridge, MA: Ballinger.
- Shimp, C. P. (1966). Probabilistically reinforced choice behavior in pigeons. *Journal of the Experimental Analysis of Behavior*, *9*, 443-455.
- Shimp, C. P. (1984a). Relations between memory and operant behaviour, according to an associative learner (AL). *Canadian Journal of Psychology*, *38*, 269-284.
- Shimp, C. P. (1984b). Timing, learning and forgetting. In J. Gibbon & L. Allan (Eds.), *Timing and time perception* (Vol. 423, pp. 346-360). New York: New York Academy of Sciences.
- Shimp, C. P. (1989). Contemporary behaviorism versus the old behavioral straw man in Gardner's *The Mind's New Science: A History of the Cognitive Revolution*. *Journal of the Experimental Analysis of Behavior*, *51*, 163-171.
- Shimp, C. P., Childers, L. J., & Hightower, F. A. (1990). Local patterns in human operant behavior and a behaving model to interrelate animal and human performances. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*, 200-212.
- Silberberg, A., Hamilton, B., Zirix, J. M., & Casey, J. (1978). The structure of choice. *Journal of Experimental Psychology: Animal Behavior Processes*, *4*, 368-398.
- Skinner, B. F. (1932). On the rate of formation of a conditioned reflex. *Journal of General Psychology*, *7*, 274-286.
- Skinner, B. F. (1957). *Verbal behavior*. New York: Appleton-Century-Crofts.
- Staddon, J. E. R., & Bueno, J. L. O. (1991). On models, behaviorism and the neural basis of learning. *Psychological Science*, *2*, 3-11.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, *59*, 433-460.
- Wearden, J. H., & Burgess, I. S. (1982). Matching since Baum (1979). *Journal of the Experimental Analysis of Behavior*, *38*, 339-348.
- Wearden, J. H., & Clark, R. B. (1988). Interresponse-time reinforcement and behavior under aperiodic reinforcement schedules: A case study using computer modeling. *Journal of Experimental Psychology: Animal Behavior Processes*, *14*, 200-211.
- Williams, B. A. (1990). Enduring problems for molecular accounts of operant behavior. *Journal of Experimental Psychology: Animal Behavior Processes*, *16*, 213-217.
- Wittgenstein, L. (1953). *Philosophical investigations* (G. E. M. Anscombe, Ed.). New York: Macmillan.

Received June 11, 1991

Final acceptance December 3, 1991