

Response stability and variability induced in humans by different feedback contingencies

J. H. R. MAES

University of Nijmegen, Nijmegen, The Netherlands

In two experiments, the behavioral effects of different response–feedback contingencies were examined with a task requiring human subjects to repeatedly type three-key sequences on a computer keyboard. In Experiment 1, the subjects first received positive feedback for response variability, followed by no feedback, or vice versa. In Experiment 2, the subjects first received positive feedback for response variability, followed by response-independent positive feedback, or vice versa. Response stability and variability were examined using different measures, such as percentage of trials meeting the variability criteria, frequency of use of the different response alternatives, and autocorrelations as an index of response randomness. The subjects' behavior in the first phase in each condition came to reflect the current feedback contingency. Depending on the measure examined, responding after each contingency change was characterized by both response stability and decreases or increases in response variability. The collective results are discussed in the framework of previous animal and human studies on behavioral stability and variability.

In a fairly large number of studies, researchers have examined the effects of extinction on conditioned responding in classical and operant conditioning procedures. Neuringer, Kornell, and Olufs (2001) found evidence in this literature, which will not be reviewed here again, of two opposing effects: response stability and response variability. An example of each of these effects can be found in the experiments reported by Neuringer et al. (2001). In their first experiment, rats were food-reinforced for emitting variable sequences of responses, using one, two, or three of three different operanda: two levers (L and R) and a key (K). Each response sequence that could be reinforced consisted of three responses. There are a total of 27 possible three-response sequences involving one, two, or three operand(a) (e.g., LLL, LLR, or LRK). Reinforcement was given for a sequence only if its relative frequency was less than or equal to .05. The purpose of this reinforcement schedule was to induce behavioral variability, and the question of major interest was whether, and how, response variability would change in a subsequent phase without reinforcement (extinction). The authors found evidence of behavioral stability, in that the ordering of the probabilities of the 27 possible sequences during extinction did not differ much from the ordering observed during reinforcement. The most preferred sequences during reinforcement were still preferred the most during extinction; the same was

true for the least preferred sequences. However, interestingly, at the same time, the relative frequencies, expressed as a ratio that was computed by dividing the probability during extinction by the probability during reinforcement, of the least preferred sequences increased. This is indicative of increased response variability. This mix of stability and variability was also obtained in two further experiments in which extinction was preceded either by response-independent noncontingent reinforcement or by consistent reinforcement of only one single sequence.

Neuringer et al. (2001) argued that the mix of stability and variability makes sense from an evolutionary perspective. On the one hand, stability implies continuation of behavior that was successful in the past; on the other hand, variability ensures the occurrence of (relatively) new behavior that, potentially, may be more successful than previous behavior.

The aim of the present experiments was to evaluate the balance between response stability and variability in humans, rather than in rats, and to examine more than one type of contingency change. As will be outlined below, the type of contingency transition may be expected to be important with respect to the extent of behavioral stability and variability. More generally, data on the conditions that are or are not conducive to the occurrence of behavioral stability and variability (adaptive or not) may be relevant to applied settings, such as educational and therapeutic contexts.

In the present experiments, I used a task that was modeled after that employed by Neuringer et al. (2001). Specifically, subjects were asked to repeatedly press a sequence of three keys on a computer keyboard. There were two conditions in the first experiment. In the first phase of the first condition, the subjects received positive feedback for a just-emitted sequence if its relative frequency did not ex-

I thank Esther Aarts and Olga Meulenbroek for acquiring the data in Experiments 1 and 2. I further thank two anonymous reviewers for their valuable comments on earlier drafts of this paper. Correspondence concerning this article should be addressed to J. H. R. Maes, Nijmegen Institute for Cognition and Information, Biological Psychology, University of Nijmegen, P. O. Box 9104, 6500 HE Nijmegen, The Netherlands (e-mail: maes@nici.kun.nl).

ceed variability criteria similar to those adopted by Neuringer et al. (2001). In the second phase, feedback was no longer presented (extinction). The reverse order of contingencies was in effect in the second condition: No feedback was followed by positive feedback for response variability. The second experiment was identical to the first, except that the extinction phase in each condition was replaced by a noncontingent feedback phase, in which positive feedback was provided independently of the subject's response. The experiments permitted both between- and within-group comparisons.

In principle, the variability criteria used in the experiments of Neuringer et al. (2001) and those used in the present experiments can be satisfied using different strategies (see also, e.g., Stokes, Mechner, & Balsam, 1999). One possibility is to behave (quasi-) randomly, as indeed was observed in rats in a similar task (Page & Neuringer, 1985). Another one is to behave systematically (nonrandomly)—for example, by orderly emitting all possible sequences in a fixed, long response sequence (e.g., repetition of the sequence 111, 112, 113, 121, 122, 123, . . . etc.; see Machado, 1992, for arguments that *patterns* of sequences may be targets of selection by reinforcement). Therefore, in the present experiments, different feedback contingencies might have a different effect on the number of trials on which specific variability criteria are met, on the behavioral strategy used, or on both. A change of strategy might or might not also imply a change in the number of instances in which variability criteria are satisfied. For this reason, both types of indices of behavioral variability were jointly evaluated.

EXPERIMENT 1

The first condition in Experiment 1 was similar to that in the rat experiments reported by Neuringer et al. (2001): A phase with variability-contingent positive feedback was followed by extinction. The questions of major interest were whether the subjects would learn to adapt their behavior to the variability contingency and whether evidence of response stability and variability similar to that observed in rats would appear as a consequence of extinction. A second condition was included to assess response stability and variability when the contingency change constitutes a transition from consistent absence of feedback to consistent positive feedback for variability. Various lines of research have indicated that a history of consistent negative or no feedback or reinforcement can have a negative effect on subsequent learning. For example, a consistent failure to solve a complicated problem can dramatically impair the ability to subsequently solve a simple problem (e.g., Levine, 1971). Similarly, the consistent failure to escape from an aversive event can result in a subsequent failure to learn to escape from that event even when the experimental situation now allows for an easy escape (a phenomenon well known as learned helplessness; Seligman, 1975). Applied to the present experiment, the consistent absence of positive feedback may have a detrimen-

tal effect on learning to adapt responding to the variability contingency in a subsequent phase.

The following indices of response variability were examined. First, the number of trials on which the variability criteria were met quantified the extent to which each subject fulfilled variability criteria that were similar to those used by Neuringer et al. (2001). This measure reflects the extent to which the subjects learned to adapt their behavior to the feedback contingencies. Second, the relative frequency of the different response options (sequences) was determined for each of the two phases of the experiment, and a U (uncertainty) value was computed (Miller & Frick, 1949). The U value reflects the relative likelihood of each of the different response options. A high U value (high uncertainty) reflects a relatively equal use of a large set of different response options; a low value reflects an uneven use of the different response options. Third, a ratio was computed for each response option by dividing its frequency during Phase 2 of the experiment by its frequency during Phase 1. A high ratio reflects a large relative increase in the use of the response option, a value of 1 signifies an equal use, and a value of <1 reflects a decreased use. As in the animal data reported by Neuringer et al. (2001), this ratio may uncover large changes in the relative use of specific response options, despite the fact that the absolute response frequencies are very small. Fourth, autocorrelations were computed to quantify the extent to which the subjects were responding systematically, using some nonrandom response strategy. The U statistic gives information about the relative uncertainty of the use of each of the 27 single behavioral elements (three-key sequences). Lag 1 to lag 27 autocorrelations were computed to assess whether the behavioral elements, for example, were used in fixed sequences (e.g., simple repetitions of one response alternative, followed by repetitions of another one, or the repetition of a fixed sequence of different behavioral elements) or whether they were used in a random sequence. This refers to variability at the level of the organization of response sequences. Each possible response sequence was given a unique code, and autocorrelations were computed for lag 1 through lag 27. The lag 1 autocorrelation refers to the correlation between the response sequence on the current trial and that on the immediately preceding trial, the lag 2 autocorrelation is the correlation between the response sequence on the current trial and that on two trials back, and so on. Figure 1 shows examples of different patterns of autocorrelations that are associated with different response strategies. Also shown for each pattern is the corresponding percentage of trials on which the variability criteria, as used in the present experiments in the first block of each novel phase, would have been met (*MetVar*; see the Results and Discussion section below).

The first five panels of Figure 1 (Simulations 1–5) depict the lag 1 to lag 27 autocorrelations (from left to right), using data that were generated by a random number generator. Each of these simulations was based on a sequence of 60 integers, with each integer falling within the range

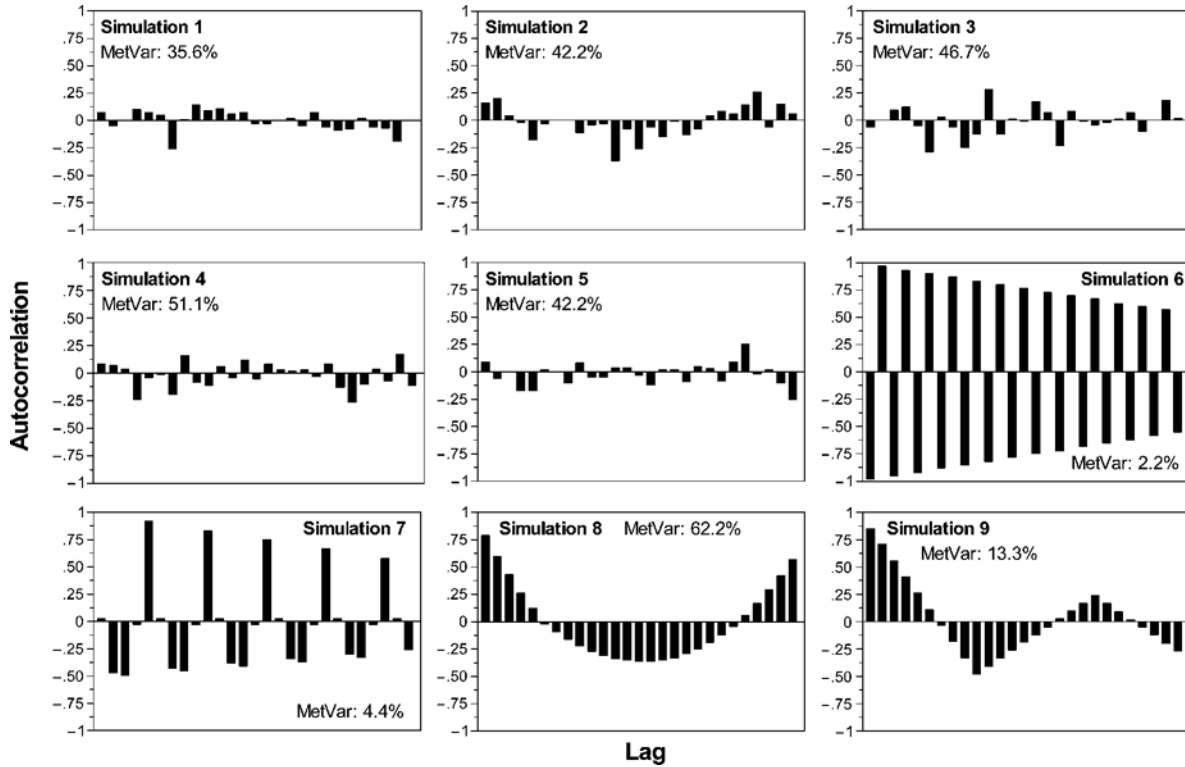


Figure 1. Lag 1 to lag 27 autocorrelations on the basis of simulated data. Each simulation is based on a sequence of 60 integers. Simulations 1–5 are computer-generated random sequences using the integers between 1 and 27. Simulation 6 is based on the consistent alternation of the integers 1 and 2, and Simulation 7 is based on the consistent repetition of a cycle of integers 1, 2, 3, 4, and 5. Simulation 8 is based on cycling the sequence 1, 2, . . . 27, and Simulation 9 is based on six cycles of a repetition of 10 identical integers, with each cycle containing a different integer.

of 1–27 (corresponding to the 27 response options in the experiments). It can be seen that none of the autocorrelations attained a high value, whereas the MetVar value was intermediate. Alternating the numbers 1 and 2 in a sequence of 60 integers generated the pattern shown in Simulation 6. The autocorrelations alternate between highly positive and highly negative, whereas MetVar is very low. A similar pattern is observed in the case of the consistent repetition of a cycle consisting of the numbers 1, 2, 3, 4, and 5, for a total of 12 times (Simulation 7). Simulation 8 displays the pattern corresponding to the consistent repetition of a cycle consisting of the integers 1, 2, 3, . . . 27, for a sequence of a total of 60 numbers. The autocorrelations show a sinus-like shape, whereas MetVar has a relatively high value (in fact, in the actual experiment, this value would be 100% in the 60-trial blocks after the very first one; see below). Finally, Simulation 9 shows the pattern corresponding with the consistent repetition of one integer for 10 times, followed by the repetition of another integer for 10 times, and so on, for a total of 60 numbers. As in Simulation 8, the pattern is sinus-like, albeit with a shorter period and a much lower MetVar. Simulations 6–9 are all examples of systematic or nonrandom behavior. These simulations illustrate that systematic behavior can be either beneficial for satisfying particular variability criteria or not.

When it is favorable, it might even result in a higher MetVar value than is the case with (quasi-) random responding.

Method

Subjects. Twenty-four undergraduate biology students from the University of Nijmegen, 13 males and 11 females, volunteered to participate. Their mean age was 22.2 years (range, 19–25 years).

Apparatus. A notebook PC was located on a table in a quiet laboratory room. The J, K, and L keys were numbered 1, 2, and 3, respectively, using stickers that covered the original letters.

Procedure. The subjects were tested individually. They first received an instruction screen informing them that they would be repeatedly asked to type a sequence consisting of three digits, using only the keys numbered 1, 2, and 3. The following examples of possible sequences were given: 123, 112, 111, and 322. They were further instructed that each time, after having typed a sequence and subsequently having pressed the Enter key, they would receive feedback as to whether the just-typed sequence would be *correct* or not. Positive feedback would consist of the word “Correct.” No feedback would be given in the case of an incorrect sequence. They were told that the purpose of the experiment was to “earn” as much positive feedback as possible. They were further informed that they would be asked to type in a total of 600 sequences and that they would be able to continuously keep track of the total number of sequences already typed. Finally, the subjects were invited, in view of the large number of trials, to work at a steady pace.

The subjects were randomly assigned to one of two experimental conditions: 7 male and 5 female subjects in Condition C–E (variability-

contingent feedback followed by no feedback [extinction]) and 6 male and 6 female subjects in Condition E–C (extinction followed by variability-contingent feedback). At the start of each trial in each condition, the expression “Sequence [x]:” appeared in white letters on the center of the screen against a black background. The term “x” designated the trial number, which was increased by one after each allowable response. This response prompt remained on the screen until the subject had entered an acceptable three-digit sequence and had pressed the Enter key.

The subjects in Condition C–E first received 300 trials in which a sequence was followed by positive feedback if it satisfied two variability criteria (Phase 1). Positive feedback consisted of a blue screen with “Correct” written in white on the center of the screen. One criterion was that the current sequence had to differ from each of the preceding two sequences. The purpose of this criterion was to prevent simple response repetition. The second criterion was that the relative frequency of the current sequence had to be less than or equal to a certain threshold. The relative frequency was computed by dividing the total number of occurrences of the current sequence by the total number of completed sequences (trials). The criterion threshold was set at 0.5 at the outset of the session and was reduced each time the two variability criteria had been satisfied: from 0.5 to 0.05 (steps: 0.5, 0.25, 0.10, 0.075, 0.05). After reaching the 0.05 threshold (typically after no more than about 16 trials), this value was maintained for the remainder of the first phase. The gradual decrease of the threshold was introduced to shape response variability. During the second 300 trials (Phase 2), the subjects in Condition C–E no longer received any feedback. However, for each sequence, the computer registered whether or not the current sequence had met the same variability criteria as those used in Phase 1. To enable a direct comparison between phases, the (dummy) threshold values employed in Phase 2 were exactly the same as those used in Phase 1. Thus, the threshold value was 0.5 at trial 301 and was gradually reduced to 0.05.

The subjects in Condition E–C received the same treatment as did those in Condition C–E, except that the order of phases was switched: No feedback was given in Phase 1, whereas positive feedback was provided in Phase 2 each time the two variability criteria had been met.

Results and Discussion

Depending on the measure examined, evidence was obtained of both response stability and variability as a result of the feedback contingency, and change thereof, in each condition. The subjects in Condition E–C were impaired in adapting their behavior to the contingency in Phase 2.

MetVar. Figure 2 shows the MetVar values for each of the first and last 60-trial blocks of each phase. MetVar was computed according to the following formula: (number of correct trials/maximum number of correct trials that could be earned) * 100. A percentage measure was used because of unequal maximum correct trials for different trial blocks. Specifically, because of the use of a ratio measure to assess fulfillment of one of the variability criteria and the use of a shaping procedure, the maximum was 45 on each of Blocks 1 and 6 (initial blocks of the two phases on which the counters were reset to zero). In contrast, the maximum was 60 on each of the remaining blocks.

The corresponding data for the individual subjects are shown in Figures 3 and 4 (see MetVar values in each panel). The individual subjects’ data indicate that all but one subject (Subject 5) in Condition C–E showed an increase in MetVar from Block 1 to Block 5; in Condition E–C, this was the case for only half the subjects. In Phase 2 (Blocks




Figure 2. Mean percentage (\pm SEM) of the trials in Experiment 1 that satisfied the variability criteria of the maximum number of trials on which a positive feedback could be presented. The subjects in Condition C–E received variability-contingent positive feedback during Blocks 1 and 5 and no feedback during Blocks 6 and 10. The order of feedback contingencies was reversed in Condition E–C.

6 and 10), 7 subjects in Condition C–E showed an increase in MetVar, whereas 5 subjects displayed a decrease. In Condition E–C, the corresponding numbers were 9 and 3 subjects, respectively.

A condition \times block \times phase repeated measures analysis of variance (ANOVA) was performed, using the data in Figure 2 in order to be able to draw conclusions at the group level. This analysis revealed significant effects for block [$F(1,22) = 5.90, p < .05$] and phase [$F(1,22) = 4.52, p < .05$] and a significant condition \times block \times phase interaction [$F(1,22) = 7.80, p < .05$; other $ps > .05$]. To examine the source of the three-term interaction, a separate condition \times block ANOVA was performed on the data from Phase 1 and Phase 2. For the Phase 1 data, the ANOVA revealed a significant main effect of condition [$F(1,22) = 5.40, p < .05$] and a significant condition \times block interaction [$F(1,22) = 6.62, p < .05$]. The latter reflected the fact that the subjects in Condition C–E had a significantly higher MetVar value than did those in Condition E–C on Block 5 [$F(1,22) = 7.76, p < .05$], but not on Block 1 ($F < 1$). Moreover, during Phase 1, there was a significant difference (increase) between blocks for the subjects in Condition C–E [$F(1,22) = 20.82, p < .01$], but not for those in Condition E–C ($F < 1$). The condition \times block ANOVA using the Phase 2 data did not reveal any significant effects [$F_s(1,22) < 1$]. The individual and group data indicate that the contingency between feedback and variability that was in effect in Phase 1 for Condition C–E resulted in a reliable adaptation of responding to that contingency. Moreover, the subjects in Condition E–C failed to show such reliable adaptation in their corresponding variability-contingent phase (Phase 2).

Further simple main effect analyses of block were performed to evaluate the statistical significance of the de-

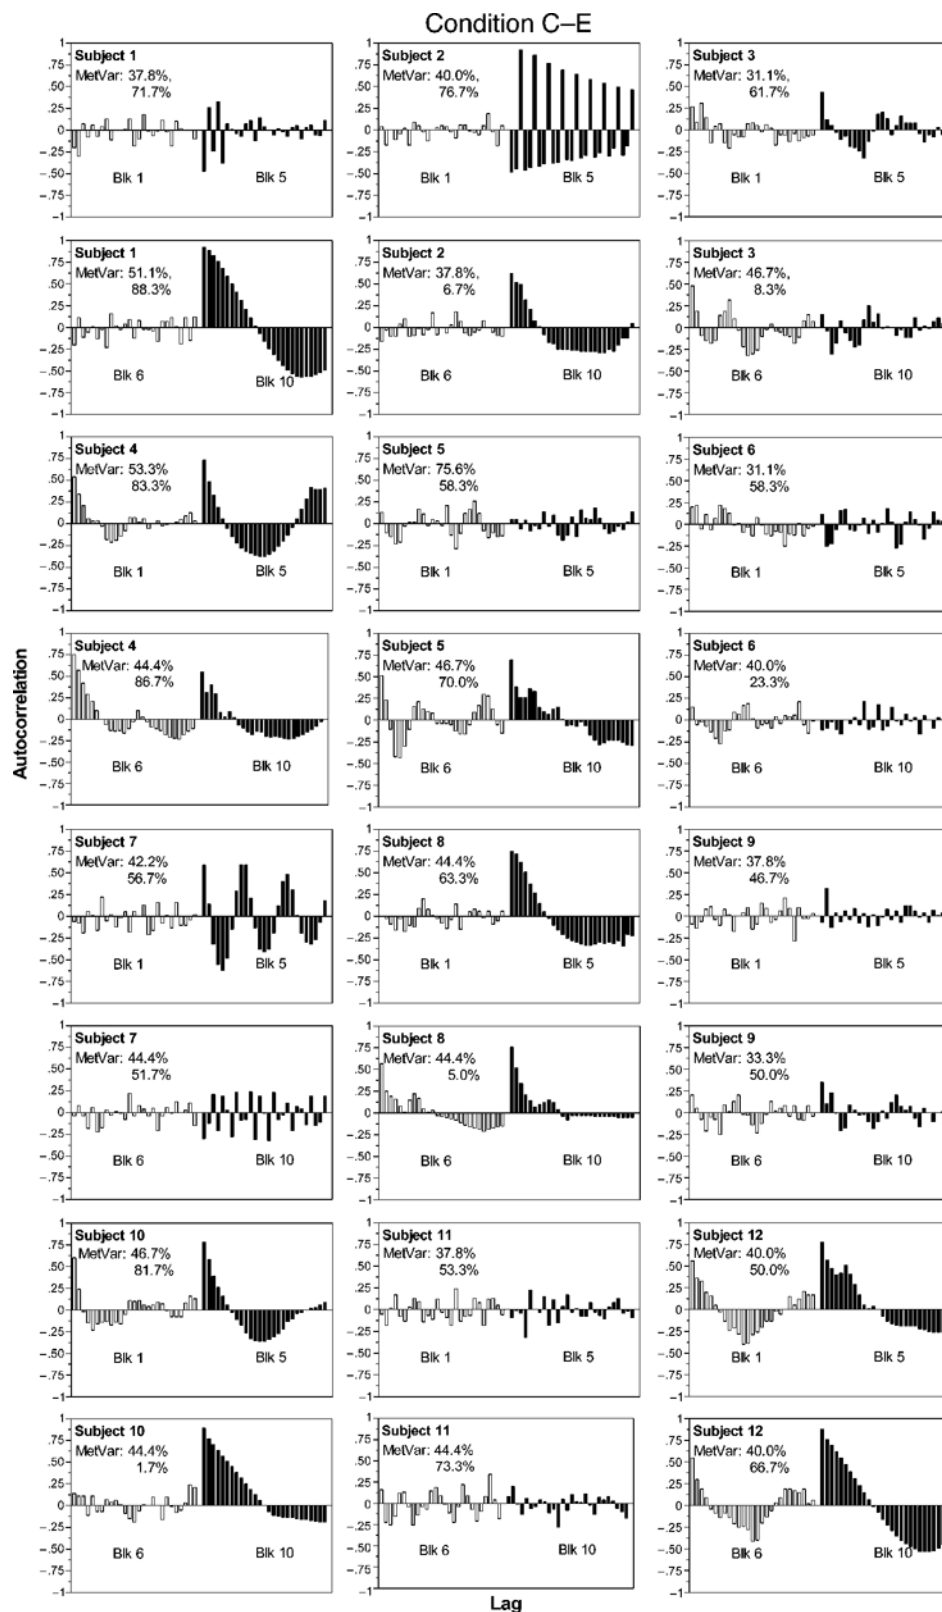


Figure 3. Lag 1 to lag 27 autocorrelations computed on the basis of the responses emitted on Trial Blocks 1, 5, 6, and 10, for the subjects in Condition C-E in Experiment 1. The lag 1 autocorrelation is depicted on the x-axis in the leftmost bar of each block; the lag 27 autocorrelation is shown in the rightmost bar of each block.

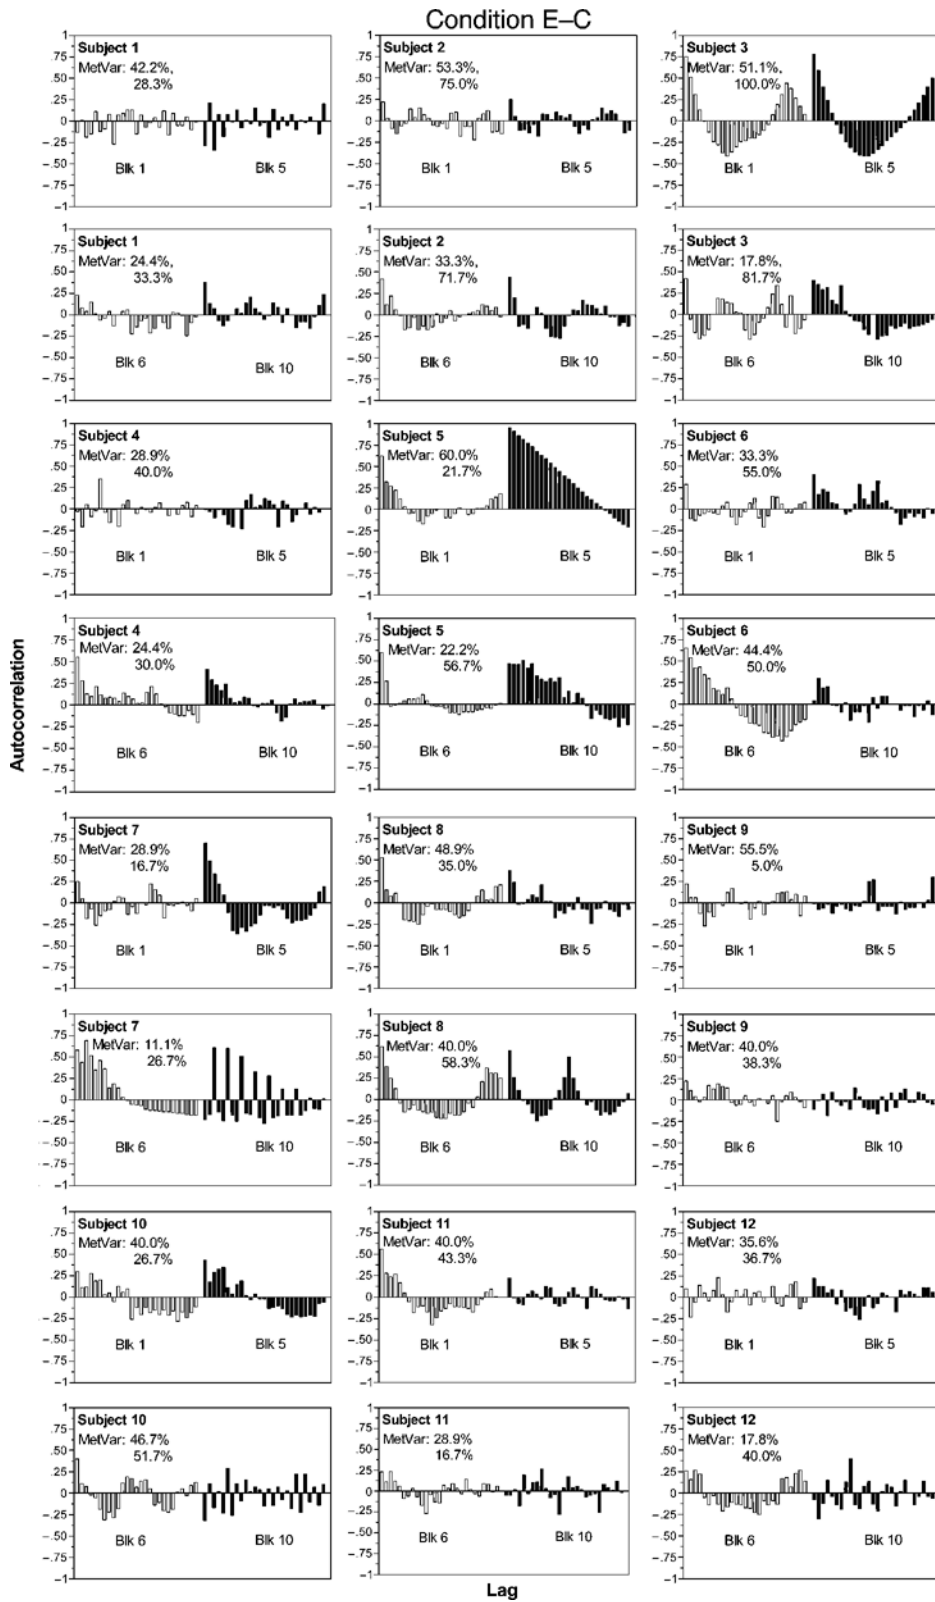


Figure 4. Lag 1 to lag 27 autocorrelations computed on the basis of the responses emitted on Trial Blocks 1, 5, 6, and 10, for the subjects in Condition E-C in Experiment 1. The lag 1 autocorrelation is depicted on the x-axis in the leftmost bar of each block; the lag 27 autocorrelation is shown in the rightmost bar of each block.

crease of MetVar that was observed in each group on the transition from one phase to the other (Block 5 vs. Block 6). These analyses revealed a highly significant decrease for Condition C–E [$F(1,11) = 37.97, p < .001$], but no significant change for Condition E–C [$F(1,11) = 1.66, p > .2$]. The significant decrease in Condition C–E on Block 6 suggests vulnerability to the changed contingency.

A final analysis was performed to evaluate the extent to which the prior nonreinforced Phase 1 in Condition E–C had a detrimental effect on earning positive feedback in Phase 2. To this end, a condition \times block ANOVA was performed using the data from Blocks 1 and 5 for Condition C–E and from Blocks 6 and 10 for Condition E–C. These blocks correspond to each condition's variability-contingent feedback phase. This analysis revealed significant effects for condition [$F(1,22) = 12.73, p < .01$] and

block [$F(1,22) = 25.24, p < .001$], but no significant condition \times block interaction ($F < 1$). The overall higher MetVar value for Condition C–E than for Condition E–C suggests a detrimental effect of no prior exposure to feedback on subsequent learning on the basis of variability-contingent feedback, at least at the beginning of training.

Frequency of response sequences. For each subject, the ordering of the frequency of response sequences was determined as they were emitted in Phase 1, from the most preferred to the least preferred. For each subject and ordered sequence, the frequency was then determined as observed in Phase 2. The frequency of the so-ordered sequences is shown in Figures 5 and 6, separately for each condition. Also shown for each individual is Pearson's correlation coefficient (r) between the frequencies observed in the two phases. This coefficient was used as an

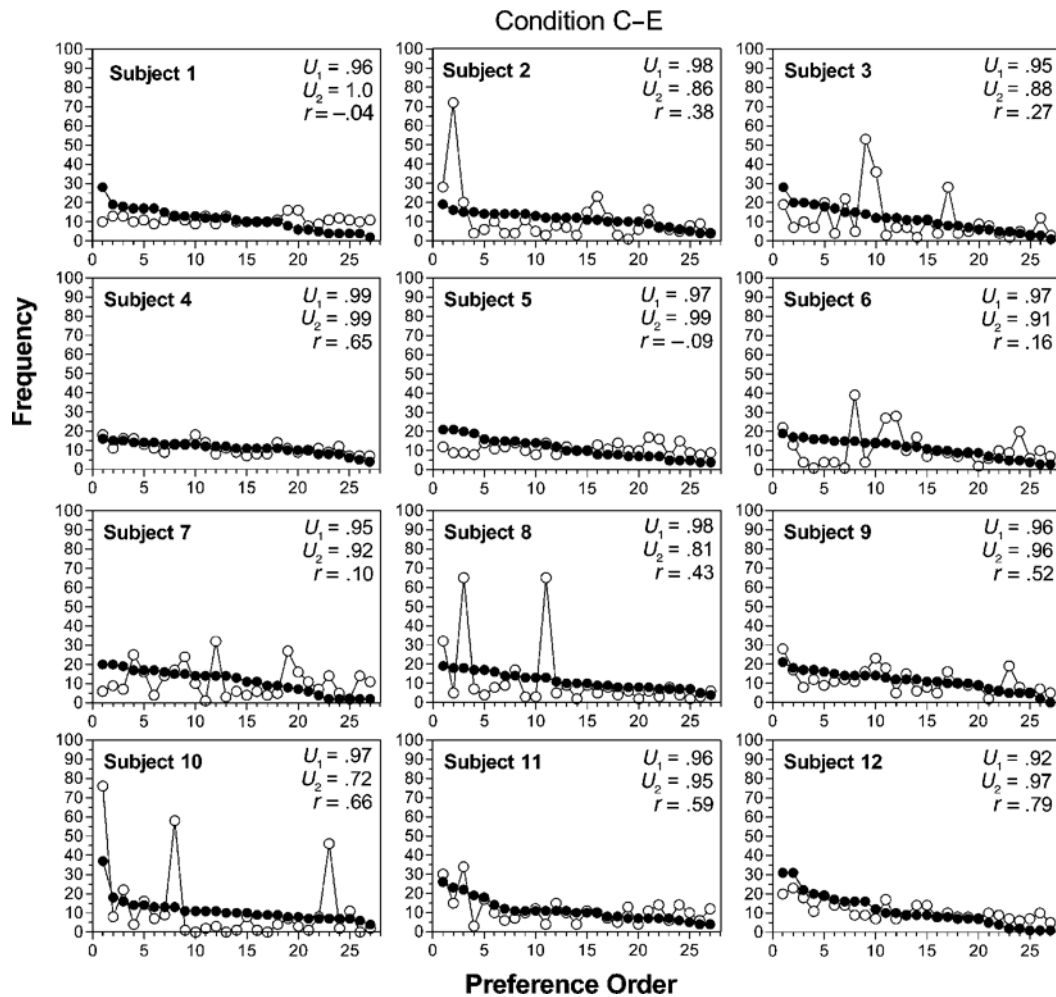


Figure 5. Frequency of the different response sequences in Phase 1 (first 300 trials, filled circles) and Phase 2 (second 300 trials, open circles) for each subject in Condition C–E in Experiment 1. The frequencies are ordered from the most preferred (Preference Order 1) to the least preferred (Preference Order 27) as emitted during Phase 1. The subjects received variability-contingent positive feedback in Phase 1 and no feedback in Phase 2. U_1 and U_2 refer to the U values (see the text) associated with Phases 1 and 2, respectively. The r value signifies Pearson's correlation between the Phase 1 and the Phase 2 frequencies.

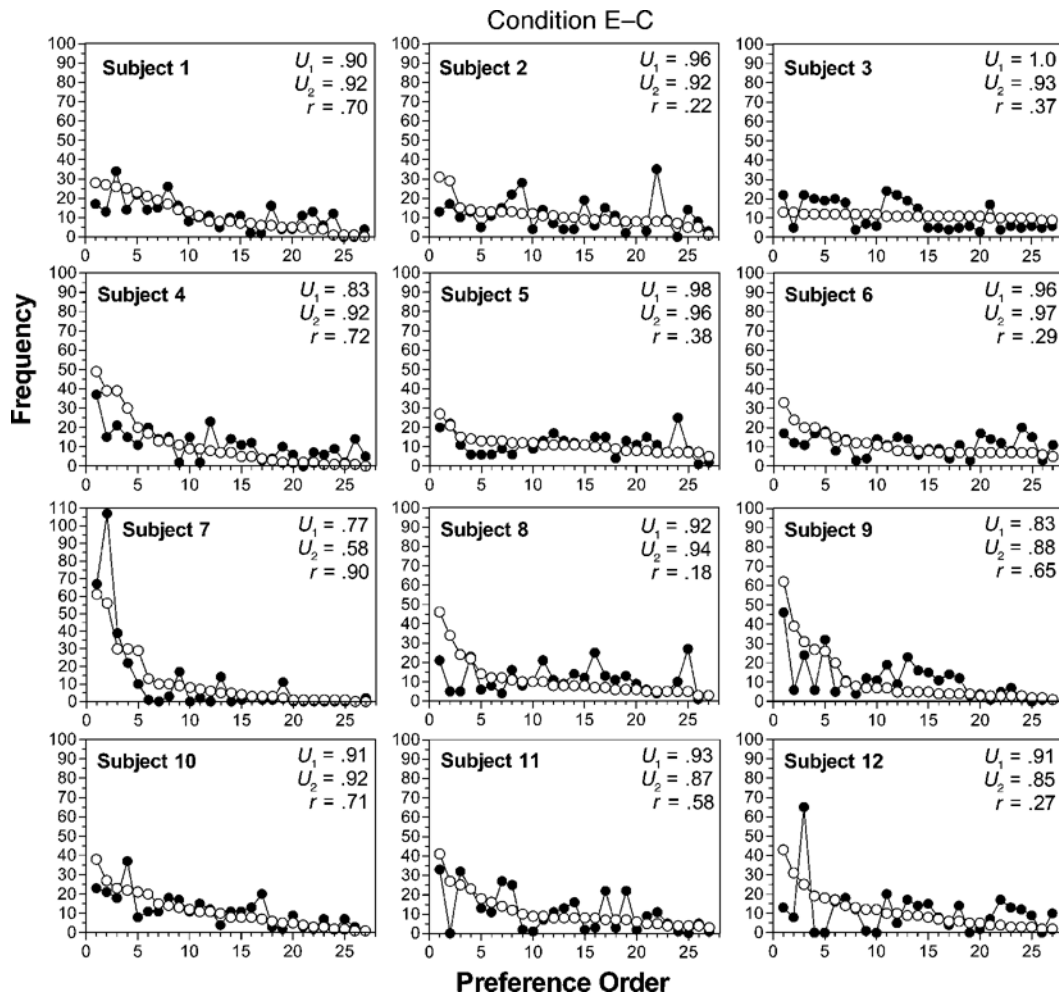


Figure 6. Frequency of the different response sequences in Phase 1 (first 300 trials, open circles) and Phase 2 (second 300 trials, filled circles) for each subject in Condition E-C in Experiment 1. The frequencies are ordered from the most preferred (Order 1) to the least preferred, as emitted during Phase 1. The subjects received no feedback in Phase 1 and variability-contingent positive feedback in Phase 2. U_1 and U_2 refer to the U value (see the text) associated with Phases 1 and 2, respectively. The r value signifies Pearson's correlation between the Phase 1 and the Phase 2 frequencies.

index of response stability. The minimum r value for significance at the $p < .05$ level is .38.

The figures further show each individual's U value associated with Phases 1 (U_1) and 2 (U_2). The U value for each phase and subject was computed according to the formula

$$\frac{-\sum \{RF_i \times [\log(RF_i)] / [\log(2)]\}}{[\log(n)] / [\log(2)]}$$

for $i = 1$ to n , with n designating the number of possible sequences (27) and RF signifying the relative frequency of each sequence. A value of 1 implies maximal uncertainty, and a value of 0 indicates minimal uncertainty.

For all the subjects in Condition C-E except Subjects 3, 6, and 7, the absolute difference in the frequency of use of the different response options was very low and/or r was

significant, implying a strong stability in the use of the different response options across the two phases. Moreover, all the subjects displayed a relatively large U value in each phase, suggesting an even use of the different response sequences in each of the two phases. Similar results were obtained for Condition E-C: Only 3 subjects failed to display both a significant r value and a small absolute difference in frequency of use of the different response alternatives (Subjects 2, 8, and 12). Moreover, the U values again were about equally high for each phase (except perhaps for Subject 7, who had a relatively low U_2 value).

I also computed ratios by dividing the mean frequency during Phase 2 by the mean frequency during Phase 1 for each preference-ordered sequence and condition. For a few subjects on a few occasions, a Phase 1 frequency had a zero

value. In these cases, this value was substituted by 1 to enable the computation of a ratio value; the same held for the corresponding data in Experiment 2. A ratio of 1 reflects no difference in the Phase 2 frequency, relative to the baseline frequency in Phase 1. For all but 3 subjects (Subjects 2, 8, and 10) in Condition C–E, the ratio value gradually increased with increasing preference order (decreasing preference), reflecting an increased use of sequences that had been relatively rare or nonpreferred in Phase 1. The mean ratios based on all 12 subjects in this condition were 1.0, 0.9, and 1.0, for the three most preferred sequences (Preference Orders 1, 2, and 3), respectively. The corresponding ratios for the three least preferred sequences (Preference Orders 25, 26, and 27) were 1.9, 3.2, and 3.0, respectively. A trend analysis using the mean ratio based on the 12 subjects for each preference order revealed a significant linear trend toward increasing ratios with increasing preference order [$F(1,11) = 12.02, ps < .01$]. In Condition E–C, except for Subjects 3, 7, 9, and 11, all the subjects tended also to show an increasing use of relatively rare sequences. The mean ratios for the three most preferred sequences were 0.7, 0.6, and 1.1; the mean ratios for the three least preferred sequences were 1.7, 1.8, and 2.1. A trend analysis on the pooled data again revealed a significant linear trend [$F(1,11) = 6.84, p < .05$]. These findings suggest that, generally, the contingency shift induced an increased use of relatively rare sequences in each condition, implying behavioral variability.

Response strategy: Autocorrelations. Figures 3 and 4 depict, for each condition, the individual subjects' pattern of lag 1 to lag 27 autocorrelations for Trial Blocks 1, 5, 6, and 10.

As can be seen, there were very large differences among the subjects in the extent to which they used a nonrandom, systematic response strategy in the various phases of the experiment. In Condition C–E, some of the subjects used a relatively random response pattern throughout all trial blocks (e.g., Subjects 3, 6, 9, and 11), whereas others showed a random pattern at the initial trial block of each phase, followed by more systematic responding on the last block of each phase (e.g., Subjects 2 and 10). Yet other subjects showed strong systematic responding only on the very last trial block (Subject 1), weak-to-strong systematic responding throughout all the phases (e.g., Subjects 4 and 12), or some other pattern of random and systematic responding. There were also strong individual differences in the nature of systematic responding, if displayed. This could take the form of consistently emitting (very) long sequences of response options (e.g., Subject 1, Block 10; Subject 4, Block 5; and Subject 12, Block 10), yielding a high MetVar. On other occasions, it reflected the repetition of relatively short sequences of response alternatives (e.g., Subject 2, Block 5, and Subject 7, Block 5), which also yielded relatively high MetVar values. A final notable response pattern was the very frequent repetition of a single response alternative (e.g., Subject 2, Block 10; Subject 8, Block 10; and Subject 10, Block 10), yielding a very low MetVar. Similar mixed behavioral patterns were

observed for the subjects in Condition E–C. A few notable patterns in this condition were a very consistent repetition of a long sequence of response options for Subject 3, especially on Block 5 (yielding a 100% MetVar), and many repetitions of single or very short sequences of response options for Subjects 5 (Block 5) and 7 (Blocks 5 and 10), yielding a low MetVar.

To be able to draw a more general conclusion about the (non-) randomness of response patterns, despite the large individual differences, a *summary statistic* was used for each subject: the MaxAut value. This measure refers to the maximum of the 27 autocorrelations computed for each subject for each trial block. The larger the MaxAut, the more the response pattern can be characterized as *systematic*. The mean MaxAut values for Condition C–E were .33, .55, .38, and .56, for Blocks 1, 5, 6, and 10, respectively. The corresponding means for Condition E–C were .39, .44, .45, and .40. The data for each condition and phase (12 subjects) were compared with 12 MaxAut values (mean MaxAut, .26) that were randomly drawn from a pool of 50 computer-generated MaxAut values (mean MaxAut, also .26). The latter values were based on sequences of 60 integers ranging from 1 to 27 (corresponding to the 27 response alternatives) that were generated by a random-number generator (see also the first 5 simulations depicted in Figure 1, which were included in the pool of 50 simulations). ANOVAs with data set (empirical vs. simulated) as a factor revealed a significantly higher empirical than simulated MaxAut value for all conditions and blocks [$F(1,22) > 5.08, ps < .05$], except for Condition C–E in the very first block [$F(1,22) = 2.69, p > .1$].

The collective data and analyses suggest that, as a result of variability-contingent positive feedback, the subjects in Condition C–E increasingly responded in a manner that was favorable for meeting the variability criteria. This result is consistent with previous research with pigeons, rats, and humans, showing that behavioral variability is sensitive to contingent reinforcement or positive feedback (e.g., Machado, 1989; Neuringer, 1986; Neuringer, Deiss, & Imig, 2000; Page & Neuringer, 1985; Ross & Neuringer, 2002; Stokes et al., 1999). As in previous studies, the high MetVar values were achieved at least partly by a relatively even use of the various response alternatives, as reflected in a high *U* value. However, in the present experiment, the increase in MetVar was also achieved at least partly by an increase in systematic responding. As indexed by the MaxAut statistic, the subjects started with a response pattern that was indistinguishable from a computer-generated random pattern in the first block. However, there was a pronounced increase in nonrandom, strategic responding from Block 1 to Block 5. Inspection of the individual data, however, shows that this was due mainly to a changed response pattern in only about one half of the subjects.

The behavioral pattern observed in Condition C–E in Phase 2 suggests a rapid adaptation to the consistent absence of feedback. There was a decrease in MetVar, which was due to the subjects' showing overall a less "MetVar favorable" random or systematic response pattern, either at

the beginning or at the end of this phase. Importantly, although for most subjects there were no large differences in the absolute frequency of the different response alternatives between phases (response stability), in Phase 2 there was a consistent increase in the use of sequences that had been nonpreferred in Phase 1. A similar result was obtained in previous studies with rats (Neuringer et al., 2001). This is indicative of behavioral variability induced by changed response–outcome contingencies at the level of the use of behavioral elements.

In Phase 1 without feedback, the subjects in Condition E–C did not display a significant change in MetVar. Although the mean MaxAut statistic revealed that, overall, the subjects in this condition responded nonrandomly on Block 5, 8 of the 12 subjects showed a rather random response pattern on this trial block (visual inspection of the autocorrelation patterns). More important, in Phase 2, the subjects in this condition had an overall lower MetVar level than that observed for the subjects in Condition C–E under an identical contingency. On Block 10, 10 of the 12 subjects in Condition E–C tended to respond either rather randomly (yielding low-to-intermediate MetVar values) or systematically, but in a manner that was highly unfavorable for attaining a high MetVar. Notably, analysis of the ratio measure suggested that the lack of learning in this condition could not be accounted for completely in terms of the subjects' simply having lost the motivation to try novel behavior, because the ratio measure suggested an increased use of sequences in Phase 2 that had been rare during Phase 1 for the majority of subjects, which is indicative of behavioral variability.

EXPERIMENT 2

Experiment 2 assessed the effects of contingent and noncontingent feedback and of a transition from the one type of contingency to the other. In this experiment, the question of whether the learning effect found for Condition C–E in Phase 1 indeed depends on the contingency between response variability and positive feedback, rather than on the presentation of positive feedback per se, was examined. Moreover, in this experiment, it was asked whether the impaired learning effect observed in Experiment 1 in Condition E–C would also occur after noncontingent feedback. On the one hand, previous animal and human research on, for example, learned helplessness has suggested that noncontingent reinforcement has a detrimental effect on learning under a subsequent condition of contingent reinforcement, just as was found in Experiment 1. On the other hand, previous research with human subjects has suggested that a phase with noncontingent reinforcement has no adverse effect on learning to vary responses in a subsequent phase in which reinforcement is contingent upon behavioral variability (e.g., Saldana & Neuringer, 1998).

Method

Subjects and Apparatus. Sixteen biology students (mean age, 22.2 years; range, 19–26 years) from the University of Nijmegen

(10 males) participated in Experiment 2. None of them had participated in Experiment 1. The students were assigned to one of two conditions (5 males and 3 females in each). The apparatus was the same as that in Experiment 1.

Procedure. The procedure was identical to that in Experiment 1, except for the following. The subjects in Condition C–NC (contingent feedback followed by noncontingent feedback) first completed 300 trials, during which each response sequence was followed by positive feedback if the two variability criteria were met. During the subsequent 300 trials, positive feedback was provided on those trials that corresponded to the trials with positive feedback in the previous phase. For example, if a subject had received positive feedback on Trials 5, 8, and 20 of Phase 1, positive feedback would be given on trials 305, 308, and 320 of Phase 2, irrespective of the response emitted on those trials. The subjects in Condition NC–C (noncontingent feedback followed by contingent feedback) received noncontingent positive feedback during the first 300 trials, using a yoking procedure. Specifically, each subject in Condition NC–C was yoked to a subject in Condition C–NC. For example, if a subject in Condition C–NC had received feedback on Trials 2, 4, and 10, his or her yoked counterpart in Condition NC–C would also receive positive feedback on Trials 2, 4, and 10, irrespective of the response emitted on those trials. During the second 300 trials, the subjects in Condition NC–C received variability-contingent positive feedback, just like those in Condition C–NC did during Phase 1. All further details were the same as those in Experiment 1.

Results and Discussion

Evidence of both response stability and variability was obtained in both conditions, depending on the measure examined. The subjects in Condition NC–C did not demonstrate the impaired learning effect in Phase 2 that was observed in Condition E–C in Experiment 1.

MetVar. Figure 7 depicts the mean MetVar value on Blocks 1, 5, 6, and 10 for each condition. The data of the individual subjects can be found in Figures 8 and 9.

All the subjects in Condition C–NC showed an increase in MetVar from Blocks 1 to 5, whereas only 3 subjects did

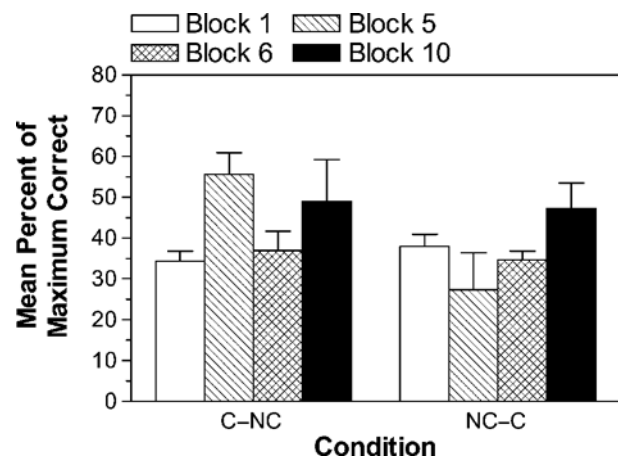


Figure 7. Mean percentage (\pm SEM) of the trials in Experiment 2 that satisfied the variability criteria of the maximum number of trials on which positive feedback could be presented. The subjects in Condition C–NC received variability-contingent positive feedback during Blocks 1 and 5 and noncontingent positive feedback during Blocks 6 and 10. The order of feedback contingencies was reversed in Condition NC–C.

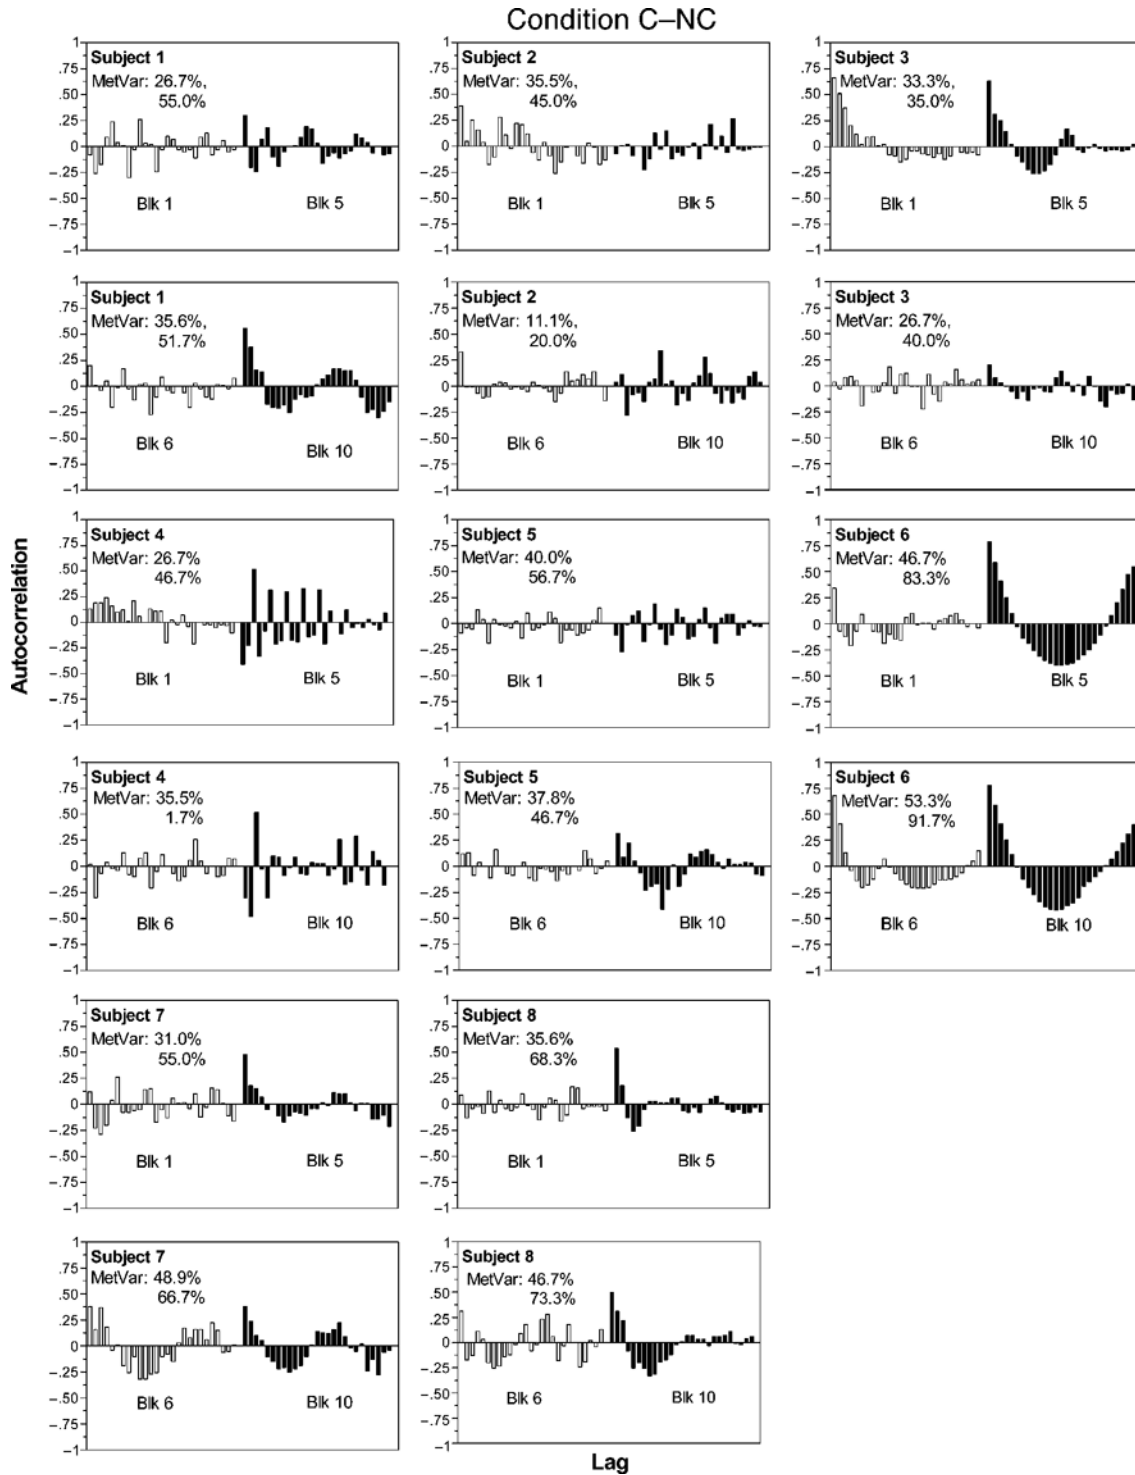


Figure 8. Lag 1 to lag 27 autocorrelations computed on the basis of the responses emitted on Trial Blocks 1, 5, 6, and 10, for the subjects in Condition C-NC in Experiment 2. The lag 1 autocorrelation is depicted on the x-axis in the leftmost bar of each block; the lag 27 autocorrelation is shown in the rightmost bar of each block.

so in Condition NC-C. The other 5 subjects in Condition NC-C showed a decrease. In Phase 2, 7 subjects in Condition C-NC and 5 subjects in Condition NC-C showed a MetVar increase from Blocks 6 to 10.

A condition \times block \times phase ANOVA on the data depicted in Figure 7 revealed a significant effect of block [$F(1,14) = 7.84, p < .05$], a significant condition \times block interaction [$F(1,14) = 6.30, p < .05$], and a significant

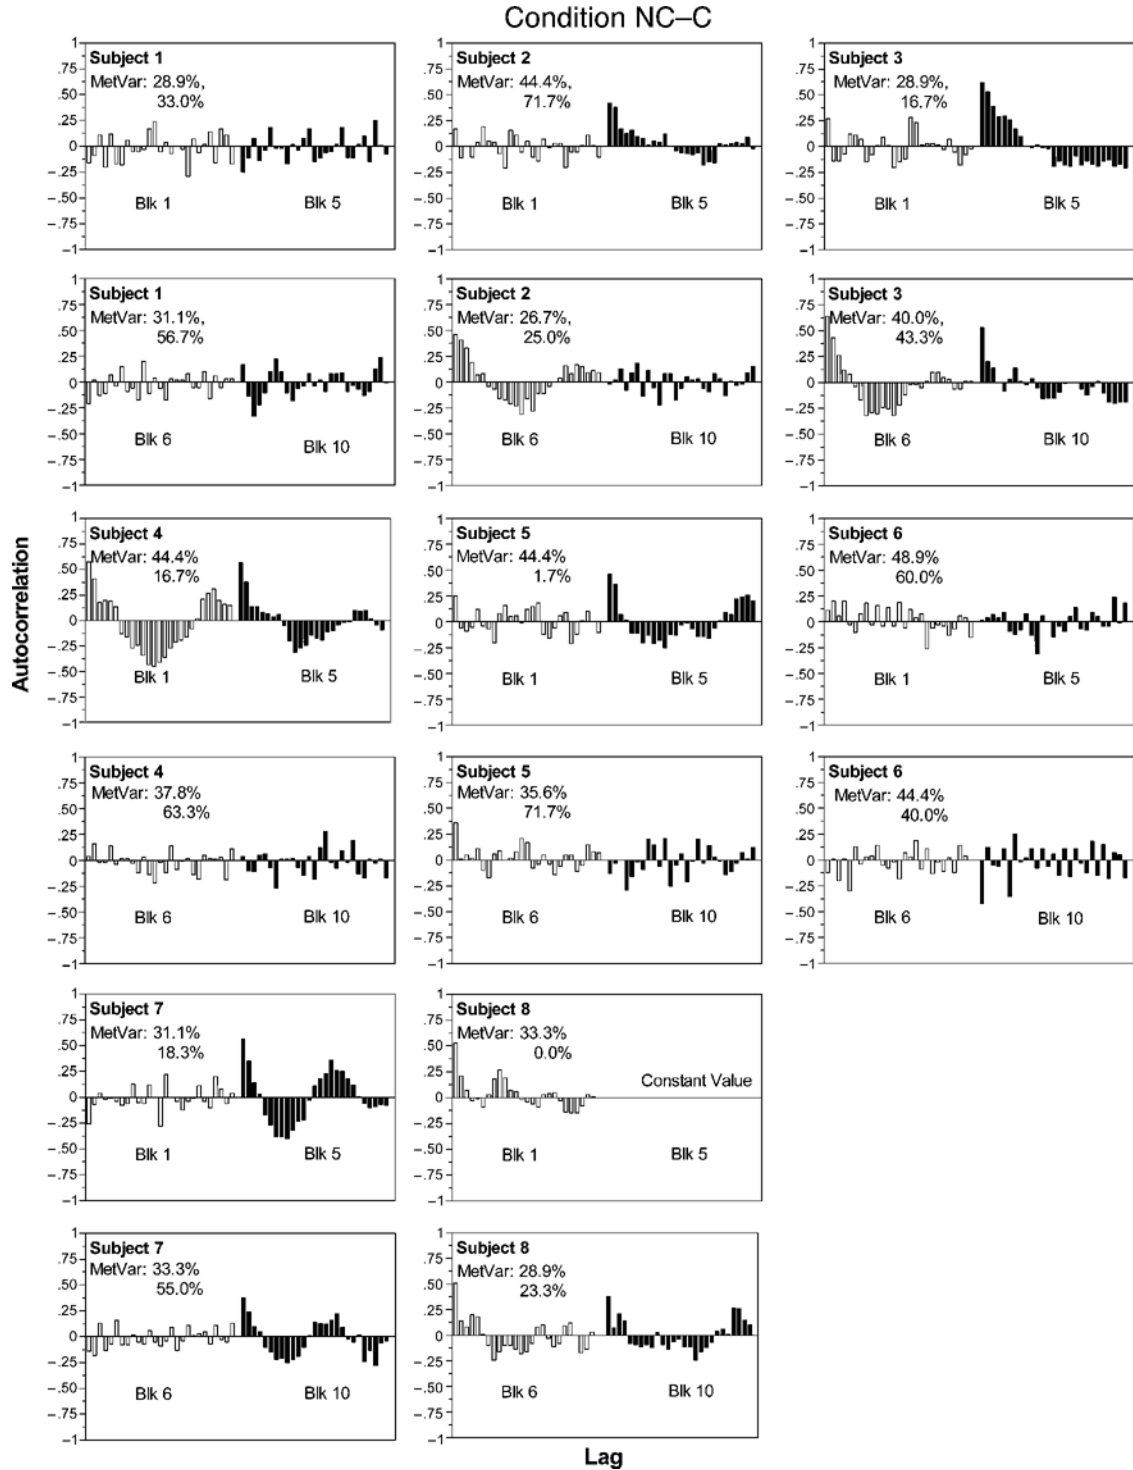


Figure 9. Lag 1 to lag 27 autocorrelations computed on the basis of the responses emitted on Trial Blocks 1, 5, 6, and 10, for the subjects in Condition NC-C in Experiment 2. The lag 1 autocorrelation is depicted on the x-axis in the leftmost bar of each block; the lag 27 autocorrelation is shown in the rightmost bar of each block.

condition \times block \times phase interaction [$F(1,14) = 5.33$, $p < .05$; other $ps > .05$]. A condition \times block ANOVA was performed for each phase separately, to examine the source of the three-term interaction. For the Phase 1 data,

a significant interaction between condition and block [$F(1,14) = 11.53$, $p < .01$; other $ps > .05$] was found, reflecting a significantly higher MetVar for Condition C-NC than for Condition NC-C on Block 5 [$F(1,14) = 7.10$,

$p < .05$], but not on Block 1 ($F < 1$), and a significant increase from Block 1 to Block 5 in Condition C–NC [$F(1,7) = 25.82, p < .01$], but not in Condition NC–C ($F = 1.63$). For Phase 2, the ANOVA revealed only a significant main effect of block [$F(1,14) = 6.82, p < .05$; other $F_s < 1$]. A separate ANOVA with block as a single factor on the data from Blocks 5 and 6 revealed a significant decrease in the percentage from Block 5 to Block 6 for Condition C–NC [$F(1,7) = 28.27, p < .01$], but not for Condition NC–C ($F < 1$). Collectively, these data and analyses imply that noncontingent positive feedback does not result in a reliable increase of MetVar, which, as in Experiment 1, was observed in the case of variability-contingent feedback. Across Phase 2, the majority of the subjects showed an increasing MetVar, suggesting that the prior noncontingent feedback in Condition NC–C did not hinder learning to adapt to the changed contingency as much as was the case for Condition E–C in Experiment 1. Finally, under noncontingent feedback, the subjects in Condition C–NC largely regained responding that was increasingly favorable for meeting the variability criteria.

A final analysis was performed on the data shown in Figure 7 to evaluate whether there was a significant difference between conditions in the increase of the number of

correct trials across contingent-feedback blocks. A condition \times block ANOVA, using Blocks 1 and 5 for Condition C–CN and Blocks 6 and 10 for Condition NC–C, revealed only a significant main effect of block [$F(4,56) = 8.99, p < .001$; other $F_s < 1.61, p_s > .1$]. This implies that the prior noncontingent positive feedback in Condition NC–C did not hinder learning in the subsequent variability-contingent phase.

Frequency of response sequences. Figures 10 and 11 display, for each condition, the mean frequency in Phases 1 and 2 of the sequence corresponding with each preference order, with the frequencies being ordered from high to low as they occurred in Phase 1. Also shown are the individuals' U and r values.

All the subjects in Condition C–NC showed a very small difference in the frequency of the different response alternatives and/or a significant r , implying large response stability. Moreover, except for Subjects 2 and 4 in Phase 2, all the subjects displayed a high U value in each phase, suggesting a similar even use of response options in the two phases. In Condition NC–C, 6 of the 8 subjects also showed a relatively small absolute difference in the frequency of the different response possibilities and a significant r value. The U measure tended to have either the same high value

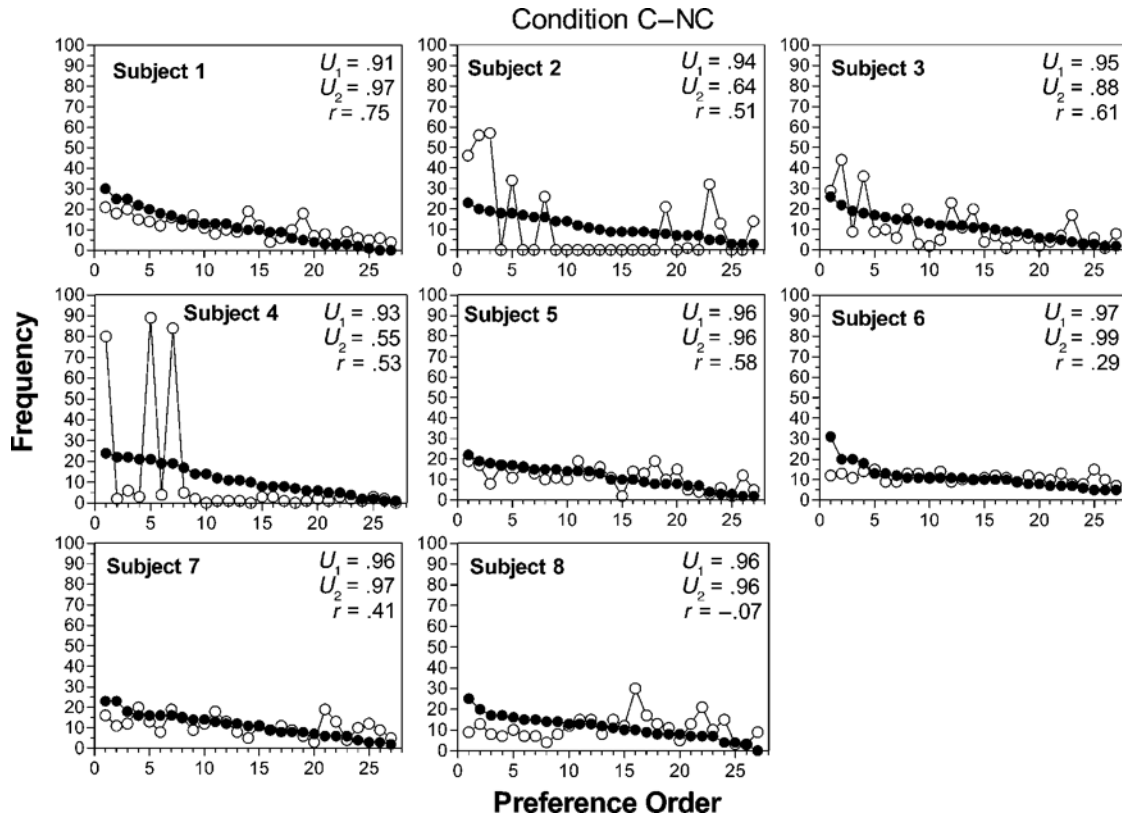


Figure 10. Frequency of the different response sequences in Phase 1 (first 300 trials, filled circles) and Phase 2 (second 300 trials, open circles) for each subject in Condition C–NC in Experiment 2. The frequencies are ordered from the most preferred (Order 1) to the least preferred, as emitted during Phase 1. The subjects received variability-contingent positive feedback in Phase 1 and no feedback in Phase 2. U_1 and U_2 refer to the U values (see the text) associated with Phases 1 and 2, respectively. The r value indicates Pearson's correlation between the Phase 1 and the Phase 2 frequencies.

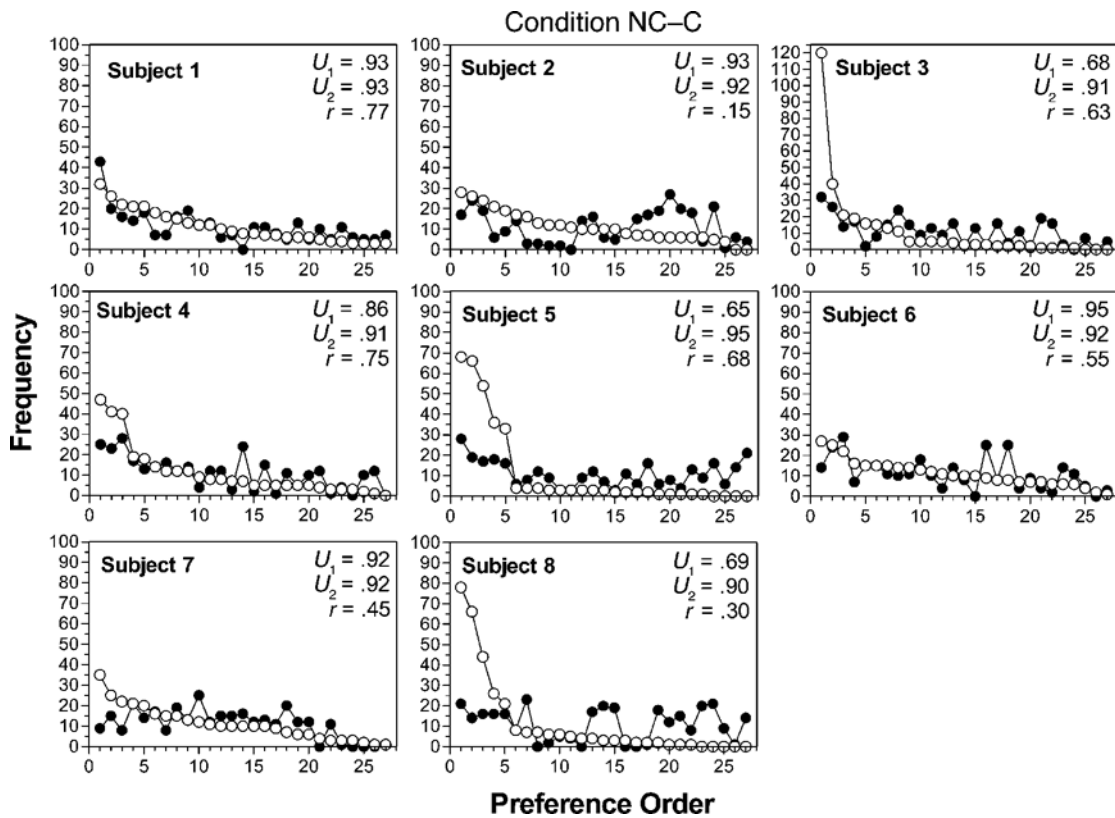


Figure 11. Frequency of the different response sequences in Phase 1 (first 300 trials, open circles) and Phase 2 (second 300 trials, filled circles) for each subject in Condition NC-C in Experiment 2. The frequencies are ordered from the most preferred (Order 1) to the least preferred, as emitted during Phase 1. The subjects received noncontingent positive feedback in Phase 1 and variability-contingent feedback in Phase 2. U_1 and U_2 refer to the U value (see the text) associated with Phases 1 and 2, respectively. The r value indicates Pearson's correlation between the Phase 1 and the Phase 2 frequencies.

in the two phases (4 subjects) or a somewhat smaller value in Phase 1 than in Phase 2 (remaining 4 subjects). These data imply response stability for the majority of the subjects and an enduring or increasing evenness in the use of the different response alternatives.

Seven of the 8 subjects in Condition C-NC mainly showed an increasing use of relatively rare sequences. One subject (Subject 4) mainly showed an increased use of relatively preferred sequences. The mean Phase 2/Phase 1 ratios corresponding to the three most preferred sequences (Orders 1, 2, and 3) were 1.2, 1.0, and 0.8, respectively; those for the three least preferred sequences (Orders 25, 26, and 27) were 2.1, 2.6, and 3.5, respectively. In Condition NC-C, all the subjects tended increasingly to use relatively nonpreferred sequences in Phase 2. The mean ratios for the three most preferred sequences were 0.5, 0.6, and 0.7; those for the three least preferred sequences were 3.8, 4.3, and 6.3, respectively. Both conditions showed a significant linear trend, with higher ratios with increasing preference order [$F_s(1,7) > 7.31, p_s < .05$], suggesting behavioral variability.

Response strategy: Autocorrelations. Figures 8 and 9 show, for Blocks 1, 5, 6, and 10, the individual subjects' patterns of lag 1 to lag 27 autocorrelations. Subject 8 in Condition NC-C continued to emit only one single-response sequence (333) on Block 5, implying the impossibility of computing autocorrelations for that trial block.

As was the case for the subjects in Experiment 1, there were large individual differences in the use of random or systematic response patterns in the different phases of the experiment. For example, in Condition C-NC, the subjects showed relatively random responding on all trial blocks (Subjects 2 and 5), random responding at the first block of each phase and more systematic responding at the last blocks (Subject 4), or some other pattern. Strong systematic responding took the form of either a repetition of fairly short sequences of response options (e.g., Subject 4, Blocks 5 and 10), yielding low to moderate MetVar values, or a repetition of long sequences of response options (e.g., Subject 6, Blocks 5 and 10), yielding a high MetVar. Similar mixed behavioral patterns were observed for Condition NC-C, although here some subjects made

many single-response repetitions on the last block of the noncontingent feedback Phase 1 (e.g., Subjects 3, 4, and 5), yielding a very low MetVar. This pattern was never observed during the corresponding noncontingent block in Condition C–NC (Block 10).

The MaxAut value, as also used in the previous experiment, was employed as an overall index of response randomness. The mean MaxAut values for Condition C–NC was .32, .48, .33, and .46 on Blocks 1, 5, 6, and 10, respectively. The corresponding means for Condition NC–C were .34, .46, .36, and .35. Each of these means was compared (ANOVA) with the mean MaxAut value of a set of eight cases (mean, .26), which were randomly chosen from the set of 50 simulated MaxAut values described in Experiment 1, to assess the (non-) randomness of the response patterns. For Condition C–NC, these analyses suggested a random response pattern in Blocks 1 and 6 [$F_s(1,7) < 1.26, p_s > .05$] and a nonrandom pattern in Blocks 5 and 10 [$F_s(1,7) > 7.66, p_s < .05$]. For Condition NC–C, nonrandom responding was observed only for Block 5 [$F(1,6) = 9.17, p < .05$; other $F_s(1,7) < 4.58, p_s > .05$].

The combined data and analyses suggest that, overall in Phase 1, the subjects in Condition C–NC developed a response strategy, some responding randomly and some nonrandomly, that promoted the occurrence of positive feedback, as was also observed in Experiment 1 for Condition C–E. Right after the contingency change, the subjects in Condition C–NC showed a decrease in MetVar, which was followed by an increase on subsequent blocks. The latter effect had not been observed for Condition C–E in Experiment 1. Condition C–NC's initial decrease and subsequent increase of MetVar in Phase 2 were accompanied by an initial decrease in strategic responding for the majority of the subjects, to the extent that the mean pattern became random, followed by a significant increase in strategic responding. Apparently, the subjects in this condition resumed the behavioral strategy that had been successful in Phase 1, despite the noncontingent nature of the positive feedback in Phase 2. Finally, after the shift, the subjects in Condition C–NC emitted more sequences that had been rare during Phase 1, just as the subjects in Condition C–E did in Experiment 1, reflecting variability at the level of individual sequences.

Overall, the subjects in Condition NC–C increasingly responded systematically across Phase 1. However, the strategy used by the majority of the subjects consisted of repeatedly emitting fixed, relatively short sequences or even simple repetitions of single sequences, which, combined with a relatively uneven use of the different response options, did not result in a significant increase in MetVar. The latter is consistent with previous animal experiments that failed to find increases in response variability with noncontingent reinforcement (e.g., Denney & Neuringer, 1998; Machado, 1989, 1992; Morgan & Neuringer, 1990; Page & Neuringer, 1985). After the shift, there was an increase in the use of rare response sequences and in MetVar. In fact, the increase in MetVar did not statistically dif-

fer from that in Condition C–NC in Phase 1, in which there had been no prior history of noncontingent feedback. Apparently, prior experience of noncontingent feedback does not negatively affect the rate of learning in a subsequent task with contingent feedback as much as is the case with the consistent absence of feedback. However, it did seem to affect the way this learning was achieved: instead of a tendency to develop systematic responding, as in Phase 1 in Conditions C–E and C–NC, the variability criteria were met using a more random response pattern for most of the subjects.

GENERAL DISCUSSION

The present experiments revealed a mix of response stability and variability evoked by different feedback contingencies and changes therein. The results for each of the four conditions will be discussed in view of the outcome of previous, analogous experiments with rats and humans.

Condition C–E

The results for the condition in which variability-contingent feedback was followed by extinction were, in many respects, similar to those for rats in a comparable experimental condition (Neuringer et al., 2001). Both rats and humans appear to be vulnerable to the contingency in the first phase, in that they display a gradual increase in the number of trials that satisfy the variability criteria (see also Denney & Neuringer, 1998; Morgan & Neuringer, 1990; Neuringer, 1986, 1991; Neuringer, Deiss, & Imig, 2000; Ross & Neuringer, 2002). Also, in both species, there is evidence of both behavioral stability and variability as a result of extinction. Stability is apparent in the ordering of the frequencies of the different response possibilities. Sequences that are relatively frequent or rare during contingent feedback are also frequent or rare during extinction. This is also reflected in a similar high U value in the two phases and suggests an overall stability in the structure of responses in terms of the sequences used in the face of changed contingencies in each of the two species. Variability is apparent in both species in the form of an increased use of rare sequences in extinction, relative to baseline frequencies during feedback or reinforcement.

However, a difference between species is that, in extinction, the rats demonstrated an increase in the number of trials that satisfied the variability criterion, whereas the humans showed a decrease. This difference may be related to the fact that, in rats, reinforcement in Phase 1 quite often resulted in a repetition of the just-reinforced response, which is not favorable for meeting the variability criteria. In extinction, no reinforcement was provided, and the immediate reinforcement effect, as present in Phase 1, was eliminated. This, in turn, increased the likelihood of meeting the variability criteria. However, the present experiments included an additional variability criterion—namely, the requirement that the current response had to be different from each of the previous two responses. This explicitly

disfavored simple response repetition in the contingent feedback phase. Consequently, the mechanism as described for the rats presumably did not play a prominent role in the present experiments. More generally, if the criteria for a positive outcome in terms of variability are relatively stringent, a decrease, rather than an increase, in variability may be expected in extinction. Evidence in favor of this claim has recently been obtained in rat studies (Grunow & Neuringer, 2002).

Condition E–C

The consistent absence of positive feedback in Phase 1 evoked a response pattern with an even distribution of the different response alternatives (high U value) and little systematic responding for the majority of the subjects. The latter resulted in a rather overall low number of trials meeting the variability criteria. That there was no increase in systematic responding in the form of simple response repetitions across Phase 1 in the large majority of the subjects might imply that, overall, the subjects were not just “giving up” but continued to try to solve the task at hand. Otherwise, it would have been more likely to see simple response repetitions. After the shift, most of the subjects showed an increasing MetVar across blocks, but the overall MetVar level remained below that acquired by the subjects in Condition C–E in Phase 1. However, the subjects did show a systematic increase in the use of relatively non-preferred sequences, which again suggests a continued motivation to solve the task. The overall lower MetVar values in Phase 2 resemble the outcomes of previous animal research on the effect of an initial failure to solve a problem on the mastery of a subsequent related task, as reflected in, for example, the learned helplessness phenomenon (see the introduction). In the animal literature, this type of phenomenon is suggested to be at least partly related to a reduced associability of responses and consequences. In the present case, the detection of a contingency between variability in responding (as an operant) and positive feedback was at least impaired in the first postshift trial block, an impairment that could not be “restored” to the level attained by Group C–E in the corresponding phase.

Condition C–NC

The results for the subjects in the first, variability-contingent feedback phase of this condition were similar to those observed for Condition C–E in Phase 1. There again was an increase in MetVar across training, which was achieved by the subjects’ using the different response sequences evenly and by (overall) increasing systematic responding. During the initial part of a subsequent noncontingent feedback phase, most of the subjects in this condition displayed a decrease in systematic behavior and MetVar. These changes had also been observed in Condition C–E in Experiment 1. However, interestingly, unlike in Condition C–E, in the course of Phase 2 all but one of the subjects in Condition C–NC again showed a MetVar increase.

This result accords with that reported by Saldana and Neuringer (1998). Using human subjects, these authors found that prior variability-contingent reinforcement resulted in higher levels of variability in a subsequent phase under a yoked (noncontingent) feedback condition than was the case without prior contingent training. These results point to sustained variability, once that variability has been established. However, this persistence appears to require at least the continued presence of some (noncontingent) feedback or reinforcement; it is absent if feedback or reinforcement is withheld altogether (Condition C–E).

Condition NC–C

At least for a small majority of the subjects, noncontingent feedback in Phase 1 failed to result in MetVar and promoted response repetition, with a corresponding relatively uneven frequency distribution (yielding a low U value). Similar results were found in rats (e.g., Denney & Neuringer, 1998; Machado, 1989, 1992; Morgan & Neuringer, 1990; Page & Neuringer, 1985). This suggests that reinforcement per se is not sufficient to evoke variable behavior. However, contrary to the effect of the prior consistent absence of feedback (Condition E–C), prior noncontingent feedback did not seem to negatively affect learning to respond according to the variability requirements in a subsequent phase of contingent feedback. Interestingly, Saldana and Neuringer (1998) reported a similar lack of negative transfer, using children in a completely different task. However, in the present study, there was a difference in the type of strategy eventually used to meet the variability criteria between the subjects in Condition NC–C and those in Conditions C–E and C–NC in Phase 1. The former subjects used a much less systematic (more random) response pattern than did the latter subjects. This suggests that a history of noncontingent feedback can have an effect on subsequent learning, albeit not necessarily on all learning parameters (see also the rat studies in Neuringer, Deiss, & Olson, 2000). At present, the reason for the difference in extent of negative transfer between Conditions E–C and NC–C is unclear and clearly requires further research.

All Conditions

For many of the subjects in each condition, there was a change in the response strategy used, at least at some point within and/or between phases. Thus, variability changes occurred at the level of behavioral *units* (or behavioral repertoire) and at the level at which these units were used in terms of response strategies. There is a fairly large body of literature on rule-governed versus contingency-governed behavior and on differences in sensitivity to contingency changes (e.g., Hayes, Brownstein, Haas, & Greenway, 1986; Hayes, Brownstein, Zettle, Rosenfarb, & Korn, 1986; Hayes & Ju, 1998). However, in most, if not all, of these studies, the strategies or rules were more or less made explicit to the subjects, rather than merely being shaped by reinforcement contingencies, as was the case in the pres-

ent experiments. Unfortunately, to my knowledge, data on contingency-induced changes in response strategies are also largely lacking in the animal conditioning literature. There is a literature on systematic response patterns in animals, but this mostly concerns systematic behavior in the early phase of discrimination problems (see, e.g., Coenders, 1997).

The present results regarding the autocorrelation measure also imply large individual differences in the type of strategies used. At present, the source of these differences is unknown. One possible candidate is a difference in memory capacity. The variability criteria in the present task can better be met by systematically emitting, one by one, all possible response options in a long sequence than by random responding. However, the former strategy is quite demanding in terms of memory load, and the subjects may have differed in their capacity or willingness to use it. The subjects may also have differed in the interpretation of the instruction to "work at a steady pace." This could have been interpreted as implying a stringent time limit or not. Under the assumption that a time pressure prompts more random responding (e.g., because it disfavors memorization of prior response sequences) than does the absence of such a pressure, the former interpretation encourages more random responding than does the latter. Clearly, more research is necessary to uncover the mechanism(s) underlying the individual differences.

In sum, with a task not used before in humans, the present study provides both replications of previous findings and novel results. It shows that behavioral variability, established either by random responding or by some systematic response strategy, may be shaped in humans by contingent positive feedback, but not by noncontingent positive feedback or by the total absence of feedback. That variability is markedly reduced if the contingent feedback is replaced by no feedback (extinction), but not when it is replaced by noncontingent feedback. Moreover, the consistent absence of feedback has a more detrimental effect on learning to behave variably in a subsequent variability-contingent feedback condition than does prior noncontingent feedback. Finally, these different contingencies, and changes therein, all induce a mix of response stability and variability at the level of the frequency of use of the different response alternatives. Some of these results correspond to those found in previous animal and human studies that often used very different procedures, such as a biologically significant reinforcement (food) instead of a biologically nonsignificant outcome (verbal feedback). In doing so, they point to the generality and robustness of these findings.

REFERENCES

- COENDERS, C. J. H. (1997). *Systematic response patterns in discrimination learning*. Unpublished doctoral dissertation, University of Nijmegen, The Netherlands. (Available on line at http://webdoc.ubn.kun.nl/mono/c/coenders_c/systrepai.pdf)
- DENNEY, J., & NEURINGER, A. (1998). Behavioral variability is controlled by discriminative stimuli. *Animal Learning & Behavior*, **26**, 154-162.
- GRUNOW, A., & NEURINGER, A. (2002). Learning to vary and varying to learn. *Psychonomic Bulletin & Review*, **9**, 250-258.
- HAYES, S. C., BROWNSTEIN, A. J., HAAS, J. R., & GREENWAY, D. E. (1986). Instructions, multiple schedules, and extinction: Distinguishing rule-governed from schedule-controlled behavior. *Journal of the Experimental Analysis of Behavior*, **46**, 137-147.
- HAYES, S. C., BROWNSTEIN, A. J., ZETTLE, R. D., ROSENFARB, I., & KORN, Z. (1986). Rule-governed behavior and sensitivity to changing consequences of responding. *Journal of the Experimental Analysis of Behavior*, **45**, 237-256.
- HAYES, S. C., & JU, W. (1998). The applied implications of rule-governed behavior. In W. T. O'Donohue (Ed.), *Learning and behavior therapy* (pp. 374-391). Boston: Allyn & Bacon.
- LEVINE, M. (1971). Hypothesis theory and non-learning despite ideal S-R reinforcement contingencies. *Psychological Review*, **78**, 130-140.
- MACHADO, A. (1989). Operant conditioning of behavioral variability using a percentile reinforcement schedule. *Journal of the Experimental Analysis of Behavior*, **52**, 155-166.
- MACHADO, A. (1992). Behavioral variability and frequency-dependent selection. *Journal of the Experimental Analysis of Behavior*, **58**, 241-263.
- MILLER, G. A., & FRICK, F. C. (1949). Statistical behavioristics and sequences of responses. *Psychological Review*, **56**, 311-324.
- MORGAN, L., & NEURINGER, A. (1990). Behavioral variability as a function of response topography and reinforcement contingency. *Animal Learning & Behavior*, **18**, 257-263.
- NEURINGER, A. (1986). Can people behave "randomly?" The role of feedback. *Journal of Experimental Psychology: General*, **115**, 62-75.
- NEURINGER, A. (1991). Operant variability and repetition as functions of interresponse time. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 3-12.
- NEURINGER, A., DEISS, C., & IMIG, S. (2000). Comparing choices and variations in people and rats: Two teaching experiments. *Behavior Research Methods, Instruments, & Computers*, **32**, 407-416.
- 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.
- ROSS, C., & NEURINGER, A. (2002). Reinforcement of variations and repetitions along three independent response dimensions. *Behavioural Processes*, **57**, 199-209.
- SALDANA, L., & NEURINGER, A. (1998). Is instrumental variability abnormally high in children exhibiting ADHD and aggressive behavior? *Behavioural Brain Research*, **94**, 51-59.
- SELIGMAN, M. E. P. (1975). *Helplessness: On depression, development, and death*. San Francisco: Freeman.
- STOKES, P. D., MECHNER, F., & BALSAM, P. D. (1999). Effects of different acquisition procedures on response variability. *Animal Learning & Behavior*, **27**, 28-41.