

# Human Choice Among Five Alternatives when Reinforcers Decay <sup>★</sup>

Jacob Rothstein, Greg Jensen, Allen Neuringer <sup>\*</sup>

*Department of Psychology, Reed College, Portland OR 97202, USA*

---

## Abstract

Human participants played a computer game in which choices among five alternatives were concurrently reinforced according to dependent random-ratio schedules. “Dependent” indicates that choices to any of the wedges activated the random number generators governing reinforcers on all five alternatives. Two conditions were compared. In the *hold* condition, once scheduled, a reinforcer — worth a constant five points — remained available until it was collected. In the *decay* condition, point values decreased with intervening responses, i.e., rapid acquisition was differentially reinforced. Slopes of matching functions were higher in the *decay* than *hold* condition. However inter-subject variability was high in both conditions.

*Key words:* Concurrent Schedules, Human Subjects, Matching, Response Dependent Schedules, Choice, Reinforcement Magnitude

---

## 1. Introduction

The goal of this research was to enable rapid assessment of choices by human participants under concurrent schedules of reinforcement. Choices by humans often show a wide range of sensitivity to differences in reinforcement frequencies, and many explanations have been offered (Bradshaw and Szabadi, 1988; Kollins et al., 1997). We devised a procedure that differed in a number of ways from those commonly employed and tested the influence of one aspect of the procedure.

Participants played a computer game in which they chose among five alternatives, each alternative associated with a different random-number generator that governed reinforcement. Every response fired all five random-number generators. This schedule differed from the more common concurrent vari-

able interval (concurrent VI-VI) in that reinforcer availability (or “set up”) depended on responses, rather than passage of time, but was similar to the concurrent VI-VI procedure in that a reinforcer could set up on a given alternative whether or not the participant was responding on that alternative (for similar schedules, see Lau and Glimcher, 2005; Macdonall, 1988). In one condition, referred to as *hold*, a reinforcer remained indefinitely after set up until collected. This condition is similar to concurrent VI-VI schedules. Under another condition, referred to as *decay*, reinforcer values (the number of points potentially received) decreased with successive choices to other alternatives. The *decay* condition differentially favored immediate (or quick) access to available reinforcers. We compared *hold* and *decay* conditions in terms of stability of choices and their sensitivity to reinforcer ratios.

---

<sup>★</sup> This study supported by NIH MH068259.

<sup>\*</sup> To whom correspondence should be addressed

*Email addresses:* [jbr@alumni.reed.edu](mailto:jbr@alumni.reed.edu) (Jacob Rothstein), [belarius@gmail.com](mailto:belarius@gmail.com) (Greg Jensen), [allen.neuringer@reed.edu](mailto:allen.neuringer@reed.edu) (Allen Neuringer).

## 2. Method

### 2.1. Participants

Twenty college-age individuals (twelve male), were told that they would be paid between \$8 and \$15, depending on their performance and signed informed-consent forms. After the experiment, all participants were paid \$15. Between one and four participants were scheduled simultaneously for 90 min experimental sessions but the experiment was performed individually at different workstations.

### 2.2. Apparatus

A Java computer program ran on Apple iMac computers in full-screen mode (no window borders, clock, or other operating system interface elements). Participants sat at a comfortable distance from the screen and used a one-button computer mouse to interact with the program. A large circle was constantly present in the center of the screen on a black background (see Figure 1). This circle was subdivided into five wedges with another, smaller circle in the center. Each of the wedges was a different color. Participants moved the mouse pointer (a cross-hair) into the center circle. After a delay of 50 ms, the center circle changed to a lighter gray, indicating that a choice response was permitted. Any mouse movement that passed outside of the center circle before 50 ms elapsed reset the timer. Movement of the cross-hair into one of the wedges marked a response to that alternative, and the wedge flashed to indicate the choice. One of two outcomes resulted from this choice. If the response yielded a reinforcer, a small pulsating dot (approximately 0.75") appeared at a randomized location within the wedge. To receive points, the participant was required to click on the pulsating dot. This led to some number of small dots moving from the location of the pulsating dot to the center circle. Each dot represented one point. As the dots completed their trajectories, they disappeared, and the number of points received was briefly displayed in the center circle (as in "+5"). If the mouse was moved into the center circle without clicking on the pulsating dot (possibly caused by hasty or inattentive responding), the dot and associated points disappeared and were lost. Most responses were not reinforced (as will be described below) and resulted only in the flash of the wedge.

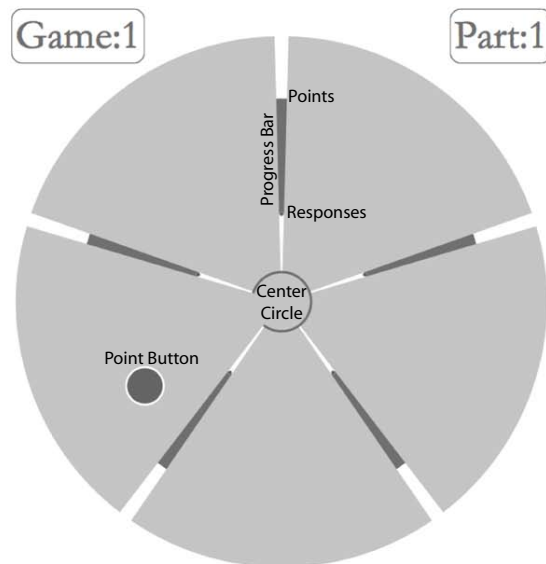


Fig. 1. Grayscale representation of computer screen.

To provide participants with feedback concerning their performances, progress bars moved along the spokes that separated the wedges. With each reinforcer, the front end of the bar moved toward the outer circumference of the circle, the amount of movement scaled such that the trial terminated at 250 points with the bar having reached the outer circumference. With each response, the back end of the bar moved 0.5% of the length of the spoke in the same direction. Thus, the difference between the front end (reinforcers) and back end (responses) indicated reinforcers per response, and the larger this area — referred to as the “progress bar” — the better the performance. Participants were told to “beat” the response end by gaining as many points as possible in as few responses as possible.

### 2.3. Procedure

The schedule of reinforcement differed from a traditional concurrent variable interval schedule in three aspects: number of response alternatives, reinforcement contingency, and, in one of the conditions, reinforcement *decay*. First, participants chose among five different alternatives rather than the more common two. Second, reinforcers were programmed by response-dependent random ratio (RR) schedules with five random-number generators scheduling reinforcers, according to Bernoulli

Table 1  
Conditions and probabilities of reinforcement on each alternative, by phase.

Phase	Condition		Alternative				
	<i>Hold-First</i>	<i>Decay-First</i>	A	B	C	D	E
1	<i>Hold</i>	<i>Decay</i>	0.03	0.09	0.18	0.07	0.00
2	<i>Decay</i>	<i>Hold</i>	0.07	0.03	0.09	0.00	0.18
3	<i>Hold</i>	<i>Decay</i>	0.09	0.00	0.03	0.18	0.07
4	<i>Decay</i>	<i>Hold</i>	0.18	0.07	0.00	0.09	0.03

processes. Each of the five governed a different probability of reinforcer availability, ranging from 0.0 to 0.18 (see Table 1 for specifics). A response to any of the alternatives activated, or “fired,” all five random generators. Thus, every response could cause reinforcers to set up on none, one, or more of the wedges.

We compared two experimental conditions, which differed only as follows: In the *hold* condition, each obtained reinforcer was worth five points, and once a reinforcer was set up, that point value was maintained until the reinforcer was collected. Thus, the *hold* condition is similar to concurrent VI-VI schedules in which reinforcers, once set up, remain available indefinitely and the amount of reinforcement is constant. In the *decay* condition, following set up of a reinforcer, the number of points decreased exponentially upon each response to a different alternative. All values were integers such that the sequence of points per reinforcer was 5, 3, 2, 2, 1, 1, 0, 0, 0 . . . Thus, if a reinforcer had set up on wedge A and A was chosen on the next response, 5 points were provided. However, if, following reinforcer set up, three responses were emitted to wedge C, then a response to A would produce only two points. In addition, as point values decayed, responses to the other wedges continued to activate all five random number generators and, therefore, if the wedge A generator had allocated a new reinforcer prior to receipt of the 2 points, the value would return to its initial 5 point maximum. In brief, in the *decay* condition, points decayed with intervening responses, but there was some probability of the points returning to their initial values.

After participants read written instructions (see Appendix A), one of the authors (JR) demonstrated the interface, explained the progress bar, emphasized that participant pay would be proportional to the size of the bar, and that participants should try to get as many points per response as possible. Ad-

ditionally, after participants were familiar with the interface, a special “demonstration mode” was enabled that allowed participants to see the current values of each of the alternatives as the demonstrator responded. On top of each of the wedges was a number, representing the current value of points in that wedge. For this demonstration, which was in the *decay* condition, the particular probabilities of reinforcement differed from those in the actual experiment but participants were able to see how point-values on a given wedge changed with intervening responses away from that wedge. This experience constituted preliminary training.

Participants then experienced four phases of the experiment, half in the *decay* condition and half in the *hold* condition. In each phase, the programmed probabilities of reinforcement were counterbalanced across the wedges, as shown in Table 1. Twelve of the participants experienced the *decay* condition first and the remainder experienced *hold* first. Thus some participants experienced *decay, hold, decay, hold* and the others experienced the opposite order. Each phase was divided into two consecutive trials, each of which terminated at the successful receipt of 250 points. Participants were informed that the two trials per phase were identical to one another, but that the phases might differ.

At the end of each trial, a score was presented on screen, based on the size of the progress bar and therefore indicating points received per response. Participants were told that they could take a break after each trial with light snacks available.

Participants took approximately 90 minutes to complete the experiment, this being accomplished in one sitting. After the experiment, participants answered a simple survey to record their thoughts about the procedure. These qualitative data were not analyzed and will not be discussed.

In overview, the experiment compared distributions of choices across five alternatives. In one condition, reinforcers worth a constant five points were held indefinitely until obtained (the *hold* condition). In the other, following set up, reinforcer values decreased from five to zero points, this decrease being a function of responses away from the set-up alternative (the *decay* condition). The main question was whether functional relationships between response and reinforcer proportions would rapidly be established under the five-alternative concurrent schedule; and whether *decay* influenced the form of the functions. The guiding hypothesis was that because rapid acquisition of reinforcers was differentially fa-

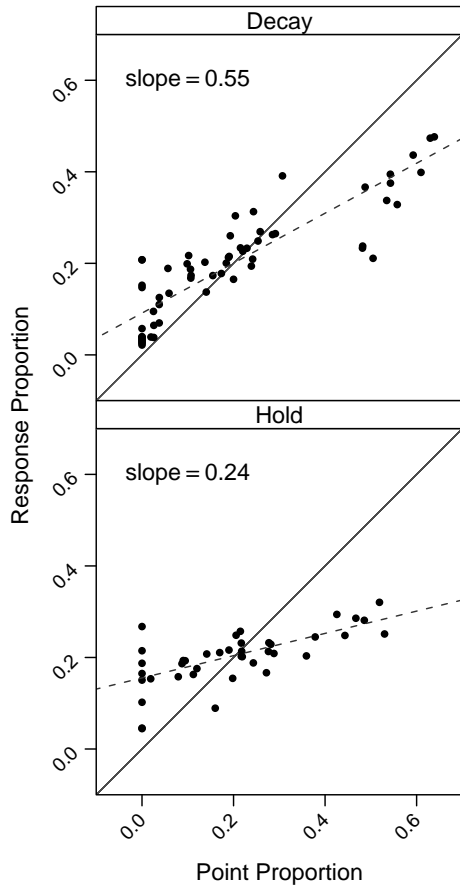


Fig. 2. Response proportions as a function of point proportions. Each participant contributes five points to one of the two subplots of this figure, depending on their initial condition (*decay/hold*). The solid black line is a strict-matching relation and the dashed line is a best fit, calculated by the robust MM-estimator method.

vored in the *decay* condition, choice distributions under *decay* would more closely approximate those from animal experiments where immediate access is presumably highly reinforcing.

### 3. Results

For all analyses, the first and second trials within each phase were combined. This article presents detailed results for Phase 1, during which some participants experienced *decay* and others *hold*.

#### 3.1. Initial Condition

There is discussion in the literature as to how to present data when more than two concurrently reinforced alternatives are available (Aparicio and Cabrera, 2001; Davison et al., 2007). With three alternatives, A, B, and C, Pliskoff and Brown (1976) presented data in the proportional form of Herrnstein (1970), namely,  $\frac{A}{A+B+C}$ , and similarly for B divided by the sum and C divided by the sum. Figure 2 shows response proportions plotted as a function of point proportions<sup>1</sup>, grouped by initial condition.

Figure 3 shows responses as a function of points, each normalized for each participant against the geometric mean of that dimension (Natapoff, 1970). Data shown are all of the participants separated into the *decay* (right) and *hold* (left) conditions. Because the data are presented on log-log coordinates, any responses to a non-reinforced alternative are, of necessity, removed prior to analysis. As a result, each participant contributed four data points (one for each wedge with a probability of reinforcement greater than zero). Obtained reinforcers, normalized to the geometric mean of all reinforced alternatives in the phase, are plotted on the abscissa and similarly normalized responses on the ordinate. Data from the Phase 1 *hold* condition are shown in the right-hand graph. Linear models have been fit to the data after transformation into log-log coordinates, testing a power model. As indicated on the plot, the slope of this function is higher under *decay* than *hold*. Thus, the *decay* condition appears to generate a higher level of sensitivity to differences in reinforcement than the more common *hold*.

Figure 4 provides data for the individual participants in the form of the Herrnstein proportions. This analysis demonstrates the consistency of correlation between response proportions and point proportions as well as the inter-subject variation. The *decay* slopes were generally higher than the *hold*, this difference reaching statistical significance (Welsh two-sample t-test for unequal variances yielded  $t = 2.2074, df = 17.65, p < .05$ .) Five of 8 participants (63%) had slopes lower than 0.3 in the *hold* condition whereas only 3 of 12 (33%) had such low slopes in the *decay* condition.

<sup>1</sup> Reinforcers were calculated in terms of numbers of received *points*, not number of reinforcing events. In the *hold* condition, each reinforcer was worth five points. In the *decay* condition, reinforcer values varied between one and five points.

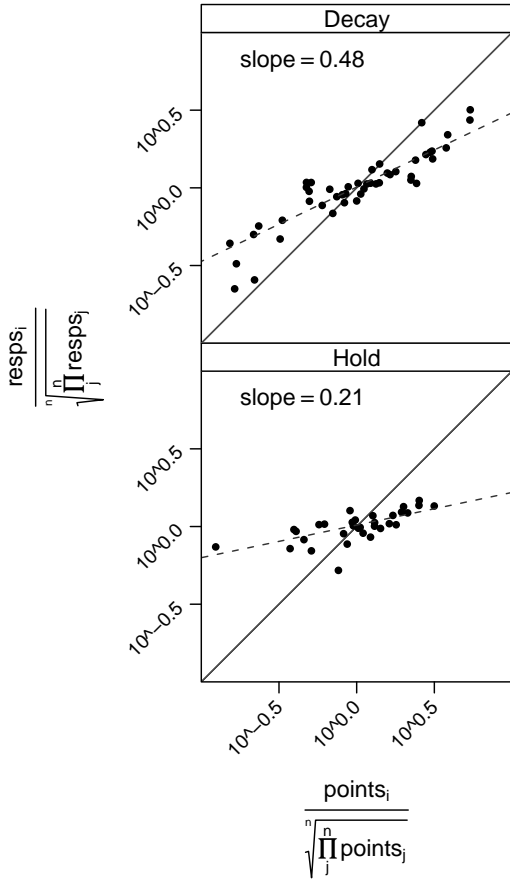


Fig. 3. Plots of geometric-mean normalized responses as a function of geometric-mean normalized reinforcers, all participants in phase 1, grouped by condition (*decay* or *hold*). Each participant contributes at most 4 points to either the left or right graph — one for each alternative that received points. Plots are shown in log-log coordinates, facilitating the examination of a power-law model. The solid black line is a strict-matching relation and the dashed line is an MM-estimator robust regression.

### 3.2. Cumulative responses across the phase

The above analyses were based on total choices and total reinforcers across the two trials of Phase 1. Figure 5 presents a finer-grained analysis: within-phase fluctuations are shown by the cumulative responses on each of the five alternatives for each participant. Stable equiprobable responding would result in five overlapping lines, as in participant *d-6*. Note the relatively constant slopes in most participants' cumulative response records, with only minor fluctuations and adjustments, which often seem to occur abruptly. However, some participants, such as




Fig. 4. Herrnstein proportions plotted for each participant in Phase 1. Response proportions are plotted on the ordinate, and response proportions on the abscissa. The closer the dashed best-fitting (least squares) line is to the solid diagonal strict-matching line, the closer a participant is to a strict (Herrnstein matching) response allocation strategy. *Decay-first* participants are indicated by a *d* and *hold-first* by an *h*.

*h-19* and *d-4*, demonstrated comparatively unstable response patterns.

### 3.3. Testing a stochastic model

Section 3.2 indicated that response generation was generally stable within the phase, and that response proportions were, in most cases, related in orderly fashion to reinforcement proportions. There is debate in the literature as to the underlying process by which such molar order is obtained (Nevin, 1969, 1979; Williams, 1988). Preliminary analyses indicated that choices in the present experiment could be modeled by a predominantly-stochastic process, but one that was biased in a readily identifiable manner. To test for stochasticity, we compared proportions of response pairs to

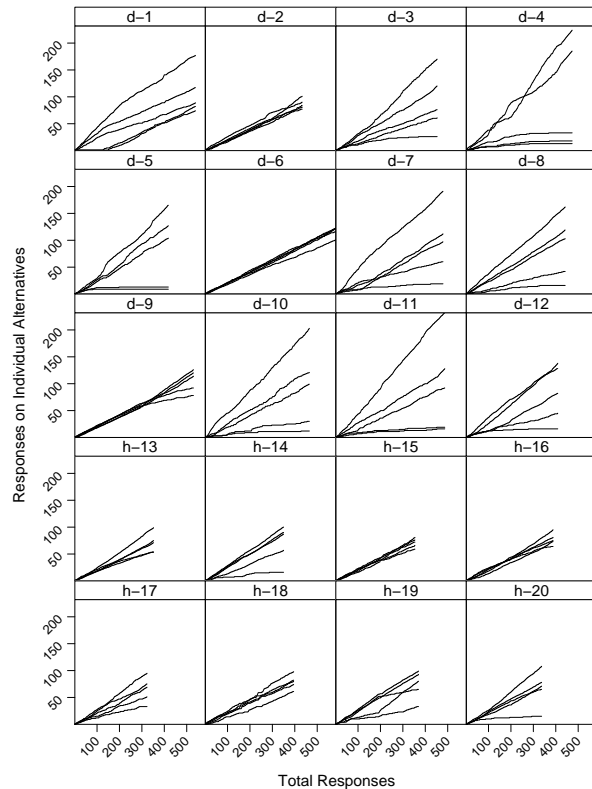


Fig. 5. Cumulative responses on each alternative as a function of total responses.

the proportions expected by a first-order Markov chain (Chatfield, 1973). Participants switched from one alternative to another more frequently than expected by a purely stochastic model. Figure 6 shows proportion of switch behavior on the ordinate, plotted against the proportion that would be expected by a weighted stochastic generator with the same first-order response distribution. The preponderance of points above the diagonal indicates that switches were more frequent than predicted by the state-independent stochastic model.

When a switch was made, the particular alternatives chosen were well modeled by a stochastic process. Figure 7 shows, for each participant separately, observed proportions of response pairs as a function of predicted proportions from the stochastic model. This analysis is based on a subset of response pairs, excluding repetition (*AA*, *BB*, *CC*, *DD*, & *EE*). The points are clustered closely around the diagonals, indicating that first-order emission proportions predicted second-order proportions as would be the case if choices were stochastic, considering only the non-

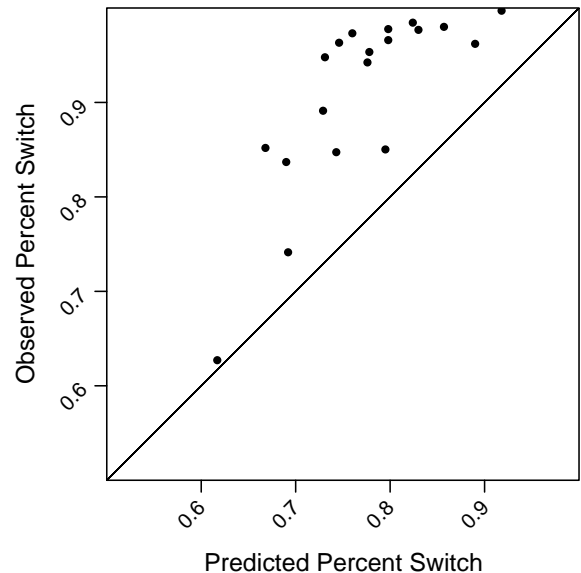


Fig. 6. Phase 1 percentage of switch behavior (response to an alternative other than the immediately previous), as compared to the percentage of such behavior expected for a simple stochastic model. Each participant contributes one point to this plot.

repeating responses.

### 3.4. Feedback Concerns

One concern was that the high programmed probabilities of reinforcement might have permitted responses to “drive,” through feedback, the proportions of obtained reinforcers. Across the whole experiment, the correlations (Pearson’s product-moment) between programmed and obtained reinforcer proportions were 0.97 for both *decay* and *hold* conditions. This high correlation indicates that the reported relationships between obtained reinforcers and choices were not an artifact of responses influencing programmed reinforcers.

## 4. Discussion

Most studies of choice in concurrent schedules of reinforcement involve binary choices and time-based schedules of reinforcement. The present study with human participants differed in a number of ways. Many of these have previously been employed but not combined as in the present case. We will discuss the main attributes of our procedure.

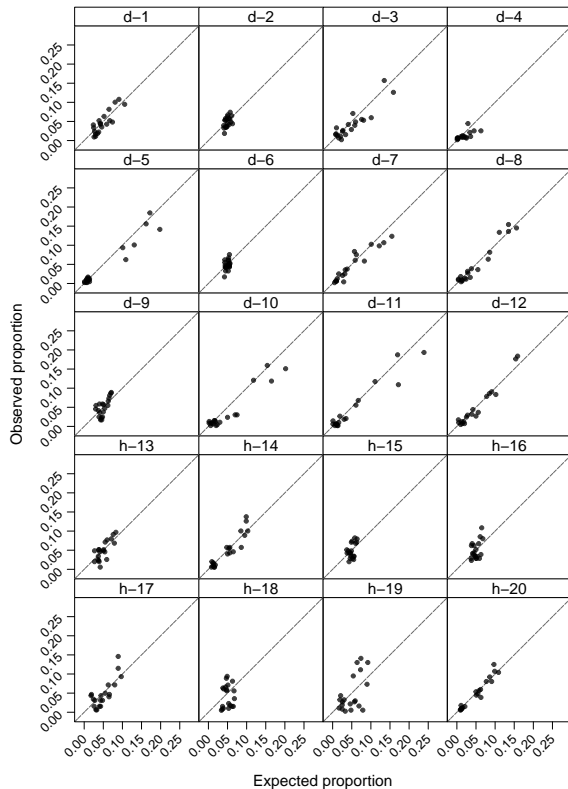


Fig. 7. Phase 1 observed proportions of individual two-response pairs as a function of expected proportion of emission of that pair by a stochastic model that never repeats. Only data from switch pairs are plotted, and proportions are relative to switch pairs only. Individual participants are shown.

Participants chose among five alternatives (implemented as on-screen wedges). Among the few other studies employing more than two operanda, Pliskoff and Brown (1976) observed close approximations to strict matching by pigeons under a three-alternative choice procedure. This study involved a Findley (1958) switch procedure in which three different colors on one response key were associated with three different VIs and responses to a second key caused random selection among the colors. Davison (1996) employed a similar procedure. Elsmore and McBride (1994) studied rats in an eight-arm radial maze, with each arm providing a different rate of reinforcement. The data were well described by power functions, as described above. Aparicio and Cabrera (2001) studied a “barrier choice” paradigm in which rats were required to select among as many as eight alternatives, with random-interval schedules of reinforcement. Again, power-function relationships were observed. In general, orderly data

have been generated under multiple alternative concurrent reinforcement studies (however, see Davison et al., 2007). One advantage of such studies is the ability to compare multiple reinforcement values at one time. Also, such procedures model real-world situations where many choices are often simultaneously available. A difficulty with the procedure (and perhaps a reason it is comparatively under-studied) is that data presentation and analysis are not as well established as with binary choice (for discussion, see Aparicio and Cabrera, 2001; Davison et al., 2007). We presented Herrnstein’s proportion analysis employed by Pliskoff and Brown (1976) and the geometric-mean normalization (Natapoff, 1970). Reassuringly, concordance between these analyses predominated.

Reinforcement availability was probabilistic and response-driven. However, unlike independent ratio schedules, in which responses to one alternative advance the counter only on that alternative (e.g. Herrnstein and Loveland, 1975), the present procedure implemented a dependent variant of this procedure. Each response activated random-number generators for all five alternatives. This type of dependent random ratio schedule, although rarely used, has resulted in orderly relationships between response and reinforcement proportions (Lau and Glimcher, 2005; Macdonall, 1988; Meisch and Spiga, 1998). The primary advantage of such response-dependent probabilistic programming is that it avoids confounding programmed reinforcement rates (as under concurrent VI-VI schedules) with rates of responding and contingency-irrelevant pauses (water-drinking, scratching, etc.). Analyses of relationships between reinforcements and responses are facilitated in terms of momentary maximization of reinforcement. However, one possible disadvantage is that real-world responding in one location rarely drives reinforcer availability in other, possibly distant locations; and, furthermore, reinforcers often become available or unavailable in time. Stated simply, concurrent VI-VI schedules may better represent real-world conditions than the present concurrent dependent random ratio schedules.

A third aspect of the present procedure was the examination of a decay contingency. After a reinforcer was activated for a particular alternative, the sooner the participant responded on that alternative, the higher the point value of the reinforcer. Extended avoidance of the alternative led to removal of the reinforcer altogether. This contingency is similar to limited-hold procedures, where reinforcers are

available only for a limited period of time following set-up (see Buskist and Morgan, 1987; Morse, 1966). Also similar are contingent magnitude-of-reinforcement schedules in which magnitudes depend upon response latencies or other attributes of the response (Buskist et al., 1988; Lippman and Tragesser, 2003). One advantage of this type of contingency is that sensitivity to reinforcement is relatively high, as demonstrated by previous work and corroborated by the present results.

The last aspect of note is that we studied the performance of human participants. Most research on concurrent schedules has been with animal models, such as pigeons, rats, mice, or monkeys, but more than 25 human studies have been described in excellent reviews by Bradshaw and Szabadi (1988) and by Kollins et al. (1997). These reviews indicate that people manifest lower power-law slopes than nonhuman animals, indicating lower sensitivity to changes in reinforcements, and greater variability in those slopes. Kollins et al. (1997) offer many hypotheses to account for such differences, but ultimately implicate procedural differences. The present method attempted to increase both attention to task and motivation for immediate reinforcement, two factors thought to be especially important.

The main results were rapid acquisition of stable performance. Within each phase, response allocations were, for the most part, quite constant after initial experience in the phase. Furthermore, response proportions (or ratios) were functionally related to reinforcer proportions. However, the slopes of the functions were relatively low (in both Herrnstein proportion and Baum power plots) and there was a high degree of inter-subject variability. Both of these aspects — slope values and variability — were closer to those in animals studies under the *decay* condition than under *hold*. Thus, consistent with conclusions by Bradshaw and Szabadi (1988), and by Kollins et al. (1997), procedural details may contribute importantly to how choice distributions relate to obtained reinforcements.

Present results may help to resolve an ongoing debate as to whether animals and people respond in a stochastic fashion under concurrent reinforcement contingencies or act in a way that momentarily maximizes the chances of reinforcement (e.g. Nevin, 1979; Williams, 1988). Participants in our study were more likely to switch among alternatives than predicted by a stochastic model. When a switch was made, the location of that switch (among the four remaining possibilities) could generally be de-

scribed by a stochastic model.

Many reasons have been offered for high switch rates in both non-human animal and human studies, but of primary importance is that probabilities of reinforcement are higher for switches than for repetitions. This is true under both concurrent VI-VI schedules and concurrent random ratio schedules of the dependent form, as in the present experiment. Thus, participants may have responded to differential reinforcement for switching, but because of uncertainty as to the location of the next reinforcer, they chose stochastically among the available options. Stated differently, momentary-maximizing responding may occur when subjects can discriminate differential reinforcement, but when such discriminations are not possible, stochastic responding results. Machado (1992) initially proposed this hypothesis and our results conform to his predictions.

## Appendix A. Instructions

Your task is to earn as many points as possible in some gambling-type games. To motivate interest, and hopefully make the games fun, you will receive a monetary bonus for your score – the better the score, the larger the bonus. You will begin with \$8 but can earn up to \$15. Your score will be indicated by a bar that grows and shrinks as you play. The LARGER the bar at the end of a trial, the more money earned.

You will play a total of five games. Some of the games will be more difficult than others. Don't be upset if points are scarce that will sometimes be the case. But in all games, try to earn as many points as possible. You'll note that points make the bar grow larger and responses without points cause the bar to shrink.

Each of the games is divided into two parts. In a given game, the parts are identical, so use your experiences in the first part to maximize points during the second.

In all of these gambling-type games, your score is partly due to chance and partly to skill. There are no tricks and no complex hidden patterns to be discovered. Your selections will determine your points, but it is not possible to be reinforced on all (or, in some cases, on many) choices. We will show you examples of similar games before you actually begin. The examples differ from the ones you will actually play.

This method is being developed with the intention of serving as an assessment tool for age- and

disability-related cognitive disorders. You are helping to provide baseline control data against which these populations will be compared. We're quite hopeful for the potential of this line of research and ask that you take it seriously, as well. If you aren't interested in focusing on this task, please feel free to either withdraw or reschedule at this point.

Take as much time as you'd like. Your speed will not influence the score. Feel free to take breaks, eat snacks, or stretch between games.

Thank you and good luck,

Allen Neuringer

## References

- Aparicio, C. F., Cabrera, F., 2001. Choice with multiple alternatives: The barrier choice paradigm. *Mexican Journal of Behavior Analysis* 27, 97–118.
- Baum, W., 1974. On two types of deviation from the matching law: bias and undermatching. *J Exp Anal Behav* 22 (1), 231–242.
- Baum, W. M., 1979. Matching, undermatching, and overmatching in studies of choice. *J Exp Anal Behav* 32 (2), 269–281.
- Bradshaw, C. M., Szabadi, E., 1988. Quantitative analysis of human operant behavior. In: Davey, G., Cullen, C. (Eds.), *Human Operant Conditioning and Behavior Modification*. John Wiley & Sons, New York, Ch. 11, pp. 225–259.
- Buskist, W., Morgan, D., 1987. Competitive fixed-interval performance in humans. *J Exp Anal Behav* 47 (2), 145–158.
- Buskist, W., Oliveira-Castro, J., Bennett, R., 1988. Some effects of response-correlated increases in reinforcer magnitude on human behavior. *J Exp Anal Behav* 49 (1), 87–94.
- Chatfield, C., 1973. Statistical inference regarding markov chain models. *Applied Statistics* 22 (1), 7–20.
- Davison, M., 1996. Stimulus effects on behavior allocation in three-alternative choice. *J Exp Anal Behav* 66 (2), 149–168.
- Davison, M., Krageloh, C. U., Fraser, M., Breier, B. H., 2007. Maternal nutrition and four-alternative choice. *J Exp Anal Behav* 87 (1), 51–62.
- Elsmore, T. F., McBride, S. A., 1994. An eight-alternative concurrent schedule: foraging in a radial maze. *J Exp Anal Behav* 61 (3), 331–348.
- Findley, J. D., 1958. Preference and switching under concurrent scheduling. *J Exp Anal Behav* 1 (2), 123–144.
- Herrnstein, R., 1970. On the law of effect. *J Exp Anal Behav* 13 (2), 243–266.
- Herrnstein, R., Loveland, D., 1975. Maximizing and matching on concurrent ratio schedules. *J Exp Anal Behav* 24 (1), 107–116.
- Kollins, S. H., Newland, M. C., Critchfield, T. S., 1997. Human sensitivity to reinforcement in operant choice: How much do consequences matter? *Psychon Bull & Rev* 4, 208–220.
- Lau, B., Glimcher, P. W., 2005. Dynamic response-by-response models of matching behavior in rhesus monkeys. *J Exp Anal Behav* 84 (3), 555–579.
- Lippman, L. G., Tragesser, S. L., 2003. Contingent magnitude of reward in modified human-operant drl-1h and crf schedules. *Psychol Rec* 53, 429–442.
- Macdonall, J., 1988. Concurrent variable-ratio schedules: Implications for the generalized matching law. *J Exp Anal Behav* 50 (1), 55–64.
- Machado, A., 1992. Behavioral variability and frequency-dependent selection. *J Exp Anal Behav* 58 (2), 241–263.
- Meisch, R. A., Spiga, R., 1998. Matching under non-independent variable-ratio schedules of drug reinforcement. *J Exp Anal Behav* 70 (1), 23–34.
- Morse, W. H., 1966. Intermittent reinforcement. In: Honig, W. K. (Ed.), *Operant Behavior: Areas of Research and Application*. Appleton-Century-Crofts, New York, pp. 52–108.
- Natapoff, A., 1970. How symmetry restricts symmetric choice. *Journal of Mathematical Psychology* 7, 444–465.
- Nevin, J., 1969. Interval reinforcement of choice behavior in discrete trials. *J Exp Anal Behav* 12 (6), 875–885.
- Nevin, J. A., 1979. Overall matching versus momentary maximizing: Nevin (1969) revisited. *J Exp Psychol Anim Behav Process* 5, 300–306.
- Pliskoff, S., Brown, T., 1976. Matching with a trio of concurrent variable-interval schedules of reinforcement. *J Exp Anal Behav* 25 (1), 69–73.
- Williams, B. A., 1988. Reinforcement, choice, and response strength. In: Atkinson, C. R., Herrnstein, R. J., Lindzey, G., Luce, R. D. (Eds.), *Stevens' handbook of experimental psychology*. Vol. 2. Wiley, pp. 167–244.