

Cumulative Effects Model: A Response to Williams (1994)

J. E. R. Staddon, D. G. S. Davis, A. Machado, and R. G. Palmer

The cumulative effects (CE) model explains free-operant choice by the ratio of total numbers of responses and reinforcements, a probability-like variable. Williams (1994) argues that the model is vulnerable to experiments that disprove melioration, a local probability model. The authors note critical differences between the nonlocal CE model and local probability models that allow the CE model to handle some data with which they are incompatible. All models are simplifications of reality; hence, a model's failures are as revealing as its successes. Williams suggests that simple models may need to be abandoned in favor of a "representational" account. The authors point out that representations must be both acquired and acted on. Acquisition requires processing of responses and reinforcers; action requires decision rules. Models are simply testable suggestions for what these rules and processes might be.

The cumulative effects (CE) model is simple in concept, but its implications are easy to misunderstand. Because of space limitations, we touch on just three experimental results and conclude with some general points.

Williams and Royalty (1989)

The Williams and Royalty (1989) article is a series of three experiments carefully designed to test the melioration hypothesis. Williams (1994) argues that the results also pose severe difficulties for the CE model. The three experiments used multiple schedules (i.e., alternating, signaled components), discrete trials, and nonchoice as well as choice conditions. The CE model was proposed for single-context, free-operant choice situations, so its applicability to these experiments cannot be taken for granted. Nevertheless, Davis, Staddon, Machado, and Palmer (1993) discussed Experiment 1, and it is obviously important to see just how far the CE model can be generalized.

Williams (1994) suggests that Experiment 3 poses the most severe challenge. The rationale for the experiment is straightforward. If (obtained) probability is the only variable that determines choice, Williams argues, then subjects should prefer a lean variable interval (VI) schedule that is sampled infrequently to a richer schedule sampled frequently—a counterintuitive and maladaptive result. The experimental procedure comprised two alternating multiple-schedule discrete-trial components. The first component involved choice between two alternatives associated with interdependent VI schedules that paid off with probabilities of .2 (Alternative 1) and .05 (Alternative 2); the second component was a single VI that paid off with probability .1. Because the subjects approximately matched response and reinforcement ratios, the obtained payoff probabilities for Alternatives 1 and 2 (in the choice component) were equal (.24); the payoff probability for the single-choice component (Alternative 3) was of course equal to the scheduled value (.1). In probe tests at the end of training, the subjects were asked to choose between Alternative 1 and Alternative 3 or between Alternative 2 and Alternative 3. The "results were that the .20 alternative . . . was strongly preferred over the .10 alternative, whereas the .10 alternative was strongly preferred over the .05 alternative . . . , [which] suggests that the scheduled probability of reinforcement, not the obtained probability, was the controlling variable" (p. 705).

Does this result contradict the CE model? It is hard to be sure. The problem has to do with the term "probability." The probabilities in the Williams and Royalty (1989) experiment were measured in the standard "steady-state" way, over the last 5 days of each condition, whereas the CE model deals with probabilities measured over the subject's entire history. Probability, like "rate" but unlike "force" or "location," is not a self-contained property of a response. One must always ask, Over what period are these things to be assessed? Williams and Royalty did not provide the historical data that would be necessary to estimate probabilities in the CE-model sense.

In fact, the CE model can predict Williams's (1994) result under some conditions. The state of the model depends on two variables for each response, the sum of responses made and reinforcers received (including initial conditions). Once we estimate the initial conditions at the beginning of an extinction test, we can compute the trajectory when different pairs of alternatives are pitted against one another in extinction. We have done this for the critical comparison between the .05 alternative and the .1 alternative in Figure 1. We chose the initial conditions in the following way. The .1 alternative was single-choice during 50% of the trials. We assumed 1,000 responses and, therefore, 100 reinforcers. The .05 alternative was paired with the .2 alternative on 50% of the trials. Given matching, we assumed 200 responses and (given the obtained .24 payoff probability) 48 reinforcers. The figure shows 1,000 responses in extinction. At

J. E. R. Staddon, D. G. S. Davis, A. Machado, and R. G. Palmer,
Department of Psychology: Experimental, Duke University.

This research was supported by grants to Duke University from the National Science Foundation and the National Institute of Mental Health.

Correspondence concerning this article should be addressed to J. E. R. Staddon, Department of Psychology: Experimental, Duke University, Durham, North Carolina 27708. Electronic mail may be sent to staddon@psych.duke.edu.




Figure 1. Performance in an extinction test following Williams and Royalty (1989), Experiment 3, according to the cumulative effects model. The graph shows a response trajectory (thick line labeled *preference*): cumulative response for the .05 alternative plotted against cumulative responses for the .1 alternative. The thin diagonal line labeled *indifference* shows the trajectory corresponding to indifference between the two alternatives. The two converging thin lines are the *V* values for the two choices (right-hand Y axis). The initial conditions are assumed to be (200, 48) for the .05 response and (1,000, 100) for the .1 response. See text for details.

first the *V* value for the .05 alternative (thin downward-sloping line) is larger than for the .1 alternative (thin horizontal line: .24 vs. .1). Thus, the .05 alternative is chosen exclusively at first, and its *V* value declines hyperbolically. However, at some point, the declining .05 *V* value equals the constant .1 *V* value, and from that point (about 480 on the X axis), both responses are made (rising, thick line labeled *preference*) with a strong preference for the .1 alternative (the slope of the *preference* line is steeper than the slope of the *indifference* line). Thus, if the averaging period in extinction is long enough, the data will show a preference not for the .05 alternative but for the .1 alternative, which is what Williams and Royalty (1989) reported. They argued that such a preference disproves all reinforcement-probability models. We think they are wrong. What is true is that the CE model is quantitatively flawed: To get the predictions in Figure 1, we had to assume initial conditions that are too small. Another way to put this is that the CE model underweights recent experience (i.e., it lacks a recency assumption). As we see in a moment, it also underweights early experience (primacy). But this is not the same thing as Williams's sweeping conclusion that all choice models that incorporate reinforcement probability are ipso facto false.

Overlooked in Williams's (1994) comment is a pattern of data that provides substantial support for the CE model. For example, in Experiment 1, Williams and Royalty (1989) commented, "There was some evidence that the degree of preference changed systematically over the course of probe testing" (p. 102). Later they summarized, "The reason that probe testing was minimized in the present study is that pilot studies (and the results of initial phase of Experiment 1) indicated a strong tendency for the higher valued schedule to be increasingly preferred with continued testing" (p. 112). Because no reinforcements were given during probes, this shift cannot be predicted by most models. As Williams and Royalty (p. 111) pointed out, melioration (and, we would add, all local models) predicts a

shift toward indifference. Among operant choice models, only the CE model predicts the shift toward the richer option in extinction testing.

Williams (1993)

We agree with Williams (1993) that the simple CE model does not predict this result. The question remains as to why is there a difference between the successive and simultaneous groups? The CE model can shed some light on this question. Consider first the difference between the concurrent and successive conditions in Williams's (1993) experiment. Let us suppose that the animals enter the experiment initially indifferent (in terms of the CE model, this just means that the initial conditions are the same for all responses). Thus, they will initially respond too much to the .05 alternative and not enough to the .2 alternative. The "successive" animals are saved from this error, because they are forced to match right from the beginning. However, the CE model draws attention to the fact that it may be this property of the animals' history—the initial disproportion between responding and reinforcement in the concurrent group—that allows the "concurrent" subjects to detect the difference in properties between the .2 and .05 conditions. The initial excess of unreinforced .05 responses in the concurrent condition may be the source of the later preference for the successive-condition .05 stimulus. It may be that a "primacy" assumption (matching the "recency" assumption Williams already concedes as a possible modification for the model) can allow the model to accommodate data like this, although we have not had time to study this possibility in detail. Nevertheless, it does offer the attractive prospect of incorporating known properties of memory into our model of learning.

Belke (1992)

The CE-model prediction is unequivocally wrong, as is the prediction of every other context-free preference theory, including the idea that subjects are guided by the scheduled (as opposed to obtained) payoff rates or probabilities.

General Comments

What are we to make of these confusing experiments? Apparently neither local reinforcement probability (melioration) nor scheduled probability can describe choice performance under all conditions. The two-variable *V* values of the CE model seem to do better than simple probabilities. The model predicts changes in preference in extinction tests that seem to match the data. However, the quantitative predictions fall short in ways that suggest modifications to allow primacy and recency effects (i.e., greater weight to events early and late in a given context). No obvious modification promises to accommodate the Belke (1992) data, however. Indeed, there is some contradiction between Belke's data and the results of Williams and Royalty's (1989) Experiment 1—and as Williams (1993) pointed out, Belke's explanation for his transfer-test results begs the question of how the original matching behavior was established.

Three methodological conclusions do seem warranted: First, transfer tests, especially brief tests in which changes in prefer-

ence throughout the test are not recorded, are a poor way to assess the mechanisms of choice. Second, the relative success of the admittedly incomplete CE model strongly implies that local models of choice such as melioration are not adequate. The portion of an animal's past history that goes to determine choice behavior is extensive, perhaps not as extensive as assumed by the CE model, but certainly more extensive than the brief period assumed by melioration and other local choice models. A corollary is that the almost exclusive focus on steady-state data gathered over relatively brief "equilibrium" periods is mistaken. Much more attention needs to be paid to the whole pattern of behavior during an animal's exposure to a given set of choice procedures if we are to understand the mechanisms that determine both steady-state and transfer performance. Third, details make a difference. "Choice" is a grand general rubric that covers experiments that share a single feature: the availability of more than one response alternative. But choice in discrete-trial experiments (like Williams's, 1993, 1994) and in free-operant experiments (like Davis et al.'s, 1993, or Belke's, 1992) may be determined by very different processes, especially at different stages of training. "Response" is another grand general term, but it makes a difference whether pigeons must peck a certain key, in a fixed location, or a certain color, in a variable location. Direct comparisons between experiments in which responses are defined differently may be problematic.

Williams's (1994) comment raises two questions about models: Should all probability models be ruled out? And, should we all go cognitive? First, we believe that the issue is not whether "models . . . that rely principally on the pattern of obtained probability of reinforcement . . . [can] account for critical features of the . . . experimental literature" (p. 707). Clearly, reinforcement probability or, more precisely, its ingredients, namely number of reinforcers and number of responses, does affect behavior. The issue, therefore, is, Do these variables affect behavior in the ways that models which use them say they do? The fact that things other than probability also have effects is not relevant to this question—one might as well dismiss pressure as a determinant of volume having found that temperature is also important. For example, most choice models take the definition of the response and the assignment-of-credit problem for granted, yet both of these are things that themselves require an explanation. Both are very difficult questions in themselves, which is precisely why they are usually set aside until simpler issues have been settled. Making models is a way to simplify the

scientific task by looking at the effects of one or two variables at a time. If we take the simplification for the reality, we are missing the point of the process.

Should we all go cognitive and "abandon the incremental response strength assumptions underlying [these] models, in favor of representational accounts that assume the subjects of conditioning experiments acquire knowledge about the sources of reward in their environment" (Williams, 1994, p. 707)? We do not think so for three reasons: (a) Representational accounts are coming under increasing attack these days from philosophers and workers in behavior-based artificial intelligence (AI; e.g., Searle, 1992; Staddon, 1993), and we give these arguments some credence. (b) Knowledge, yes, but how gained? The pigeon can get information only from the responses it makes, the reinforcers it gets, and the times and places at which these events occur. No doubt it does "acquire knowledge," in some sense. The scientific question is, How does it do so? Exploring models that incorporate measureable variables is an essential part of the effort to understand how "knowledge" is acquired and what it is. (c) "Representation" is all very well, but the animal's problem is always what to do about it. Without action, knowledge is worthless. Without some kind of "strength" assumption, it is hard to see how knowledge can ever lead to action.

References

- Belke, T. W. (1992). Stimulus preference and the transitivity of preference. *Animal Learning and Behavior*, 20, 401–406.
- Davis, D. G. S., Staddon, J. E. R., Machado, A., & Palmer, R. G. (1993). The process of recurrent choice. *Psychological Review*, 100, 320–341.
- Searle, J. R. (1992). *The rediscovery of the mind*. Cambridge, MA: MIT Press.
- Staddon, J. (1993). *Behaviorism: Mind, Mechanism and Society*. London: Duckworth.
- Williams, B. A. (1993). Molar versus local reinforcement probability as determinants of stimulus value. *Journal of the Experimental Analysis of Behavior*, 59, 163–172.
- Williams, B. A. (1994). The role of probability of reinforcement in models of choice. *Psychological Review*, 101, 704–707.
- Williams, B. A., & Royalty, P. (1989). A test of the melioration theory of matching. *Journal of Experimental Psychology: Animal Behavior Processes*, 15, 99–113.

Received March 31, 1994
 Revision received May 2, 1994
 Accepted May 15, 1994 ■