BRIEF REPORTS

Learning to vary and varying to learn

ALICIA GRUNOW and ALLEN NEURINGER
Reed College, Portland, Oregon

We compared two sources of behavior variability: decreased levels of reinforcement and reinforcement contingent on variability itself. In Experiment 1, four groups of rats were reinforced for different levels of response-sequence variability: one group was reinforced for low variability, two groups were reinforced for intermediate levels, and one group was reinforced for very high variability. All of the groups experienced three different reinforcement frequencies for meeting their respective variability contingencies. Results showed that reinforcement contingencies controlled response variability more than did reinforcement frequencies. Experiment 2 showed that only those animals concurrently reinforced for high variability acquired a difficult-to-learn sequence; animals reinforced for low variability learned little or not at all. Variability was therefore controlled mainly by reinforcement contingencies, and learning increased as a function of levels of baseline variability. Knowledge of these relationships may be helpful to those who attempt to condition operant responses.

Variability is thought to be necessary for the shaping of operant responses (Baum, 1994; Glen, Ellis, & Greenspoon, 1992; Palmer & Donahoe, 1992; Skinner, 1984). For example, "Variability in behavior provides the means by which a totally new behavior, never performed by an individual before, can gradually be developed" (Mazur, 1998, p. 131). An important question for those who attempt to modify behaviors in the world outside of the laboratory concerns sources of such variability. More precisely, How can one generate the variability necessary for shaping to succeed?

There is inherent variability in all behavior, but it does not always suffice to generate reinforceable instances, and, moreover, a trainer of new behavior does not have control over this "endogenous" source.

An alternative is for trainers to withhold reinforcement; many studies have documented increased variability in extinction (Antonitis, 1951; Balsam, Deich, Ohyama, & Stokes, 1998; Neuringer, Kornell, & Olufs, 2001). However, responding decreases during extinction, sometimes to zero levels, and previously reinforced responses often persist. During shaping, it is necessary that responding be both variable and maintained, and, therefore, reinforcement must be provided, at least occasionally. It is more relevant to ask, What happens when reinforcement is temporarily withheld or when reinforcement frequency is lowered, but not

to zero levels? The evidence here is mixed, some studies reporting increased variability when reinforcement frequencies decrease (Boren, Moerschbaecher, & Whyte, 1978; Tatham, Wanchisen, & Hineline, 1993; Tremont, 1984), but others showing small or no effects (Eckerman & Lanson, 1969; Herrnstein, 1961; Machado, 1989).

A third source of the requisite response variations might be reinforcement contingent on variability itself. Although direct reinforcement of variability has only rarely been discussed as a contributor to operant shaping (see Neuringer, 1993; Stokes & Balsam, 1991), there is substantial evidence showing that variability can be reinforced (Blough, 1966; Bryant & Church, 1974; Machado, 1989; Page & Neuringer, 1985; Pryor, Haag, & O'Reilly, 1969).

A question of importance is, Which of these putative sources of variability contributes most to the variability required to train new operants? Two studies compared the variability generated by different reinforcement frequencies with that generated by variability-reinforcing contingencies (Blough, 1966; Machado, 1989). Different reinforcement frequencies had little effect in one case (Blough, 1966) and inconsistent effects in the other (Machado, 1989). In marked contrast, both studies found that the contingencies exerted strong and consistent influence: As the contingencies required more response variability, levels of variability increased. Thus, direct reinforcement of variability could contribute more to the variations necessary for the conditioning of new operant responses than the intermittency of reinforcement. This is an important claim for those who would apply operant conditioning principles, and its generality therefore warrants testing.

In Experiment 1, we employed contingencies, procedures, and parameters not previously studied. One group of rats was reinforced with food only if response sequences were

This research was supported in part by grants from the National Science Foundation to A.N. Data were reported by A.G. in a Reed College undergraduate thesis. The authors thank Nate Kornell for his critically important technical assistance and theoretical contributions to this research. The authors also thank Gene Olson for his excellent animal care. Correspondence should be addressed to A. Neuringer, Psychology Department, Reed College, Portland, OR 97202 (e-mail: allen.neuringer@reed.edu).

highly variable. The contingencies were parametrically less demanding for three other groups, with the last requiring minimal variation. All groups were first trained under continual reinforcement (CRF) conditions, in which every success in meeting the contingency was reinforced, following which reinforcement frequency was decreased to a variable interval of 1 min (VI 1), in which meeting the respective contingencies was reinforced on the average of no more than once per minute, and then to a variable interval of 5 min (VI 5), with a final return to VI 1. The goal was to test whether reinforcement contingency in fact exerts greater control over response variability than does reinforcement frequency.

Since we found greater control by contingencies than by reinforcement frequencies in Experiment 1, we further studied the effects of parametric differences in variability contingencies on the acquisition of a particular sequence of responses in Experiment 2. The main question was, How do different levels of baseline variability affect learning?

EXPERIMENT 1

Method

Subjects

Forty male Long-Evans rats, 4 months old at the start, were housed in pairs, with free access to water and a 12:12-h light:dark cycle. The rats had previously been trained to press levers for food. Experimental sessions were provided 5 or 6 days a week, with free access to food for 1.5 h after each session and for the same duration during nonexperimental days.

Apparatus

Ten modified Gerbands operant chambers were used (details in Neuringer et al., 2001). On the front wall, two levers, each 5.1 cm wide, 5 cm above the floor, with 8.9 cm between them, were designated *right* (R) and *left* (L), with a 28-V DC light above each. A pellet tray between the two levers provided access to 45-mg Noyes pellets. Three pigeon keys, each 2 cm in diameter and 9 cm above the floor, were located on the rear wall; the middle key (K) illuminated with 28-V DC lights (the other two keys not being used in this experiment). A tube protruded 10 cm below the center of the chamber ceiling, providing continuous access to water, and an overhead 28-V bulb served as the houselight.

Procedure

A trial consisted of three responses across the three operanda. Trials ended with reinforcement if a variability contingency had been met, and there was a brief timeout otherwise. A different level of sequence variability was required to meet the contingencies for each of the four groups. In the first phase, reinforcement was provided each time the variability contingencies were satisfied. In later phases, reinforcement frequency was decreased.

Variability contingencies. With three operanda (two levers and one key) and three responses per trial, 27 different sequences were possible (e.g., KKK, RKL, LLR). A sequence was reinforced only if its weighted relative frequency was less than a threshold value. The relative frequency of a sequence was calculated by dividing the number of times the sequence had been emitted by total trials. All sequence frequencies were multiplied by a weighting coefficient of .98, applied after each reinforcement (for details, see Denney & Neuringer, 1998). This weighting caused recent sequences to contribute more to the assessment of relative frequencies than did sequences emitted earlier. Data for these calculations were carried over from one session to the next.

The weighted relative frequency of the emitted sequence was compared with a threshold to determine whether the given sequence had occurred sufficiently infrequently to be reinforced. One group of animals experienced a very demanding threshold, .037, where a given sequence was reinforced only if the current value of its weighted relative frequency was less than .037 (i.e., 1/27). Threshold values assigned to the three other groups were .055, .074, and .37, the last being very permissive. A simulating random generator, emitting each of three responses with equal probability, satisfied the same contingencies (i.e., was "reinforced") on 53% of trials under .037 contingencies, 84% under .055 contingencies, 96% under .074, and 100% under .37 contingencies.

Reinforcement frequency. In Phase 1, every trial that met the variability contingency was reinforced. This condition is referred to as *CRF*. In Phase 2, a VI 1 min (intervals between 10 sec and 3 min) schedule limited reinforcement to no more than once per minute. When a programmed interval had elapsed, reinforcement was delivered following the next sequence that met the variability contingency. After the reinforcer, a new interval was randomly chosen, and timing began again. Trials that met the variability contingency but occurred before the interval had elapsed resulted in the same timeout (see below) as trials that did not meet the contingency. Phase 3 provided a VI 5 min schedule with the same intervals as under VI 1 multiplied by 5. Each phase consisted of ten 45-min sessions.

Procedural details. A trial began with illumination of the lights above the levers and the keylight. Responding on any of the three operanda resulted in darkening of the corresponding light for 0.33 sec, during which responses reset the interval. (Responses during this interval did not count toward sequence completion.) Each trial, consisting of three effective responses, terminated with reinforcement or timeout. Trials that met the variability contingency appropriate for the group and, later, the VI requirement as well resulted in a 2000-Hz tone of 0.5-sec duration, followed by a 0.1-sec 3100-Hz tone, with the delivery of one Noyes food pellet coinciding with the 3100-Hz tone. A 1-sec intertrial period, during which the lever and keylights were dark but the houselight remained illuminated, followed the delivery of the food pellet and was reset if a response occurred. If the sequence failed to meet the variability contingency or occurred before the VI interval elapsed, the chamber was darkened for 1 sec, with responses during this timeout resetting the 1-sec period.

Preliminary training. Since the animals had previously learned to press levers, only the key response was shaped, following which a "light chase" procedure was instituted for one 30-min session. At first, only one of the lights above the levers or behind the key was randomly selected and illuminated. Responding on the illuminated operandum resulted in reinforcement, while all other responses were ignored. After 50 reinforcements, trial length was increased to two (first one operandum was illuminated, to which a response caused illumination of one of the other two operanda, randomly selected) with reinforcement following two correct responses. After another 50 reinforcements, three correct responses were required for reinforcement. One baseline session followed, during which all subjects were reinforced for varying their sequences, just as in the experimental phases to follow, but with the threshold equal to the most permissive .37 for all groups.

Conditions. After the baseline session, the subjects were divided into four experimental groups by a quasi-random process. Using the baseline data, the subjects were ranked in order of the percentage of trials containing keypresses (they were least frequent). The highest ranked 4 subjects were randomly assigned, 1 to each of the four groups. The same process was repeated for the next group of 4, and so on. Levels of sequence variability during baseline were also evaluated, as well as response rates and experimental chambers; to equalize the groups, minor changes were made to group assignments. The four threshold contingencies were then randomly assigned to the groups. These contingencies were constant throughout the experiment (i.e., one group was reinforced for meeting the .037

variability contingency, the others were reinforced for meeting the .055, .074 and .37 contingencies, respectively). Following the VI 5 condition and after 1 month with no experimentation, the subjects were returned to VI 1 for five sessions, in order to see whether performances would return to previous VI 1 levels.

Measures

Sequence variability was evaluated with the U -value statistic, computed as follows:

$$\frac{-\sum_{i=1}^n [RF_i \times \{\log(RF_i) / \log(2)\}]}{\log(n) / \log(2)},$$

where n equals the number of sequences possible (in this case, 27) and RF refers to the unweighted relative frequency of each of the n sequences (see Denney & Neuringer, 1998, for details). U -value measures the distribution of sequence frequencies, with equal frequencies (indicating high variability) yielding a high U -value (1.0 being the highest) and with unequal frequencies (indicating low variability) generating a low U -value (0.0 being the lowest). U -values were computed for each session, and all statistical analyses were based on arithmetic averages across the last three sessions of each phase. Because the number of reinforcers obtained during a session was sometimes high (especially during the CRF phase), to avoid satiation effects, U -values were calculated on the basis of trials prior to the 150th reinforcement in a session or before 45 min had elapsed, whichever occurred first. This same procedure was followed for all data and analyses.

Results and Discussion

Figure 1 (left), shows that behavioral variability was precisely controlled by the variability contingencies [$F(3,36) =$

206.435, $p < .0001$], as indicated by the four functions, but only weakly and inconsistently influenced by reinforcement frequency (not significant), as shown by changes in the functions along the x-axis. Shown are average U -values during the last three sessions of CRF, VI 1, and VI 5 phases, for each of the four groups. A significant contingency $\times$ frequency interaction [$F(6,72) = 8.610$, $p < .0001$] is explained by the following: For the low-variability group (.37), variability increased as reinforcement frequency decreased [$F(2,72) = 11.385$, $p < .0001$], a result consistent with some previous reports; for the high-variability groups (.037 and .055), variability decreased together with reinforcement frequency [$F(2,72) = 11.272$, $p < .0001$, and $F(2,72) = 3.583$, $p = .033$, respectively], a result not previously reported; for the intermediate group (.074), variability did not change with reinforcement frequency. Thus, variability contingencies exerted primary control over levels of behavioral variability, and effects of reinforcement frequency were small and depended on baseline variability levels.¹

Session-by-session changes in U -values are shown in Figure 2. During the initial 10 CRF sessions, U -values rapidly diverged, the variability contingencies generating large effects. When reinforcement frequency decreased to VI 1, U -values increased in only the most permissive group (.37), and most of that effect was transient. When reinforcement again decreased to VI 5, the .37 group showed the same transient increase, with the high-variability .037 and .055 groups now showing small decreases. Upon return to VI 1, U -values tended to return

Figure 1. (Left) U -value, an index of behavioral variability, as a function of reinforcement frequencies (CRF, reinforcement every time variability contingencies were met; VI 1, reinforcement for meeting variability contingencies no more than once per minute, on average; VI 5, reinforcement no more than once every 5 min). Each line represents a different group: .037 = very high variability required for reinforcement; .37 = very low variability required; and the other two groups, .055 and .074 = intermediate levels required. (Right) Trials per minute, an index of response rate, as a function of reinforcement frequencies for each of the four groups described above.

Figure 2. *U*-values during each session of the experiment for each of the four groups. B indicates initial baseline. See Figure 1.

toward the original VI 1 levels, indicating that the results were due to reinforcement rather than an artifact of order of experience. These transition effects therefore support the conclusions based on across-session averages presented above.

Figure 3 provides further support for these conclusions. Shown are the relative frequencies of each of the 27 sequence patterns (frequencies of each sequence divided by total sequences per session), averaged over the last three sessions. The ordering is from most frequent on the left to least on the right. As the contingencies became more demanding across groups (from top to bottom panels), the distributions broadened and flattened, indicating increasingly equal emission of sequence patterns and correspondingly increased variability. The three curves in each graph represent the three reinforcement frequencies, and it is generally difficult to distinguish them. Decreasing reinforcement had greatest effects on the .37 group (Figure 3, top), with probabilities of the most favored LLL and RRR sequences decreasing. For the .037 group (Figure 3, bottom), the change to VI 5 increased the KKK sequence, a result previously observed when responding was extinguished (Neuringer et al., 2001).

Rates of responding showed an altogether different pattern (Figure 1, right). Trials per minute (the measure of response rate) decreased significantly as reinforcement frequency decreased [$F(2,72) = 239.647, p < .0001$], with no significant differences among the four contingency groups. Thus, response variability depended mainly on variability contingency and only little on reinforcement

frequency, but response rate depended exclusively on reinforcement frequency and not at all on variability contingency, a clear double dissociation. Stated differently, withholding reinforcement caused response rates to fall but did not consistently increase response variability.

As expected, the different contingencies influenced how frequently the subjects actually were reinforced under CRF conditions. Figure 4 shows obtained reinforcers per trial (reinforcement probabilities) as a function of the contingencies and of the programmed reinforcement frequencies. Under CRF, each of the groups differed significantly from the others ($p < .05$), except for the .074 and .055 groups. Under VI 1 and VI 5 conditions, however, the groups were reinforced with approximately equal probabilities (.12 to .15 under VI 1; and .05 to .06 under VI 5). Under VI 1, only the .37 and .037 groups differed statistically; under VI 5, there were no significant differences. Thus, it is unlikely that differences in obtained reinforcers can explain the observed control by variability contingencies.

EXPERIMENT 2

Given the results of Experiment 1, the next goal was to describe the relationship between variability and learning. Two previous studies showed that reinforcement of variability facilitated acquisition of difficult-to-learn response sequences (Neuringer, 1993; Neuringer, Deiss, & Olson, 2000). (Easy sequences were readily learned whether or not variability was reinforced.) In these studies, sequence

Figure 3. Relative frequencies (or probabilities) of each of the 27 possible response sequences under CRF (squares), VI 1 (triangles), and VI 5 (circles) phases. The separate graphs represent performances by each of the groups, from lowest required variability on top (.37 group) to greatest variability on bottom (.037).

variability was intermittently reinforced, whereas a particular sequence of responses (referred to as the *target sequence*) was concurrently reinforced whenever it occurred. The goal of these studies was to find out whether reinforcement of variable response sequences facilitated acquisition of a particular target sequence. The subjects learned to emit the target response sequence faster than did control subjects who were never reinforced for variations. Little is known, however, about the parameters im-

portant to the observed facilitation. It is not known, for example, whether the levels of baseline variability affect speed of acquisition. The present experiment employed the same variability contingencies used in Experiment 1. Reinforcement was provided intermittently (VI 1 min) for variations and, concurrently, whenever a preselected target sequence was emitted. The main question was whether levels of reinforced variability would affect the learning of operant sequences.

Figure 4. Reinforcement obtained per session divided by trials per session (or reinforcement probability) as a function of programmed reinforcement schedules for the four contingency groups.

Method

Subjects and Apparatus

The subjects and apparatus were the same as those in Experiment 1.

Procedure

The subjects remained in the same four contingency groups as in Experiment 1. After baseline variability was reestablished under a VI 1 schedule (five sessions), an "easy" target sequence, LRL, was concurrently reinforced (10 sessions): Whenever a trial contained the LRL sequence, the target, three pellets were provided. The LRL target was one that the subjects emitted frequently when variability contingencies were most permissive. Emission of the target sequence caused the VI timer for the variability reinforcement to reset (i.e., to restart timing of the current interval). This decreased the likelihood of a reinforcement for variability immediately after a target. After a period of remedial training to reestablish the original baseline levels of variability,² followed by 5 additional sessions of VI 1 variability training, a difficult target sequence, LKK, was concurrently reinforced in the same manner as just described (10 sessions). LKK was rarely emitted during baseline except under the most demanding variability contingencies (.037).

The same procedures were used as in Experiment 1, the only difference being the addition of concurrent reinforcement for the sequence of responses designated as the target. The target sequence resulted in a sequence of tones (a 0.1-sec 3100-Hz tone, followed by a 0.4-sec 2000-Hz tone, then a 0.1-sec 3100-Hz tone repeated three times), with a food pellet delivery coinciding with the second of the 3100-Hz tones, resulting in the delivery of three pellets whenever the target was emitted. Reinforcement for meeting the variability contingency was the same as before (one pellet).

Results and Discussion

The easy LRL target sequence was learned by all contingency groups—that is, LRL frequencies increased significantly across sessions [$F(10,360) = 106.25, p < .0001$] (Figure 5, top). The group $\times$ session interaction was significant [$F(30,360) = 1.922, p < .005$], with LRL frequen-

cies increasing more rapidly (using the difference between the last session and baseline as the measure of increase) in the high-variability group (.037) than in low-variability groups (.055 and .37) ($p < .05$). (Note that baseline levels of LRL sequences were lowest in the .037 group and highest in the .37 group.) However, all groups achieved the same high levels of target emissions during the last three sessions [$F(3,36) = 1.866, n.s.$]. Thus, for the easy LRL sequence, the high-variability group learned fastest, but all groups learned to the same level of proficiency.

Differences in learning were amplified for the difficult LKK target sequence. Speed of acquisition clearly depended on the variability contingencies, shown by a highly significant group $\times$ session interaction [$F(30,360) = 8.394, p < .0001$] (Figure 5, bottom). Again using the difference between the last session and baseline as a measure of learning speed, each group differed from every other except .074 and .37 ($p = .21$) and .037 and .055, which approached significance ($p = .06$). In general, the more variability required by the contingencies, the faster the learning of the difficult target. Group contingencies also affected the level of target performance during the last three sessions [$F(3,36) = 15.410, p < .0001$], with all of the individual comparisons significant except for .074 versus .37. This is the first study to show the importance of the variability-contingency parameter in governing speed of acquisition of operants. Learning of a difficult sequence was maximized by concurrent reinforcement of high baseline variability.

Differences in learning were further demonstrated by the distributions of relative frequencies across the final three sessions of training (Figure 6). When LRL was the target (Figure 6, top), it was emitted most frequently by all groups, with closely related sequences LLL, RLL, LLR, and RLR also occurring more frequently than others. When LKK was the target (Figure 6, bottom), it, together with KKK, a close approximation, increased in both of the high-variability groups, .037 and .055. In marked contrast, there was only small indication of learning in the intermediate-variability .074 group, with KKK most frequent, followed by LRL and LKK, the target. No learning occurred for the lowest-variability .37 group, which continued to emit the previously reinforced LRL most frequently and only rarely emitted the LKK target sequences. Thus, reinforcing high variability facilitated acquisition of a difficult target sequence, and reinforcing low variability resulted in continuation of a previously learned, but no longer functional, response. Levels of variability were critically important in determining the success or failure of conditioning.

GENERAL DISCUSSION

Blough (1966) and Machado (1989) showed that levels of behavioral variability were precisely controlled by the reinforcement contingencies but that rates of reinforcement exerted only minor or no control. In Experiment 1, we tested the generality of these findings with rats rather than pigeons, contingencies different from those used pre-

Figure 5. The top graph shows relative frequencies (or probabilities) of an easy target sequence (LRL) as a function of sessions for each of the four variability groups. The bottom graph shows the same for a difficult target (LKK).

viously, different and more complex response sets, and a wider range of parameter values. The results confirmed the Blough and Machado findings and extended them in important ways. Contingencies controlled sequence variability: Level of emitted variability was a direct function of the level required by the reinforcement contingencies. Frequency of reinforcement exerted only small effects, the direction of these effects depending in part on the contingencies.

Some previous studies reported that behavioral variability increases as reinforcement frequencies decrease (Boren et al., 1978; Tatham et al., 1993; Tremont, 1984). We found this to be the case only when initial levels of vari-

ability were low, as characteristic of previous studies. The opposite was observed when initial variability was high, variability significantly *decreasing* together with reinforcement rates. These differing effects may be important to those who apply operant conditioning principles. Withholding reinforcement will not necessarily result in the variations necessary to shape a new response. Furthermore, withholding or decreasing reinforcement results in slowed and less motivated responding, this also working against effective shaping of new behaviors. The take-home message is that direct reinforcement of variations may be the most effective method to produce the variability necessary to condition new operant responses.

Figure 6. Relative frequencies (or probabilities) of each of the 27 possible response sequences when an easy target was reinforced (LRL, on top) and a difficult one (LKK, on bottom). Each line represents a different variability group.

Experiment 2 supported this hypothesis by showing that reinforced variability facilitated learning of difficult-to-learn operant response sequences (Neuringer, 1993; Neuringer et al., 2000). As in Experiment 1, four different variability contingencies were employed, requiring high, intermediate, and low variability, but now, in addition, a particular sequence was identified as the to-be-learned target, and it was reinforced whenever it occurred. All subjects, whether their variations were high or low, learned

an easy target sequence, with the high-variability group (.037) learning fastest. When a difficult sequence served as the target, learning was directly correlated with variability: The higher the variability, the faster the learning. Indeed, the group that experienced the most permissive contingencies (and therefore varied least) continued to emit the previously reinforced, but no longer functional, sequence (i.e., the low-variability subjects never learned the difficult target). These results indicate that reinforced

variability helps subjects both to learn new response sequences and not to emit previously established, but no longer functional, sequences. The results are also consistent with other findings that learning is correlated with periods of high behavioral variability—for example, when learning motor skills (Manoel & Connolly, 1995) and cognitive strategies (Siegler, 1996).

Why does reinforcement of high variability facilitate operant learning? One possibility is that high variability expands the set from which instances are emitted: The high-variability groups were more likely to emit the difficult target and, therefore, to be reinforced for it. Reinforcement of variability may generate reinforceable instances and, for that reason alone, may promote learning. Another possibility was suggested by the finding that, despite starting at a disadvantage, the high-variability group acquired the easy target sequence at a faster rate than did the other groups. Reinforcement of variability may increase sensitivity to environmental change, at least within the domain of the variations. In other words, reinforcement of variability may have a general facilitative effect on learning. Further study is necessary to test this conjecture.

In either case, explicit reinforcement of behavioral variations should be considered by those attempting to condition operant responses under real-world conditions. Combining reinforcement of variations with reinforcement of approximations to a specified goal response may be an effective way to shape new responses.

REFERENCES

- 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). Needham Heights, MA: Allyn & Bacon.
- BAUM, W. M. (1994). *Understanding behaviorism: Science, behavior and culture*. New York: Harper Collins.
- 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.
- BRYANT, D., & CHURCH, R. M. (1974). The determinants of random choice. *Animal Learning & Behavior*, **2**, 245-248.
- DENNEY, J., & NEURINGER, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior*, **26**, 154-162.
- ECKERMAN, D. A., & LANSON, R. N. (1969). Variability of response location for pigeons responding under continuous reinforcement, intermittent reinforcement and extinction. *Journal of the Experimental Analysis of Behavior*, **12**, 73-80.
- GLEN, S. S., ELLIS, J., & GREENSPOON, J. (1992). On the revolutionary nature of the operant as a unit of behavioral selection. *American Psychologist*, **47**, 1329-1336.
- HERRNSTEIN, R. J. (1961). Stereotypy and intermittent reinforcement. *Science*, **133**, 2067-2069.
- MACHADO, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, **52**, 155-166.
- MANOEL, E. J., & CONNOLLY, K. J. (1995). Variability and the development of skilled actions. *International Journal of Psychophysiology*, **19**, 129-147.
- MAZUR, J. E. (1998). *Learning and behavior* (4th ed.). Upper Saddle River, NJ: Prentice Hall.
- NEURINGER, A. (1993). Reinforced variation and selection. *Animal Learning & Behavior*, **21**, 83-91.
- NEURINGER, A., DEISS, C., & OLSON, G. (2000). Reinforced variability and operant learning. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 98-111.
- NEURINGER, A., KORNEILL, N., & OLUF, M. (2001). Stability and variability in extinction. *Journal of Experimental Psychology: Animal Behavior Processes*, **27**, 79-94.
- PAGE, S., & NEURINGER, A. (1985). Variability is an operant. *Journal of Experimental Psychology: Animal Behavior Processes*, **11**, 429-452.
- PALMER, D. C., & DONAHOE, J. W. (1992). Essentialism and selectionism in cognitive science and behavior analysis. *American Psychologist*, **47**, 1344-1358.
- 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.
- SIEGLER, R. S. (1996). *Emerging minds: The process of change in children's thinking*. New York: Oxford University Press.
- SKINNER, B. F. (1984). Selection by consequences. *Behavioral & Brain Sciences*, **7**, 477-510.
- STOKES, P. D., & BALSAM, P. D. (1991). Effects of reinforcing preselected approximations on the topography of the rat's bar press. *Journal of the Experimental Analysis of Behavior*, **55**, 213-231.
- 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.

NOTES

1. To further test these findings, we computed *U*-values based on exactly the same number of sequences, a total of 150, for each subject in each condition. (The analysis given in the text was based on equal reinforcements, but the number of sequences necessary to obtain reinforcement varied. *U*-value may be sensitive to number of instances.) The same pattern of results was obtained with contingency clearly influencing variability [$F(3,36) = 189.550, p < .0001$] and other effects, as reported in the text.

2. Since learning of the LRL sequence in Phase 1 had greatly decreased key responses for all groups, the light chase program used in Experiment 1 was reinstated for one session, following which all subjects were exposed to their original contingencies on a CRF schedule for six sessions and, as indicated in the text, an additional five sessions of VI 1 variability training, as in the preliminary phase.

(Manuscript received July 14, 2000;
revision accepted for publication August 15, 2001.)