Operant variability when reinforcement is delayed

KATIE WAGNER and ALLEN NEURINGER
Reed College, Portland, Oregon

Operant responses are often weakened when delays are imposed between the responses and reinforcers. We examined what happens when delayed reinforcers were contingent upon operant response *variability*. Three groups of rats were rewarded for varying their response sequences, with one group rewarded for high variability, another for middle, and the third for low levels. Consistent with many reports in the literature, responding slowed significantly in all groups as delays were lengthened. Consistent with other reports, large differences in variability were maintained across the three groups despite the delays. Reinforced variability appears to be relatively immune to disruption by such things as delays, response slowing, prefeeding, and noncontingent reinforcement. Furthermore, the small effects on variability depended on baseline levels: As delays lengthened, variability increased in the low group, was statistically unchanged in the middle group, and decreased in the high group, an interaction similar to that reported previously when reinforcement frequencies were lowered. Thus, variable operant responding is controlled by reinforcement contingencies, but sometimes differently than more commonly studied repetitive responding.

When reinforcement is plentiful, operant responding often becomes repetitive in form and pattern. When conditions worsen, as when reinforcement is degraded or withheld, responding tends to vary more, this serving possibly to enable the organism to discover new routes to reward (Balsam, Deich, Ohyama, & Stokes, 1998; Neuringer, Deiss, & Olson, 2000). In addition, variability is sometimes engendered by reinforcers that are contingent upon variability itself. This is a strange idea for some, but when reinforcers depend on high variability, highly variable responding is generated and maintained. Such reinforcement of variability has been reported for response topographies (see, e.g., Goetz & Baer, 1973; Pryor, Haag, & O'Reilly, 1969), response rates (e.g., Blough, 1966), and most commonly studied, response sequences (e.g., Machado, 1989; Stokes, 1999; Stokes & Harrison, 2002). As an example, when rats were rewarded for varying their patterns of four left and right leverpresses, highly variable response sequences were maintained across many sessions (Cherot, Jones, & Neuringer, 1996). Related studies show similar effects with human participants (e.g., Miller & Neuringer, 2000). Furthermore, reinforcement contingencies can specify particular levels of variability—high, intermediate, or low (Grunow & Neuringer, 2002; Jensen, Miller, & Neuringer, 2006; Machado, 1989). These, together with other studies demonstrating discriminative control—animals learn to vary in the presence of one stimulus and

repeat in another (Cohen, Neuringer, & Rhodes, 1990; Denney & Neuringer, 1998)—have led to the conclusion that variability is an operant dimension of behavior, controlled by reinforcers in much the same manner as are other operant dimensions, such as response speed, force, location, and topography (Neuringer, 2002).

However, operant variability may have some characteristics that differ from those of more commonly studied repetitive responses. A case in point is that reinforced variability appears to be relatively immune to interference from events that weaken repeated operants. For example, Doughty and Lattal (2001) showed that prefeeding pigeons prior to experimental sessions had less of an effect on operant variability than on operant repetitions, and the same held when noncontingent reinforcers were presented during the sessions. Similarly, alcohol severely interfered with rats' performances when a particular sequence of responses had to be repeated but not when response variations were required (Cohen et al., 1990; McElroy & Neuringer, 1990). Another example: When responses in a sequence were systematically slowed by introduction of interresponse time-outs, ability to repeat a particular sequence was increasingly impaired; but when variability was required, performance was either improved or maintained at high level (Neuringer, 1991).

In a related study—one that led to the present experiment—reinforcement rates and variability contingencies were both manipulated in order to compare their effects (Grunow & Neuringer, 2002). Each of four groups of rats was reinforced for a different level of response-sequence variability, from low to high. Large differences in variability were maintained by these contingencies, even when reinforcement rates were systematically lowered—from continuous reinforcement (CRF; every sequence that met the respective variability contingency was reinforced)

This research was supported in part by Research Grant MH068259 from the National Institutes of Health. A portion of the data was submitted by K.W. as a senior thesis to Reed College. We thank Greg Jensen for help with programming and data analysis, and Gene Olson for care of animals. Correspondence concerning this article may be directed to A. Neuringer, Department of Psychology, Reed College, Portland, OR 97202 (e-mail: allen.neuringer@reed.edu).

through variable interval (VI) 1 min (reinforcement could occur no more than once per minute) to VI 5 min (reinforcement no more frequent than once per 5 min). However, rates of responding decreased in all groups as reinforcements became less frequent. Thus, response speed seemed to be controlled by a different aspect of reinforcement than response variability; *speed* was controlled by reinforcement rates (and not reinforcement contingencies), whereas *variability* was controlled mainly by reinforcement contingencies (and not reinforcement rates). We say “mainly,” because a significant interaction was observed: As reinforcers became less frequent, variability decreased in the high group of rats and increased in the low group. Although these effects were small, they were statistically significant. Thus, when baseline variability was high, degrading reinforcement (by lowering its rate) caused variability to fall, but when baseline variability was low, degrading reinforcement caused the opposite: Variability increased.

In the present study, we attempted to assess the effects of *delayed reinforcers* on operant variability and to test the generality of the Grunow and Neuringer (2002) findings. In another recent study, Odum, Ward, Barnes, and Burke (in press) observed an effect consistent with those described above—namely, that delaying reinforcers adversely affected operant repetitions but not operant variability. In the present case, we tested whether the particular levels of variability maintained by the reinforcement contingencies influenced the effects. Each of three groups of rats was reinforced for a different level of response variability. Delays of reinforcement—intervals between the last response in a trial and reinforcement—were systematically manipulated. We predicted that rates of responding would slow as delay durations increased (see, e.g., Reilly & Lattal, 2004; Richards, 1981) but that the reinforcement contingencies would maintain different levels of variability in the three groups (see, e.g., Doughty & Lattal, 2001; Grunow & Neuringer, 2002). Finally, we sought to test for an interaction between reinforcement delay and reinforcement contingency.

Since aspects of the procedure are unusual, we provide here a brief preview. The chamber contained two levers on the front wall and three keys—only two were used in the present experiment—on the rear. A *trial* consisted of a sequence of three responses across these four operanda. Variability of sequences was required for reinforcement, with different minimum levels across three groups of rats. For the high-variability group, reinforcement was provided only for the sequences that had occurred relatively infrequently, as defined by a percentile variability contingency (see below), whereas for the low-variability group, the contingencies were much more permissive. A middle group experienced contingencies of intermediate difficulty. In the delay condition of the experiment, each reinforcement was preceded by a delay—a blackout period during which the chamber was dark and responses ineffective (systematically manipulated between 0.5 and 16 sec). In a control time-out condition, reinforcers were given immediately at the end of a trial, with blackouts presented *after* each

reinforcement. Thus, in the delay condition, a blackout *preceded* the reinforcement, whereas in the control, the blackout *followed* the reinforcement. We expected that reinforcement would be degraded (as indicated by, e.g., changes in response rates) only by the delays.

EXPERIMENT 1

Method

Subjects

Thirty experimentally naive male Long-Evans rats, obtained from Simonson Labs in Gilroy, CA, approximately 11 weeks old at start of training and 25 weeks old at start of the experimental manipulations, were housed in pairs (one cage housed 3) and maintained on a 12:12-h light:dark cycle with water continuously available. The rats were food deprived for approximately 21–22 h, with 1 h of free food after each experimental session. Experimental sessions were conducted on 6 or 7 days per week.

Apparatus

Ten Apple Macintosh eMac computers were connected to 10 modified Gerbrands (Cambridge, MA) operant chambers via Meta-Research (Portland, OR) Benchtop interfaces. The chambers measured 27 cm deep $\times$ 30.5 cm wide $\times$ 29 cm high, with a front, back, and ceiling made of Plexiglas, sidewalls of aluminum, and a floor of stainless steel rods. Each chamber was equipped with two levers, each 5.1 cm wide, situated side by side, 8.9 cm apart. Between the two levers, on the same sidewall, was a pellet tray into which 45-mg BioServe (Frenchtown, NJ) pellets were presented as reinforcers. On the opposite wall were three response keys, each 2 cm in diameter, but only the right and left keys were lit and activated during the experiments. The keys were opaque plastic lenses that were lit with a 28-V DC light bulb when active. When the levers were active, a 28-V DC light bulb recessed behind an opaque plastic lens that was flush with the chamber wall was lit. An additional 28-V DC bulb served as the houselight. A speaker was located outside of the chamber but within the sound- and light-attenuating outer box that surrounded the chamber. The controlling computer program was written in True BASIC.

Procedures

Preliminary training. Each rat was shaped to respond on the keys and levers and then introduced to a percentile variability reinforcement contingency, these being identical for all animals during the preliminary phase.

Trial definition. Each trial consisted of a sequence of three responses across the four operanda: left lever (L), right lever (R), left key (X), and right key (Z). At the beginning of a trial, lights behind X and Z and above L and R, as well as the houselight, were illuminated, indicating that all four operanda were *activated* and that a response could occur on any of them. A response to an operandum was recorded and resulted in inactivation of that operandum together with darkening of its corresponding light. The rat could now respond to any of the remaining three activated (and lit) operanda, following which that operandum became inactive with the other three lit and active. Thus, a trial consisted of one of $4 \times 3 \times 3 = 36$ possible patterns (e.g., LRX, RZR, XZL), and so on. It was possible for two responses to be emitted on the same operandum within a given trial, but not consecutively (e.g., LRL was an acceptable sequence, but LLR and LRR were impossible). Thus all responses, other than the first in a trial, required switching from one operandum to another. Preventing repetitions on a given operandum avoided the problem of incidental repetitions or “chatter” without the necessity of inserting arbitrary time-out pauses following each response. The topographically different operanda assured different sequence probabilities at baseline (see Neuringer, Kornell, & Olufs, 2001).

During the preliminary training phase, trials ended with one of two consequences (see Figure 1). If the sequence met the variability

Figure 1. A diagram of the procedure, showing the sequence of events for a reinforced trial in the delay-of-reinforcement condition (top); a reinforced trial in the postreinforcement time-out condition (middle); and for both, a nonreinforced trial (bottom). The larger ovals indicate a trial, the smaller circles indicate individual responses within the trial, a smiling face indicates a food pellet reinforcement, and the narrow rectangle indicates a blackout during which the chamber was dark and the response operanda inactivated.

contingencies (to be described below), then immediately following the third response, a reinforcer paired with a 1200-Hz, .25-sec tone was delivered and directly followed by a 2-sec time-out, during which the chamber was dark. Responses during the time-out had no effect, except that within 1 sec of the end of the time-out period, responses delayed termination of the time-out for 1 sec. If the sequence failed to meet the variability contingency, the 2-sec time-out occurred without reinforcer or tone. The rats generally withheld responses during these time-out periods.

One session was provided in which the rats were reinforced for each of their first 25 sequences, no matter the sequence (i.e., variability was not required). In all remaining sessions of the experiment, percentile variability contingencies were in force.

Percentile variability contingencies. At first, all rats were reinforced for varying their sequences under a .50 percentile schedule based on the previous 25 sequences as will be described below (see Galbicka, 1994, for a description of the percentile reinforcement contingencies). This will be referred to as a percentile variability 50% schedule. Thirty-six counters were maintained by the com-

puter, 1 for each of the 36 possible sequences. Whenever a sequence was emitted, a count would be added to the appropriate counter, and relative frequencies were computed by dividing the frequency of a given sequence by the sum of the frequencies of all of the sequences. If the relative frequency of the emitted sequence was lower than a particular *threshold*, the trial would end with reinforcement. The threshold was a variable whose value changed dynamically in the following manner. On any given trial, the 50th percentile of the preceding 25 sequences was defined as the threshold for reinforcement. This means that the rats were reinforced for sequences that they emitted that were less frequent, as determined by relative frequencies, than 50% of their preceding 25 sequences. All sequence frequencies were maintained across sessions and loaded into the computer at the beginning of each session.

Because probabilities of emitting different sequence patterns were likely to change across time, a weighting coefficient of .97 was applied as follows. At the end of each trial, the cumulative count in each of the 36 sequence counters was multiplied by the coefficient. Therefore, the most recently emitted sequence contributed 1 to the count, the second

most recently emitted sequence contributed .97 to the count, and the third most recently emitted sequence contributed a count of $.97 \times .97$ (i.e., .94). These weighted counts were placed into the percentile "stack" and were used to compute whether a given relative frequency was sufficiently low for reinforcement, as will be described.

During this preliminary training (two to five sessions), several rats began to emit exclusively only the two sequences that contained leverpresses RLR and LRL. This was problematic, because if responses to the keys extinguished, it would be impossible to observe changes in variability. Therefore, we introduced a minimum relative frequency requirement of .40 in addition to the percentile variability criterion. Now, in order for a sequence to be reinforced, its relative frequency not only had to meet the threshold set by the .50 percentile criterion, but also had to have been emitted in fewer than 40% of the trials (weighted by the weighting coefficient). This resulted in all rats responding on all four operanda, at least occasionally.

Sessions terminated after either 100 reinforcers had been presented or approximately 2 h, whichever came first.

Low-, middle-, and high-variability reinforcement contingencies. The rats were quasi-randomly divided into three groups, each of which was based on the final three sessions under percentile variability 50% contingencies, with groups equalized as much as possible in terms of *U* values, a measure of variability, and speed of responding. (Owing to an error in subject designation, group sizes were 9, 10, and 11.) After group division, the percentile for the middle group ($n = 9$) remained at 50%, the percentile for the low group ($n = 11$) was changed to 75%, and the percentile for the high group ($n = 10$) was gradually changed to 20% across 10 sessions. For the low group, reinforcement was possible if the relative frequency of the current trial was among the least frequent 75% of the previous 25 sequences, generally an easy criterion to meet. The high group's current sequence had to be among the least frequent 20%, a much more difficult criterion. After a few sessions, the percentile stack was reduced from 25 sequences to 15, with this value maintained throughout the remainder of the experiment.

Tables 1A and 1B provide an example of the percentile procedure and show how it determined the thresholds for reinforcement. Table 1A shows the 20 trials preceding the current sequence, ZXI. The most recent sequence is at the top; the least recent at the bottom. Note that a moving window consisting of only the 15 sequences prior to the current trial was used to establish the percentile "stack" located in Table 1B. The same data are shown there, now ordered in terms of relative frequencies at time of emission, from least frequent to most frequent. Also shown are the current criteria values for the high, middle, and low contingency groups. The high threshold was computed by multiplying $.2 \times 15 = 3$. Therefore, if the current sequence's relative frequency was lower than or equal to the third value from the top, reinforcement was provided. The middle was $.5 \times 15 = 7.5$. An "integer" operator changed this to 7, so that for the middle group, if the current relative frequency was less than or equal to the 7th from the top, then reinforcement was provided; and low was computed similarly, but here it is 11 from the top. Therefore, if these data represented an animal in the high group, then the current trial (ZXI, with a current relative frequency value equal to 0.0556) would *not* be reinforced (since its value was higher than the high threshold of 0.0321), nor would an animal in the middle group be reinforced. On the other hand, if the data were from a low animal, reinforcement would be provided.

Because the low-variability group's contingencies were more permissive (relatively frequent sequences could be reinforced, as well as rare ones) than those for the middle group, which, in turn, were more permissive than the stringent requirement for high variability in the remaining group, reinforcement rates were expected to differ: The high-variability group would be reinforced much less frequently than the middle-variability group, which would be reinforced less than the low-variability group. In order to avoid this confound, the middle- and low-variability groups were reinforced probabilistically for meeting their assigned contingencies, with these probabilities

Table 1A
Percentile Variability Calculations Showing Most Recent 21 Trials in Order of Emission

How Many Trials Ago	Sequence	Relative Frequency
Current	ZXI	0.0556
1	RXI	0.0452
2	ZXI	0.0533
3	XRI	0.031
4	XLR	0.0748
5	ZRI	0.063
6	XIX	0.0193
7	XRI	0.0442
8	XLR	0.0854
9	LXI	0.1006
10	ZRI	0.0734
11	LXI	0.0828
12	RXI	0.0321
13	RZI	0.0677
14	XLR	0.093
15	RXI	0.0516
16	XLR	0.1052
17	XLR	0.1187
18	RZI	0.0829
19	LXI	0.0884
20	LXI	0.0281

Table 1B
Percentile Variability Calculations Showing Most Recent 15 Trials Ordered from Lowest to Highest Relative Frequencies

How Many Trials Ago	Sequence	Relative Frequency
6	XIX	0.0193
3	XRI	0.031
12	RXI	0.0321
High Variability Threshold		
7	XRI	0.0442
1	RXI	0.0452
15	RXI	0.0516
2	ZXI	0.0533
Middle Variability Threshold		
5	ZRI	0.063
13	RZI	0.0677
10	ZRI	0.0734
4	XLR	0.0748
Low Variability Threshold		
11	LXI	0.0828
8	XLR	0.0854
14	XLR	0.093
9	LXI	0.1006

intended to equate overall frequencies of reinforcement across the three groups. Namely, the middle group was reinforced for meeting the percentile variability contingency with a probability of .4, and the low group with a probability of .3, with these values assigned after responding had been simulated under the different contingencies. Trials that did not meet the probability criterion, even if they did satisfy the variability contingency, resulted in a blackout.

Twenty-nine sessions of training were provided, during which some of the parameters, such as time-out durations, were varied to generate stable responding.

Experimental Conditions

Postreinforcement time-out. This condition served as a control for the delay-of-reinforcement phase to follow. Pilot research in our

lab had indicated that effects of delays on variability occurred rapidly (an observation that was supported in Experiment 2, below) and in this and the following phase, each value of time-out (or delay) was presented for two consecutive sessions. As during the training phase described above, a time-out occurred at the end of every trial. On reinforced trials, the time-out was presented immediately after the reinforcement (see Figure 1). The time-out durations, in the order experienced, were 0.5, 16, 2, 8, and 32 sec. (The 16-sec value was run for three sessions because of a computer error during the first session.)

Prereinforcement delay. The prereinforcement delay condition followed immediately after the postreinforcement phase and was identical to it except that on reinforced trials, the reinforcer was administered *after* the blackout—here to be referred to as a delay so as to be consistent with the literature—as shown in Figure 1. Therefore, reinforced trials ended with a delay that was followed by food, and nonreinforced trials ended only with a blackout that was identical in all ways to the delay. This method assured that no differential stimulus cues were associated with the delay-to-reinforcement period; that is, the delay could not itself serve as a reliable cue for reinforcement. Except for the location of blackouts, this delay condition was identical to the time-out condition described above. Four delay durations were studied: 0.5, 8, 2, and 16 sec, each for two sessions. The 32-sec interval was omitted because response rates fell to low levels during the 16-sec delays. Since within-subjects statistical comparisons will be made across conditions, we did not include the postreinforcement 32-sec time-outs in any of the analyses.

Overview of procedure. Three groups of rats were reinforced for sequence variations, one group for low variations, another for middle levels, and a third for highly variable sequences. Reinforcement depended on a sequence meeting the variability contingency (low, middle, or high, respectively) as well as a probability generator (to maintain equal reinforcement frequencies across groups). The variability contingency served as one independent variable (between groups). For each group, all nonreinforced trials ended with a blackout period during which the chamber was dark and all operanda were inactivated. During the delay phase of the experiment, the same blackouts *preceded* reinforcers, whereas during the time-out control phase, immediate reinforcers were *followed* by the blackouts, forming the second independent variable (within-groups repeated measures). The duration of the blackouts was the third independent variable (also within-groups repeated measures). The main questions were how changing the duration of the blackouts would affect variability, and whether an interaction between blackouts and contingency would be found.

Measures. U values served as a measure of variability and were computed as follows:

$$-\sum_{i=1}^n \{ \text{RF}_i [\log(\text{RF}_i)] \} / \log(n),$$

where n equals the number of possible sequences, and RF refers to the unweighted relative frequency of each of the n sequences. U values range from 1 (*most variable*) to 0 (*least variable*) and measure the distribution of the sequence frequencies (see Denney & Neuringer, 1998). U values were computed separately for each condition for each rat. Because U values are sensitive to number of instances, numbers of trials were equalized across each of the delay and time-out values, this being done individually for each subject but separately for the delay and time-out conditions. Thus, within each of these conditions, the lowest number of trials for a given duration determined the number used for all durations. The data between the beginning of each session and this number were employed for the U -value analyses. For example, suppose that a rat emitted 150 sequences in the first session of the 16-sec delay, 200 sequences in the first session of the 8-sec delay, 250 sequences in the first session of the 2-sec delay, and 300 sequences in the first session of the .5-sec delay. Because only 150 sequences were generated at the 16-sec delay, only the first

150 sequences from each of the other delay lengths were used for analyses. The data from the second session of each delay length were similarly truncated. The two sessions' data were then combined and one U value was calculated from the total, with this value serving as the measure of variability. Median rates of responding were also calculated from the same data set. These rates were computed by dividing the total number of trials emitted in a session by the time during which responses were possible—that is, the time during which operanda were activated.

Results

Variability

We first consider U values as a function of (1) the variability-reinforcing contingencies (reinforcement of high, middle, or low levels of variability); (2) the delay condition (prereinforcement delay vs. postreinforcement time-out); and (3) the durations of delays and time-outs (0.5, 2, 8, and 16 sec). The results are shown in Figure 2, with least-squares lines fit to the data points. An omnibus ANOVA ($3 \times 2 \times 4$) yielded a number of significant effects. First, variability increased as the reinforcement contingencies became increasingly demanding [$F(2,27) = 56.62, p < .001$], as can be seen by comparing the three panels of Figure 2, from low to high contingencies. U values under each of the contingencies differed significantly from each of the other two ($ps < .01$); for example, high differed from both middle and low, etc. Also, when the two conditions were considered individually, U values increased significantly as the contingencies became more demanding, both for the delay condition, represented by the filled circles [$F(2,27) = 64.80, p < .001$], and for the time-out condition, represented by the open circles [$F(2,27) = 39.22, p < .001$]. Thus, the variability-reinforcing contingencies clearly controlled levels of sequence variability.

Two other results were indicated by a significant contingency $\times$ condition $\times$ duration interaction [$F(6,81) = 5.91, p < .001$]. First, variability was higher when reinforcement was delayed than when time-outs followed reinforcements, but this was true only under low (Figure 2, left panel) and middle (center panel) reinforcement contingencies [$F(1,27) = 5.24, p = .030$, and $F(1,27) = 16.69, p < .001$, respectively] and not under high [$F(1,27) = 0.49$, n.s.]. Second, both condition (delay vs. time-out) and contingency (low, middle, and high) influenced directions of change. Consider first the delay condition. With increases in delays, variability increased under low contingencies [$F(3,81) = 2.70, p = .051$] (left panel, filled circles), increased under middle contingencies, but not significantly [$F(3,81) = 1.51$, n.s.] (middle panel, filled circles), and *decreased* under high contingencies, an effect that approached significance [$F(3,81) = 2.63, p = .056$] (right panel, filled circles). Note especially the reversal between the low and high groups: The effects of delaying reinforcements depended on whether low or high levels of variability were being reinforced. In contrast, with increases in postreinforcement time-outs, variability *decreased* at all contingency levels (open circles), the effects being statistically significant for low [$F(3,81) =$

Figure 2. *U* values, a measure of sequence variability, as a function of four durations of prereinforcement delays (filled circles) or postreinforcement time-outs (open circles), respectively. Shown separately are performances for the groups of rats that were reinforced for low levels (left panel), middle levels (middle panel), and high levels (right panel) of response variability. Best-fitting least-squared lines are drawn to indicate directionality of change.

10.31, $p < .001$] and middle [$F(3,81) = 15.47$, $p < .001$] contingencies, but not high [$F(3,81) = 1.56$, n.s.].

As noted in the Method section, numbers of trials were equalized, individually for each subject, across all of the delay conditions, with this being done so as not to bias the *U*-value analyses. However, to ensure that this truncation procedure did not misrepresent the results, we repeated the *U*-value analysis on the complete data set, using all of the trials emitted in all of the sessions, and found even stronger evidence for statistical significance. For example, overall differences across the groups' *U* values yielded $F(2,27) = 70.9$, $p < .001$; this *F* ratio was larger than that reported above for the truncated data set, as were all of the other full data set *F* ratios.

Thus, (1) variability-reinforcing contingencies exerted strong control over levels of sequence variability; (2) delays and time-outs affected variability differently and in some unexpected ways (see below); and (3) as delay durations increased, variability increased when the contingencies were lenient (low) but tended to decrease when the contingencies were stringent (high).

Response Rate

Figure 3 depicts rates of responding, defined as trials per minute. The main effects were these. First, as expected, delaying reinforcement consistently resulted in lower response rates than did postreinforcement time-outs [$F(1,81) = 113.41$, $p < .001$]. This difference is indicated by filled circles (delays) falling below open circles (time-outs) at each duration. Second, as delay durations increased, rates of responding decreased [$F(3,81) = 83.62$, $p < .001$], as is shown by the steep downward sloped best-

fitting lines under each of the three contingencies (filled circles). Although postreinforcement time-out durations also affected response rates [$F(3,81) = 11.08$, $p < .001$], the effect was due mainly to an unexplained low rate at the 2-sec value, and overall response rates were relatively unchanged, with the best-fitting lines flat or upward sloping (open circles). Third, as the variability contingencies became increasingly demanding, from low to high, overall response rates decreased [$F(2,27) = 8.39$, $p = .001$], with this being true under each of the conditions individually ($ps \leq .01$). The decrease may have resulted from the time taken to traverse the chamber, which occurred increasingly with contingency demands. Thus, in summary, when a delay period preceded reinforcement, response rates decreased appreciably: the longer the delay, the slower the response rates. In comparison, when identical blackouts were presented after reinforcements, responding was not consistently affected, and responding was significantly faster than when reinforcement was delayed.

Reinforcement Rate

Because the low percentile reinforcement requirements were much more easily satisfied than the middle, and the middle more than the high, the low group was reinforced with a .3 probability for meeting the percentile requirement, and the middle group with a probability of .4. The high group was reinforced whenever its percentile requirements were satisfied. To test whether these procedures were successful in equating overall reinforcements, we calculated the ratios of reinforced trials to total trials and found that the overall average reinforcement probability (across all delays and conditions) for the high group was

Figure 3. Trials per minute, a measure of response speed, as a function of four durations of prereinforcement delays (filled circles) or postreinforcement time-outs (open circles), respectively. Shown separately are performances for the groups of rats that were reinforced for low levels (left panel), middle levels (middle panel), and high levels (right panel) of response variability. Best-fitting least-squared lines are drawn to indicate directionality of change.

.19, for the middle group .17, and for the low group, .18, a difference that was not statistically significant [$F(2,27) = 2.37, p > .10$]. Thus, we succeeded in approximately equalizing reinforcements.

Four main conclusions can be reached. First, response sequence variability was controlled by reinforcement contingencies, or more precisely, by whether low, middle, or high levels of variability were required for reinforcement. Contingency-controlled variability was found under both prereinforcement delay and postreinforcement time-out conditions and across all durations of delays and time-outs. Second, delays of reinforcement had smaller effects on variability than might be predicted from previous research on delay, and the effects depended on the variability contingencies: When low levels of variability sufficed for reinforcement, responding became increasingly variable as delay durations increased, but when reinforcement depended on highly variable outputs, opposite changes were observed (i.e., variability tended to decrease). Third, at least under low and middle contingencies, delayed reinforcement led to higher variability than did postreinforcement time-outs. Last, prereinforcement delays—but not postreinforcement time-outs—caused rates of responding to decrease. Each of these conclusions will be considered in the Discussion section, but we first describe two control experiments.

EXPERIMENT 2

In Experiment 1, all rats first experienced the time-out condition and then the delay. Experiment 2A controlled for order of experience: Using only the 8-sec delay and time-

out durations, we alternated single sessions of delay with single sessions of time-out. The main question was whether variability would differ under these conditions, and whether delays would again lead to higher variability than would time-outs when the contingencies were permissive.

Experiment 2B also controlled for order of experience but compared five consecutive sessions under each of the two conditions: The order of presentation (delay vs. time-out) was counterbalanced across animals. In addition, in this experiment, we tested whether the observed effects in Experiment 1 were transitory.

Method

Subjects, Apparatus, and Procedure. Experiment 2A was preceded by a comparison of 0.5- and 16-sec delays and time-outs, but these data are not presented because rates of responding in a number of animals decreased to very low levels under the 16-sec duration. Experiment 2B followed directly after Experiment 2A. The subjects, group allocations, and procedures in both experiments were the same as in Experiment 1, except for the following.

Experiment 2A. Sessions with 8-sec prereinforcement delays alternated with 8-sec postreinforcement time-outs, one session per condition. There were six cycles, with the last five cycles used in analyses. Sessions terminated when either 2 h had passed or 60 reinforcers had been administered.

Experiment 2B. Five consecutive sessions of 8-sec delay-of-reinforcement preceded five consecutive sessions of postreinforcement time-outs for a randomly selected half of the subjects, with the order reversed for the other half. In all other ways, the procedure was the same as in Experiments 1 and 2A. The few instances of missing data—owing to computer or chamber malfunctioning—were replaced with an average of the data from the prior and following sessions so that ANOVAs could be run. In the one instance where the data from the last session of a condition were missing, data from the fourth session of that condition were inserted.

Results

Variability. Figure 4 (top panels) shows that essentially the same effects as observed in Experiment 1 were found in Experiment 2A. Variability increased significantly as the contingency demands increased from low to high [$F(2,27) = 60.38, p < .001$]. The condition $\times$ contingency interaction was significant [$F(2,27) = 4.67, p = .018$], with higher delay than time-out U values under low and middle reinforcement contingencies [$F(1,27) = 21.04, p < .001$, and $F(1,27) = 4.38, p = .046$, respectively], but no difference under the high contingency [$F(1,27) = .02$, n.s.]. Thus, when a single session of delayed reinforcement alternated with one session of time-out, the results supported and generally replicated those from Experiment 1. As can be seen in Figure 2, clear differences in variability were observed at the 8-sec duration in Experiment 1, but only for the low and middle contingency groups, an effect replicated in the present experiment.

Essentially the same pattern of results was obtained in Experiment 2B (see Figure 4, bottom panels), where five

sessions of 8-sec delay preceded five sessions of 8-sec postreinforcement time-out for half of the animals, and the reverse for the others. Again, contingencies exerted strong and statistically significant influence [$F(2,27) = 55.39, p < .001$], with a significant interaction again between condition and contingency [$F(2,27) = 6.08, p = .007$]; that is, under low contingencies, delaying reinforcement resulted in higher variability than did postreinforcement time-out [$F(1,27) = 15.26, p < .001$], whereas the differences between these two conditions were not statistically significant under the middle- and high-variability contingencies [$F(1,27) = 1.57$ and $F(1,27) = 1.20$, respectively; both n.s.].

Response rate. Response rate effects in Experiments 2A and 2B were also essentially the same as in Experiment 1 (see Figure 5). Of most importance, response rates were higher in the postreinforcement time-out condition (open circles) than in the prereinforcement delay condition (filled circles) [$F(1,27) = 168.63, p < .001$ in Experiment 2A and $F(1,27) = 174.41, p < .001$ in Experiment 2B].

Figure 4. The top graph shows U values as a function of sessions in Experiment 2A in which delay-of-reinforcement conditions (filled circles) alternated with time-out conditions (open circles). Shown separately are data from the rats reinforced for low levels of variability (left panel), middle levels of variability (middle panel), and high levels of variability (right panel). The bottom graph shows the same for Experiment 2B when the rats received five sessions of experience under delay conditions (or time-out, counterbalanced) and then five sessions under the other.

Figure 5. Trials per minute in Experiment 2A (top) and 2B (bottom). See Figure 4 for additional explanation.

DISCUSSION

There were four main findings in the present study.

1. Delaying reinforcements caused responding to slow appreciably—the longer the delay, the lower the response rate. This was found whether low, middle, or high variability was reinforced. This result is consistent with those from many other studies (e.g., Odum et al., in press; Reilly & Lattal, 2004; Richards, 1981; Shahan & Lattal, 2005). As expected, postreinforcement time-outs had small and inconsistent effects on response rate. Response rates were also affected by the variability contingencies, but that may have been because of the particular configuration of operanda. With increasingly stringent demands for variability, the rats had to traverse the chamber increasingly, and that required more time.

2. Sequence variability was primarily controlled by reinforcement contingencies, with large differences maintained among the three groups throughout the experiment: The high group responded most variably, the low group least, and the middle group responded at an intermediate level. These results are consistent with other studies that show robust control by reinforcement contingencies over operant variability (Blough, 1966; Grunow & Neu-

ringer, 2002; Machado, 1989; Page & Neuringer, 1985). They also indicate a dissociation similar to that described by Grunow and Neuringer (2002): Response *rates* were primarily controlled by reinforcement frequencies (in Grunow & Neuringer, 2002) and delays (in the present experiment), whereas response *variability* was primarily controlled by reinforcement contingencies in both experiments. The results are consistent as well with previous studies showing that reinforced variability is relatively immune to adverse impact from other reinforcement-weakening events (Cohen et al., 1990; Doughty & Lattal, 2001; Grunow & Neuringer, 2002; McElroy & Neuringer, 1990; Neuringer, 1991).

3. Although delays produced only small changes in variability, there was a statistical interaction between delay and contingency. More precisely, in the group reinforced for high variability, as delays increased, responding tended to become less variable, but in the low group, responding became more variable. Variability in the middle group also increased with delay, but the effect did not reach statistical significance. These results parallel those from Grunow and Neuringer (2002), where reinforcement frequency interacted with contingencies to influence variability of responding.

4. Postreinforcement time-outs produced *lower levels* of variability in two of the three groups (the low and middle) than did comparable delays, an unexpected finding. Also unexpected was that the longer the postreinforcement time-outs, the lower the variability, this being true across all groups.

To summarize, many of the findings were predicted on the basis of Grunow and Neuringer (2002), where lowering reinforcement rates caused large, consistent, and significant decreases in response rates, suggesting a weakening of reinforcement control. The same was found in the present experiment when delay durations were lengthened. However, in both cases, there was little or inconsistent weakening with respect to control of variability: The reinforcement contingencies maintained a high degree of control over sequence variability despite the lowered rates of or increased delays to reinforcement. Furthermore, statistical interactions were the same in both cases: Variability decreased slightly when reinforcement was degraded in the high-variability groups but increased in the low-variability groups. These results raise basic questions. In particular, if variability is an operant, why was it not more affected by reinforcement delays and rates, and how can increases in variability—in the low groups—be explained?

One obvious possibility is that the results challenge the operant nature of variability. The argument is as follows: Since delays are known to *decrease* control by reinforcement in other cases, the observed *increases in variability* with increasing delay—at least when low or middle baseline levels of variability were reinforced—suggest that the variability must have other causes. That is, the variability was *not* generated by contingent reinforcement. The main difficulty with this interpretation is that the three reinforcement contingencies—low, middle, and high—maintained large and significant differences in variability throughout the experiment, despite the delays. Thus, sensitivity to variability-reinforcing contingencies appears to have been maintained.

Elicitation Effects

To explain the results, we hypothesize that other influences were in fact involved and, in particular, that schedule-induced, elicited, or noncontingent influences (these being used here as synonyms) worked in concert with the contingencies. Many studies have shown that intermittent reinforcers elicit activities such as drinking, gnawing, wheel-running, and pacing (Falk & Kupfer, 1998). These behaviors, often referred to as schedule-induced, are characterized as “elicited” because their production does not depend on a particular contingency relating behaviors to reinforcers. A characteristic of adjunctive responses is that they change as a function of temporal proximity to the intermittently presented reinforcer (see, e.g., Staddon & Simmelhag, 1971).

Variability may also be elicited when reinforcement is intermittent (Antonitis, 1951; Balsam et al., 1998; Boren, Moerschbaecher, & Whyte, 1978; Tatham, Wanchisen, & Hine, 1993; Tremont, 1984). From the results of such studies, Gharib, Gade, and Roberts (2004) derived a gen-

eral hypothesis—namely, “that reducing reward expectation increases variation” (p. 275). These authors provided additional convincing evidence of elicited variability in their studies of barpress hold-down durations. Their rat subjects experienced a multiple schedule in which reinforcers were more probable in one stimulus condition than in another. Barpress durations varied significantly more during the low-reinforcement stimulus than during the high.

Evidence consistent with the Gharib et al. (2004) expectancy hypothesis was provided by Cherot et al. (1996), who showed that levels of sequence variability—in rats rewarded for varying their sequences—decreased as the reinforcer was approached. Other rats showed that the decrease in variability was elicited. These rats were rewarded for sequence repetitions, and they showed the same change—decreased variations as the reinforcer was approached. Of importance for the present discussion were the large differences maintained by the two contingencies: The “vary” animals varied significantly more than did the “repeat.” Thus, in brief, reinforcement of variability maintained significantly higher levels of variability than did reinforcement of response repetitions, but, superimposed on these large group differences, variability tended to decrease in both groups as the reinforcer was approached.

The previous discussion suggests that as reinforcers are approached in time, variability tends to decrease and, conversely, repetitions to increase, with these results being independent of the reinforcement contingencies. Stated differently, when expectation or anticipation of reinforcement is low, elicited variability is high; and as expectation increases with approach, variability decreases—the Gharib et al. (2004) hypothesis. This relationship is illustrated in Figure 6, where hypothesized elicited variability (along the *y*-axis) is represented as a function of time to reinforcement (along the *x*-axis). The particular function is hyperbolic, but apart from its decreasing monotonicity, the function’s particular form is not critical for our discussion. (Killeen [1995] and Killeen & Sitomer [2003] provide much evidence for a related theory.) As can be seen in the top panel, when responding is relatively distant from a potential reinforcer, elicited variability is high; and as the reinforcer is approached, elicited variability decreases. Since the delay occurs *at the end of the interreinforcement interval*, the time available for responding is relatively early in the interval, where elicited variability is high. Shown in the bottom panel is the same function, but in this case, the postreinforcement time-out condition is represented. The time-out occurs after reinforcement and therefore *early in the interval* prior to the next reinforcement. Thus, here, responding is limited to the period late in the interval—preceding reinforcement, where elicited variability is relatively low. All else being equal, therefore, elicited variability will be higher under delay-of-reinforcement conditions than under postreinforcement time-out conditions.

Figure 6 represents 8-sec delay and time-out conditions, used as an illustrative example. Similar figures would represent the other blackout durations—both delay

Figure 6. Description of how elicited variability, along the y-axis, changes with interreinforcement time, along the x-axis, under the delay-of-reinforcement condition (top) and time-out condition (bottom). The shaded portions indicate blackouts, when responses are not possible, and the open portions indicate the periods when responses can be emitted. The drawn curve was computed from $y = \{1 - [1/(1 + x)]\}$.

and time-out—but the proportion of the interval allocated to responding would differ, depending on the blackout and response durations *relative* to the total interreinforcement interval. This is because responding within an interreinforcement interval—both contingency controlled and elicited—is a function of the relative time to reinforcement, an observation often referred to as “scalar timing” (Gibbon, 1977; Killeen & Fetterman, 1988). Delays of reinforcement similarly exert their effects relative to the interreinforcement interval; they are scalar as well (e.g., Schaal, Odum, & Shahan, 2000; Schaal, Schuh, & Branch, 1992; see, however, Shahan & Lattal, 2005, regarding possibly different effects for unsignaled delays).

Explanation of Results

1. The overall differences in variability across the three groups can be accounted for in large part by the differ-

ent contingencies: low, middle, and high. All other effects on variability were small relative to the main effects of contingency.

2. Elicitation can account for the higher levels of variability under delay conditions than under time-out. (We refer here to the low and middle contingencies. The high will be discussed below.) The reason is that, as explained above, delay periods occurred at the end of each interreinforcement interval (*before* the reinforcer was received), whereas time-out periods occurred at the beginning of each interval (*after* the previous reinforcer). Since elicited variability is higher early in an interreinforcement interval rather than late, it accounts for the higher overall variability with delays than with time-outs.

3. Elicitation also explains the directions of change in low and middle groups. The reasoning is as follows. In the delay condition, as durations increased, the *relative time* allotted to responding—in comparison with the total interreinforcement interval—*decreased*; that is, the proportion of the total trial time (including delay) in which responding was possible was greater with the 0.5-sec delay than with the 16-sec delay. Thus, the opportunity to respond when delays were short included both high and low elicited variability periods (see Figure 6), whereas with the longer delay durations, responding was mainly confined to the high-variability periods—that is, to a time period relatively early in the overall interreinforcement interval. The consequence was that variability *increased* with increasing delay durations. The same analysis also explains the opposite postreinforcement time-out effects: Increasing time-out durations caused variability to *decrease*, because responding was increasingly confined to the later portions of the interreinforcement intervals.

4. We now compare the low and middle contingency groups with the high group. First, it appears that as the contingencies became more demanding, reinforcement exerted increasing control of variability, and, consequently, the relative effects of elicitation decreased. This hypothesis was supported by the fact that levels of variability under delay and time-out conditions were similar for the high group, where the reinforcement contingencies were most demanding, but differed for the middle and low (observed in Experiments 1, 2A, and 2B, as well as in Grunow & Neuringer, 2002). A second finding is more problematic, and we do not currently have a complete explanation. That is, with increasing delay durations, variability in the high group tended to decrease, a direction *opposite* to that for the low and middle groups. The difference between high and low was also observed in the four experiments just cited. One possible explanation is that because the contingencies exerted preponderant control when high variability was required, the weakening effects of delay on contingency could be discerned. Under the more permissive middle and low contingencies, elicited effects overshadowed this weakening. However, this is an ad hoc explanation, and additional study is required.

The elicitation-plus-contingency hypothesis may also explain why delays often interfere more with repeated op-

erant responses than with variable operants (e.g., Doughty & Lattal, 2001; Odum et al., in press). According to this hypothesis, delays constrain responding to a period early in the interreinforcement interval, the result being that elicited variability will *enhance* performance under an operant variability contingency, as discussed above, but interfere with performance under an operant repetition contingency (see Figure 6). Furthermore, again as indicated above, the longer the delay, the greater the effects—facilitative for operant variability and interfering for operant repetition.

Implications

Three findings may have important implications. First, when reinforcement is based on a particular level of response variability, that level is produced and maintained. This finding, which has considerable general support, may be important for training procedures where variability is functional. Examples include situations where new operant responses are being shaped, new environments explored, difficult problems confronted, and creative activities engaged in (see Neuringer, 2002). In each of these cases, explicit reinforcement of variations may contribute to performance. Regarding creativity, some research has indicated that expectations of rewards can have detrimental effects (e.g., Amabile, 1983). To the extent that behavioral variability contributes to creativity (Campbell, 1960), such detrimental effects must be compared with those produced by the precise contingencies that are applied to creative behaviors. Although expectancies may, indeed, sometimes result in lowered creativity, the contingency-plus-elicitation hypothesis indicates that expectancy-produced effects will be relatively minor, and, in particular, when highly creative behaviors are explicitly reinforced, expectation of reinforcement may play little or no role. The hypothesis predicts, therefore, that detrimental reward effects will be small or nonexistent when reinforcers are contingent upon a high level of creativity.

Second, baseline levels of variability may influence effects of weakening operations. In particular, when high variability is maintained by reinforcement, the level of variability may be lowered somewhat when reinforcers are weakened. Quite differently, when low variability is sufficient for reinforcement, the same weakening events may lead to higher levels of variability. These differences may be relevant to the cases mentioned above—problem-solving, exploration, and creativity—where high variability may be functional, and therefore often reinforced, as well as to opposite cases where repetitive behaviors, such as in repetitive office or factory work, are maintained.

Third, variability-reinforcing contingencies appear to partly immunize variability from adverse impacts. The present study, as well as a similar experiment by Odum et al. (in press), has shown preponderant control by reinforcement contingencies and relatively small influence of reinforcement delays. Other examples have been described above. Many of these reinforcement-weakening manipulations

caused response rates to decrease and repeated operants to be degraded, but contingency-controlled variability was conserved. This may point to important differences between operant variation and operant repetition. The functionality of these differences should be explored in future research.

REFERENCES

- AMABILE, T. M. (1983). *The social psychology of creativity*. New York: Springer.
- ANTONITIS, J. J. (1951). Response variability in the white rat during conditioning, extinction, and reconditioning. *Journal of Experimental Psychology*, **42**, 273-281.
- BALSAM, P. D., DEICH, J. D., OHYAMA, T., & STOKES, P. D. (1998). Origins of new behavior. In W. O'Donohue (Ed.), *Learning and behavior therapy* (pp. 403-420). Boston: Allyn & Bacon.
- BLOUGH, D. S. (1966). The reinforcement of least frequent interresponse times. *Journal of the Experimental Analysis of Behavior*, **9**, 581-591.
- BOREN, J. J., MOERSCHBAECHER, J. M., & WHYTE, A. A. (1978). Variability of response location on fixed-ratio and fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **30**, 65-67.
- CAMPBELL, D. T. (1960). Blind variation and selective retention in creative thought as in other knowledge processes. *Psychological Review*, **67**, 380-400.
- CHEROT, C., JONES, A., & NEURINGER, A. (1996). Reinforced variability decreases with approach to reinforcers. *Journal of Experimental Psychology: Animal Behavior Processes*, **22**, 497-508.
- COHEN, L., NEURINGER, A., & RHODES, D. (1990). Effects of ethanol on reinforced variations and repetitions by rats under a multiple schedule. *Journal of the Experimental Analysis of Behavior*, **54**, 1-12.
- DENNEY, J., & NEURINGER, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior*, **26**, 154-162.
- DOUGHTY, A. H., & LATTAL, K. A. (2001). Resistance to change of operant variation and repetition. *Journal of the Experimental Analysis of Behavior*, **76**, 195-215.
- FALK, J. L., & KUPFER, A. S. (1998). Adjunctive behavior: Application to the analysis and treatment of behavior problems. In W. O'Donohue (Ed.), *Learning and behavior therapy* (pp. 337-351). Boston: Allyn & Bacon.
- GALBICKA, G. (1994). Shaping in the 21st century: Moving percentile schedules into applied settings. *Journal of Applied Behavior Analysis*, **27**, 739-760.
- GHARIB, A., GADE, C., & ROBERTS, S. (2004). Control of variation by reward probability. *Journal of Experimental Psychology: Animal Behavior Processes*, **30**, 271-282.
- GIBBON, J. (1977). Scalar expectancy theory and Weber's law in animal timing. *Psychological Review*, **84**, 279-335.
- GOETZ, E. M., & BAER, D. M. (1973). Social control of form diversity and emergence of new forms in children's blockbuilding. *Journal of Applied Behavior Analysis*, **6**, 209-217.
- GRUNOW, A., & NEURINGER, A. (2002). Learning to vary and varying to learn. *Psychonomic Bulletin & Review*, **9**, 250-258.
- JENSEN, G., MILLER, C., & NEURINGER, A. (2006). Truly random operant responding: Results and reasons. In E. A. Wasserman & T. R. Zentall (Eds.), *Comparative cognition: Experimental explorations of animal intelligence* (pp. 459-480). New York: Oxford University Press.
- KILLEEN, P. R. (1995). Economics, ecologies, and mechanics: The dynamics of responding under conditions of varying motivation. *Journal of the Experimental Analysis of Behavior*, **64**, 405-431.
- KILLEEN, P. R., & FETTERMAN, G. J. (1988). A behavioral theory of timing. *Psychological Review*, **95**, 274-295.
- KILLEEN, P. R., & SITOMER, M. T. (2003). MPR. *Behavioural Processes*, **62**, 49-64.
- MACHADO, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, **52**, 155-166.
- MCLEROY, E., & NEURINGER, A. (1990). Effects of alcohol on reinforced

- repetitions and reinforced variations in rats. *Psychopharmacology*, **102**, 49-55.
- MILLER, N., & NEURINGER, A. (2000). Reinforcing variability in adolescents with autism. *Journal of Applied Behavior Analysis*, **33**, 151-165.
- NEURINGER, A. (1991). Operant variability and repetition as functions of interresponse time. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 3-12.
- NEURINGER, A. (2002). Operant variability: Evidence, functions, and theory. *Psychonomic Bulletin & Review*, **9**, 672-705.
- NEURINGER, A., DEISS, C., & OLSON, G. (2000). Reinforced variability and operant learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 98-111.
- NEURINGER, A., KORNELL, N., & OLUFS, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 79-94.
- ODUM, A. L., WARD, R. D., BARNES, C. A., & BURKE, K. A. (in press). The effects of delayed reinforcement on variability and repetition of response sequences. *Journal of the Experimental Analysis of Behavior*.
- PAGE, S., & NEURINGER, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, **11**, 429-452.
- PRYOR, K. W., HAAG, R., & O'REILLY, J. (1969). The creative porpoise: Training for novel behavior. *Journal of the Experimental Analysis of Behavior*, **12**, 653-661.
- REILLY, M. P., & LATTAL, K. A. (2004). Within-session delay-of-reinforcement gradients. *Journal of the Experimental Analysis of Behavior*, **82**, 21-35.
- RICHARDS, R. (1981). A comparison of signaled and unsignaled delay of reinforcement. *Journal of the Experimental Analysis of Behavior*, **35**, 145-152.
- SCHAAL, D. W., ODUM, A. L., & SHAHAN, T. A. (2000). Pigeons may not remember the stimuli that reinforced their recent behavior. *Journal of the Experimental Analysis of Behavior*, **73**, 125-139.
- SCHAAL, D. W., SCHUH, K. J., & BRANCH, M. N. (1992). Key pecking of pigeons under variable-interval schedules of briefly signaled delayed reinforcement: Effects of variable-interval value. *Journal of the Experimental Analysis of Behavior*, **58**, 277-286.
- SHAHAN, T. A., & LATTAL, K. A. (2005). Unsignaled delay of reinforcement, relative time, and resistance to change. *Journal of the Experimental Analysis of Behavior*, **83**, 201-219.
- STADDON, J. E. R., & SIMMELHAG, V. L. (1971). The "superstition" experiment: A reexamination of its implications for the principles of adaptive behavior. *Psychological Review*, **78**, 3-43.
- STOKES, P. D. (1999). Learned variability levels: Implications for creativity. *Creativity Research Journal*, **1**, 37-45.
- STOKES, P. D., & HARRISON, H. M. (2002). Constraints have different concurrent effects and aftereffects on variability. *Journal of Experimental Psychology: General*, **131**, 552-566.
- TATHAM, T., WANCHISEN, B., & HINELINE, P. (1993). Effects of fixed and variable ratios on human behavioral variability. *Journal of the Experimental Analysis of Behavior*, **59**, 349-359.
- TREMONT, P. J. (1984). Variability of force and interresponse time under random interval reinforcement schedules. *Behavioural Processes*, **9**, 413-420.

(Manuscript received July 29, 2005;
revision accepted for publication October 28, 2005.)