Concurrent-chain performance in transition: Effects of terminal-link duration and individual reinforcers

JAMES E. MAZUR

Southern Connecticut State University, New Haven, Connecticut

Pigeons responded on concurrent-chain schedules with variable-interval initial links and equal delays as terminal links. The terminal-link delays were 1 sec in some conditions and 20 sec in other conditions. The percentages of reinforcers delivered for responses on the left key were 10%, 30%, 70%, or 90%, and this percentage was switched every five to nine sessions. The rate of change in the pigeons' response percentages after a switch was the same whether the terminal-link delays were 1 sec or 20 sec. Analysis of the effects of individual reinforcers showed that after a response on one key had been reinforced, response percentages on that key were higher for at least the next 100 responses. Small effects of individual reinforcers were evident after eight or nine additional reinforcers had been delivered. The effects of individual reinforcers were about equally large during times of transition and during periods in which overall response percentages were relatively stable.

In operant research on choice, one popular procedure is the concurrent-chain schedule. A typical concurrent-chain schedule includes a pair of schedules called *initial links*, each of which occasionally leads to its own *terminal link*. Each terminal link is another reinforcement schedule that leads to food. For example, the initial links might be two identical variable-interval (VI) schedules, and the terminal links might be two different fixed-interval (FI) schedules. Response percentages in the initial links are used as measures of a subject's preference for one terminal-link schedule versus the other.

Concurrent-chain schedules have proven to be useful evaluative tools, and they have helped to distinguish among competing theories of choice (e.g., Fantino, 1969; Grace, 1994; Mazur, 2000, 2001; Savastano & Fantino, 1996). However, nearly all of the published research on concurrent-chain schedules has reported only steady-state performance, collected after subjects had been exposed to the same schedules for many daily sessions. These studies provide no information about how choice behavior changes when there is a change in the reinforcement schedules for the different alternatives.

In recent years, there has been increasing interest in how choice behavior changes in response to a change in the reinforcement contingencies, and a number of different models of choice behavior in transition have been proposed (e.g., Couvillon & Bitterman, 1985; Davison & Baum,

2000; Dragoi & Staddon, 1999; Mazur, 1992). Studies on this topic have examined how such factors as the overall rate of reinforcement, the types of reinforcement schedules, and the frequency of the changes in the reinforcement contingencies affect how animals adjust to the change. In this body of research, both global aspects of transitional choice behavior (such as the gradual change in response percentages when one alternative starts to deliver more reinforcers) and local effects (such as the effects of each individual reinforcer on subsequent choices) have been examined. Most of these studies have used simple concurrent reinforcement schedules (e.g., Davison & Baum, 2000; Dreyfus, 1991; Mark & Gallistel, 1994; Mazur, 1992) or discrete-trial procedures (Bailey & Mazur, 1990; Couvillon & Bitterman, 1985; Fischer, Couvillon, & Bitterman, 1993). Only a few published reports on choice behavior in transition have used concurrent-chain schedules (Grace, 2002; Mazur, Blake, & McManus, 2001). Because of this scarcity of data, very little is known about how preference develops with concurrent-chain schedules and how preference changes when there is a change in the schedules.

In the present experiment, the choice responses of pigeons on fairly simple concurrent-chain schedules were examined in order to address some basic questions about transitional choice behavior. Throughout the experiment, the terminal links were identical fixed-time (FT) schedules, which required no responses but delivered food after either 1 sec (in some conditions) or 20 sec (in other conditions). The initial links were VI schedules that delivered terminal-link entries at different rates: In different conditions, the VI schedule associated with the left response key delivered 10%, 30%, 70%, or 90% of the terminal-link entries. The left-key reinforcement percent-

This research was supported by Grant MH 38357 from the National Institute of Mental Health. I thank Diane Coe, Heidi Goodby, Phong Voong, and Joel Yudt for their help in various phases of the research. Correspondence concerning this article should be sent to J. E. Mazur, Psychology Department, Southern Connecticut State University, New Haven, CT 06515 (e-mail: mazur@southernct.edu).

age was changed every few sessions, and the resultant changes in the pigeons' choice responses were examined.

One basic question was how the duration of the terminal links (FT 1 sec vs. FT 20 sec) would affect the overall rate of change in the pigeons' choice percentages after a switch in the reinforcement percentages. A common view about concurrent-chain schedules is that the terminal links are conditioned reinforcers, because they are paired with a primary reinforcer (food), and that the strengths of these conditioned reinforcers is inversely related to their durations (e.g., Killeen, 1982; Squires & Fantino, 1971; Vaughan, 1985). One reasonable extrapolation from this view (although certainly not the only one) is that the rate of learning will depend on the strength of these conditioned reinforcers. If so, the rate of change in the pigeons' choice responses should be much greater when the terminal links are FT 1 sec than when they are FT 20 sec. A different approach, however, is to view each terminal-link entry as one piece of information about the current reinforcement percentages. From this perspective, the rate of change in choice responses should be the same with FT 1-sec and FT 20-sec terminal links, because they provide the same amount of information about the reinforcement percentages. These are only two of many possible hypotheses about how varying terminal-link duration might affect rates of change in choice behavior, but they do suggest two very different outcomes for the present experiment.

A second basic question addressed in this study was how the occurrence of individual reinforcers would affect subsequent choice responding. Using a similar concurrent-chain procedure, Mazur et al. (2001) found that the percentage of responses on the *rich key* (the key currently associated with a shorter terminal link or a higher rate of reinforcement) was higher immediately after a rich-key reinforcer than after a lean-key reinforcer. The effect of an individual reinforcer persisted for at least the next 30 responses if no additional reinforcer had occurred. However, once another reinforcer had been delivered, no residual effect of the previous reinforcer could be detected in most cases. In contrast, Davison and Baum (2000) found that the effects of individual reinforcers could be detected even after six or more subsequent reinforcers had been delivered, and they concluded that "every reinforcer counts." The procedures used by Davison and Baum differed from those of Mazur et al. in at least two major ways. First, Davison and Baum used concurrent VI schedules, rather than concurrent-chain schedules. Second, their procedure featured very frequent changes in the reinforcement schedules: Seven different pairs of VI schedules were presented in a random order each session, with a 10-sec blackout signaling the start of each new pair of schedules. It could be that individual reinforcers have more pronounced effects on subsequent responding when the reinforcement schedules change more frequently.

In the present experiment, the reinforcement schedules changed only once every five to nine sessions, and

the effects of individual reinforcers were examined both immediately after a change and several sessions later. The effects of the 1-sec and 20-sec terminal links were also compared, to determine whether reinforcers delivered after shorter delays had larger aftereffects on the pigeons' choice responses.

METHOD

Subjects

The subjects were 4 White Carneaux pigeons maintained at about 80% of their free-feeding weights. All had previous experience with a variety of experimental procedures.

Apparatus

The experimental chamber was 30 cm long, 30 cm wide, and 31 cm high. The chamber had three response keys, each 2 cm in diameter, mounted in the front wall of the chamber, 24 cm above the floor and 8 cm apart, center to center. The center key was not used in this experiment. A force of approximately 0.15 N was required to operate each key. Each key could be transilluminated with lights of different colors. A hopper below the center key provided controlled access to grain, and when grain was available, the hopper was illuminated with a 2-W white light. Pairs of 2-W lights (two white, two green, and two red) were mounted above the Plexiglas ceiling of the chamber. The chamber was enclosed in a sound-attenuating box containing a ventilation fan. All the stimuli were controlled and the responses recorded by an IBM-compatible personal computer using the Medstate programming language.

Procedure

Experimental sessions were usually conducted 6 days a week. Throughout the experiment, a concurrent-chain procedure was used, in which a single VI 45-sec schedule operated in the initial links and each terminal link was an FT schedule that consisted of a fixed delay followed by a 3-sec food presentation. In the initial links, the white houselights were lit, and the two side keys were illuminated, the left key green and the right key red. The VI 45-sec schedule assigned terminal links to the two response keys, using a pseudorandom sequence that ensured that the actual percentage of terminal links for each key was close to the assigned percentage for that session (cf. Stubbs & Pliskoff, 1969). Once a terminal link was assigned to one key, the VI timer stopped and did not restart until that terminal link was entered and completed. A 3-sec changeover delay (COD) was in effect, so that no reinforcer was delivered until at least 3 sec had elapsed after a switch from one key to the other. When a terminal-link entry was assigned to a key, the next post-COD peck on that key extinguished all keylights, and the terminal link began.

The terminal link for the left key was a fixed delay with the green houselights on, and the terminal link for the right key was a fixed delay with the red houselights on. Each terminal link ended with a 3-sec presentation of grain, and only the white light above the grain was lit during the reinforcement period. After each food presentation, the keylights and white houselights were again illuminated, and the next initial link began. Each session ended when the cumulative initial-link time reached 40 min.

The experiment consisted of three phases, with 13 conditions in Phase 1, 13 conditions in Phase 2, and 7 conditions in Phase 3. All the terminal links were 1 sec in duration in Phases 1 and 3, and all the terminal links were 20 sec in Phase 2. In the different conditions within each phase, the percentages of reinforcers assigned to the left key were 10%, 30%, 70%, or 90%. The number of sessions per condition varied randomly between five and nine sessions, and the left-key reinforcement percentage was changed at the start of each new condition.

The data from the first condition of each phase were not used in the data analysis, because there was no transition from a previous

condition with the same terminal-link duration. The remaining 12 conditions from Phases 1 and 2 included every possible transition among the four reinforcement probabilities (e.g., a transition from 10% to 30%, from 10% to 70%, from 10% to 90%, and so on). The last 6 conditions of Phase 3 included only half the possible transitions for any given bird, but the reinforcement percentages were counterbalanced across subjects, so all 12 possible transitions were equally represented in Phase 3 as well.

RESULTS

Unless otherwise noted, all the analyses were based on the first five sessions of each condition. This was done in order to keep the comparisons across conditions comparable, because individual conditions varied in length from five to nine sessions.

Response Rates and Transition Rates

The overall response rates in the initial links averaged 71.5 and 76.0 responses per minute in the two phases with 1-sec terminal links (Phases 1 and 3, respectively) and 36.4 responses per minute with 20-sec terminal links (Phase 2). A repeated measures analysis of variance (ANOVA) showed a significant difference across phases [$F(2,6) = 23.07, p < .01$]. These results show that terminal-link durations had the expected effect on the strength of initial-link responding.

One of the main questions of this experiment was whether the rate of behavioral change at a point of transition would be different with terminal links of different

durations. To answer this question, Figure 1 shows the percentage of responses on the key with the higher reinforcement percentage after a change in percentages, for each of the three phases, averaged across all subjects and all conditions. The data are separated into 4-min blocks of initial-link time, including the last three 4-min blocks of the session before a new condition, and the first three sessions of each new condition. (Throughout this article, *rich response key percentage* will refer to responses on the key that has the higher reinforcement percentage after a transition, even though this key frequently has the lower response percentage before the transition. This is why the rich response key percentage is less than 50% in the left panel of Figure 1.) This figure shows that, despite the large differences in overall response rates with 1-sec versus 20-sec terminal links, the changes in the rich response key percentage after a switch in the reinforcement percentages were very similar in all three phases. A repeated measures ANOVA was conducted on the data from the first two sessions after each transition, which, as Figure 1 shows, was when most of the changes in response percentages occurred. Rich response key percentages were calculated for successive 20-min blocks. There was a significant change in rich response key percentages across blocks [$F(3,9) = 88.73, p < .001$] and a significant effect of reinforcement percentage [$F(5,15) = 50.32, p < .001$] but no significant effect of phase [$F(2,6) = 0.41$] and no significant phase $\times$ reinforcement percentage interaction [$F(10,30) = 1.45$]. This analysis

Figure 1. For each of the three phases of the experiment, rich response key percentages are shown in 4-min blocks of initial-link time. The vertical lines indicate breaks between sessions. The first three data points from each phase are from the last 12 min of sessions that preceded a transition in reinforcement percentages, and the remaining data points are from the first three sessions that followed a transition. *Rich response key percentage* refers to responses on the key with the higher reinforcement percentage after the transition.

therefore suggests that there was no significant difference in transition rates across the three phases of the experiment.

How did the rate of change in response percentages at the start of a new condition vary as a function of the size of the difference between the old and the new reinforcement percentages? For example, when there was a large change in the responses percentages (e.g., from 10% to 90%), did response percentages change more rapidly than, more slowly than, or at the same rate as when there was a small change in the reinforcement percentages

(e.g., from 10% to 30%)? To answer this question, Figure 2 shows the response percentages for each of the six possible transitions from a lower to a higher reinforcement percentage. Because Figure 1 showed no differences between the FT 1-sec and the FT 20-sec conditions, the data from all three phases were pooled to produce the results shown in Figure 2. This figure shows the rich response key percentage in 4-min blocks, beginning with the last 12 min of the session preceding a switch in the reinforcement percentages, followed by the results from the first three sessions after a switch. There

Figure 2. For the six different combinations of reinforcement percentages before and after a transition in reinforcement percentages, rich response key percentages are shown in 4-min blocks of initial-link time, pooled across all three phases of the experiment. The vertical lines indicate breaks between sessions. The first three data points are from the last 12 min of sessions that preceded a transition, and the remaining data points are from the first three sessions that followed a transition. *Rich response key percentage* refers to responses on the key with the higher reinforcement percentage after the transition.

are two ways to compare the results from the different conditions. In absolute terms, it is clear that there were much larger changes in the rich response key percentage after a large change in the reinforcement percentages than after a small change. However, in relative terms, it appears that it took the pigeons roughly the same amount of time to switch from the previous response percentage and arrive at a relatively stable new response percentage. Notice that in all six cases, response percentages changed the most in the first session after a switch, much less in the second session, and very little in the third session.

To analyze the rate of change in relative terms, rich response key percentages in the different conditions were calculated as a proportion of the total change from the pretransition response percentage to the asymptotic response percentage after a transition. For each subject and each condition, the mean rich response key percentage in the last session before a transition was used as the starting point, and the mean rich response key percentage from the fifth session after a transition was used as an estimate of the asymptote. The proportion of change from the starting point to asymptote was calculated for the first four 20-min blocks after a transition. These proportions therefore measured the rate of change toward the new asymptotes, while factoring out the absolute size of

the change from the starting point to the asymptote. A repeated measures ANOVA conducted on these rescaled data showed a significant change across blocks [$F(3,9) = 54.35, p < .001$] but no significant effect of reinforcement percentage [$F(5,15) = 2.73$], no significant effect of phase [$F(2,6) = 0.07$], and no significant interactions. In summary, when response percentages were rescaled as a proportion of the total change toward a new asymptote, rates of change were not significantly different with the different reinforcement percentages or with the different terminal-link durations.

Effects of Individual Reinforcers

Several analyses were conducted in order to examine the effects of individual reinforcers on subsequent responding. The first analysis divided the data from all the subjects from the first five sessions of all the conditions into four groups, on the basis of the rich reinforcement percentage (70% or 90%) and the terminal-link duration (1 or 20 sec). Figure 3 shows, on a response-by-response basis, the rich response key percentage on the 100 successive responses that followed a rich-key reinforcer, as compared with the 100 responses that followed a lean-key reinforcer. For these calculations, a sequence ended whenever a terminal link occurred, and a new sequence began

Figure 3. Rich response key percentages are shown for the 100 responses that followed rich-key or lean-key reinforcers, for conditions with rich-key reinforcement percentages of 90% (top panels) and 70% (bottom panels). FT, fixed time.

at the start of each new initial link. Because terminal-link entries often occurred before 100 responses, the number of data points used to calculate the percentages decreased as the serial position of a response increased. For example, in the top left panel of Figure 3, the response percentage for the 1st response after a rich-key reinforcer was based on 6,771 cases, whereas the response percentage for the 100th response after a rich-key reinforcer was based on 1,063 cases. The fewest cases occurred for the lean-key reinforcers in the 90% conditions with FT 20-sec terminal links (top right panel of Figure 3), which is why the data points are more variable, especially many responses after a reinforcer.

All four panels of Figure 3 show that rich response key percentages were substantially higher after a rich-key reinforcer than after a lean-key reinforcer, and this difference persisted for at least 100 responses. In other words, responding increased on whichever key had just delivered a reinforcer. For about the first 5 responses, this effect was much larger after 1-sec terminal links than after 20-sec terminal links, but for the remaining responses, the size of the effect was similar for the two terminal-link durations. More specifically, for the first 5 re-

sponses after a reinforcer, the mean difference between rich-key and lean-key response percentages was 28.3% with the 1-sec terminal links and 12.5% with the 20-sec terminal links. However, for the next 95 responses after a reinforcer, the mean difference between rich-key and lean-key response percentages was 5.6% with 1-sec terminal links and 5.7% with 20-sec terminal links. In summary, the greater effect of a reinforcer in the conditions with 1-sec terminal links was short-lived and lasted for only about 5 responses after a reinforcer.

Did individual reinforcers have a greater effect in the first session after a change in the reinforcement percentages (when response percentages were changing rapidly, as can be seen in Figure 1) and less effect several sessions later (when response percentages were relatively stable)? To answer this question, the analyses described in the preceding paragraph were conducted on data from only the first session after a transition and on the data from the fifth session after a transition. Figure 4 shows the pooled results from all conditions with FT 1-sec terminal links (top panels) and FT 20-sec terminal links (bottom panels). With both terminal-link durations, the main differences between the first and the fifth sessions

Figure 4. Rich response key percentages are shown for the 100 responses that followed rich-key or lean-key reinforcers, for conditions with 1-sec terminal links (top panels) and 20-sec terminal links (bottom panels). The left panels show the results from the first session after a switch in reinforcement percentages, and the right panels show the results from the fifth session after a switch.

were the overall response percentages. Rich response key percentages were, of course, higher in the fifth session after a change in the reinforcement percentages than in the first session. However, the differences in rich response key percentages after a rich-key versus after a lean-key reinforcer were just about as large in the fifth session as in the first session, despite the fact that rich response key percentages were changing rapidly in the first session but showed no systematic changes in the fifth session. Averaged across the 100 responses that followed a reinforcer in the FT 1-sec conditions, the mean difference between response percentages after a rich-key reinforcer and after a lean-key reinforcer was 8.2% in the first session and 9.3% in the fifth session. In the FT 20-sec conditions, the difference was 6.8% in the first session and 8.2% in the fifth session. Thus, the rapid in-

creases in rich response key percentages observed in the first session of each new condition (Figures 1 and 2) were evidently not the product of an increased effectiveness of individual reinforcers.

Davison and Baum (2000) found that the effects of individual reinforcers could be seen even after several subsequent reinforcers had been delivered. For instance (letting L and R stand for left-key and right-key reinforcers, respectively), they found a higher left-key response percentage after the sequence LLL than after the sequence RLL. This shows that a right-key reinforcer continued to have an effect on responding even after two subsequent left-key reinforcers had occurred. Davison and Baum's procedure featured very rapid changes in the reinforcement percentages (with at least seven different reinforcement percentages per session), and the following analyses

Figure 5. Left-key response percentages after different sequences of reinforcers in the conditions with left-key reinforcement percentages of 70% and 30%. The sequences of reinforcers are plotted from left to right, so the sequence represented by each data point two or three reinforcers back can be determined by following the lines that connect the point to the open or closed circles to its right.

were conducted to see whether similar effects of previous reinforcers could be observed in the present experiment, which had transitions every five to nine sessions.

To measure the effects of past reinforcers, left-key response percentages were calculated for every possible sequence of one, two, or three reinforcers. Separate calculations were done for each of the four reinforcement percentages used in this experiment. Figure 5 shows the left-key response percentages after different sequences of reinforcers in the conditions with left-key reinforcement percentages of 70% and 30%. Filled circles represent left-key reinforcers, open circles represent right-key reinforcers, and the sequences of reinforcers are plotted from left to right. The data points plotted three reinforcers back represent the eight possible sequences of three reinforcers (LLL, LLR, etc.). Similarly, the four data points plotted two reinforcers back represent the

four possible sequences of two reinforcers (LL, LR, RL, and RR). The sequence of reinforcers for each data point can be determined by following the lines that connect the points. For example, the top-left data point in the upper panel is a filled circle, and it is followed by two more filled circles, so this data point shows the average left response percentage after all LLL sequences (i.e., after three left-key reinforcers). The open circle directly below this data point is connected to the same two filled circles, so it represents the left response percentage after all RLL reinforcer sequences.

Figure 5 shows that for both the 70% and the 30% conditions, individual reinforcers had measurable effects both two and three reinforcers back. Notice that every filled circle three reinforcers back, representing a left-key reinforcer, has a higher left-key response percentage than does the corresponding open circle, representing a

Figure 6. Left-key response percentages after different sequences of reinforcers in the conditions with left-key reinforcement percentages of 90% and 10%. The sequence of reinforcers represented by each data point two reinforcers back can be determined by following the line that connects the point to the open or closed circle to its right.

right-key reinforcer. On average, the difference between left-key and right-key reinforcers was largest one reinforcer back and smallest three reinforcers back, showing a progressive decline in the effect of a reinforcer as additional reinforcers occurred. Figure 6 shows similar results from the conditions with left-key reinforcement percentages of 90% and 10%. In this case, the results

from three-reinforcer sequences are not shown, because with these more extreme reinforcement percentages, several of the three-reinforcer sequences occurred too rarely to provide an adequate sample. The results are similar to those in Figure 5: Each left-key reinforcer two reinforcers back led to a higher response percentage, but the effects were smaller than those for one reinforcer back.

Figure 7. For each of the four left-key reinforcement percentages, left-key response percentages are compared for cases in which a left-key or a right-key reinforcer occurred from 1 to 10 reinforcers back. In every case, the data are pooled across all sequences of left-key and right-key reinforcers that occurred after the reinforcer in question.

Examining response percentages after sequences of more than 2 or 3 reinforcers becomes problematic because of the smaller sample sizes associated with the many different possible reinforcer sequences. However, there is another way to examine the effects of past reinforcers that avoids the problem of reduced sample sizes. The left response percentage can be calculated for all cases in which a left-key reinforcer occurred 5 reinforcers back, for example, and can be compared with the percentage for all cases in which a right-key reinforcer occurred 5 reinforcers back, pooling the data across all the different reinforcer sequences that may have occurred for the next 4 reinforcers. Figure 7 presents the results from this type of analysis for the four different left-key reinforcement probabilities, showing data from up to 10 reinforcers back. The patterns for the four different reinforcement probabilities are similar. Immediately after a left-key reinforcer, left response percentages averaged about 8% higher than immediately after a right-key reinforcer. For reinforcers that occurred previously, the difference between left-key and right-key reinforcers gradually diminished, and by 9 or 10 reinforcers back, there was little or no difference in current response percentages. Although the differences in percentages after left versus right reinforcers were not large, they were statistically significant. An ANOVA showed a significant effect of reinforcer location [left or right; $F(1,3) = 77.51$, $p < .01$] and a significant interaction between reinforcer location and the number of reinforcers back [$F(9,27) = 11.11$, $p < .001$].

DISCUSSION

Many previous studies on concurrent-chain schedules have shown that choice responses are strongly affected by the durations of the terminal-link schedules (e.g., Fantino & Davison, 1983; MacEwen, 1972). In the present experiment, with equal terminal-link delays for the two alternatives, initial-link response rates were much higher with FT 1-sec terminal-link schedules than with FT 20-sec schedules. Nevertheless, the changes in rich response key percentages after a switch in the reinforcement percentages were very similar in the three phases of the experiment, regardless of terminal-link duration (Figure 1). One possible conclusion is that it is not the conditioned reinforcing strength of a terminal link that controls the rate of behavioral change but, rather, the information about the current reinforcement percentages that each terminal-link entry provides. Regardless of its duration, each terminal-link entry provides the same information about what initial-link percentages are currently in effect.

There are, however, several ways in which the present findings could be reconciled with a conditioned reinforcement approach. For example, according to delay reduction theory (Squires & Fantino, 1971), the conditioned reinforcing strength of a terminal link depends on the reduction in time to reinforcement that its onset sig-

nals. Regardless of whether the terminal links were 1 or 20 sec in the present experiment, they signaled a 45-sec reduction in the average time to reinforcement (because the initial-link schedule was VI 45 sec in all the conditions). Thus, it could be argued that the rates of behavioral change were similar because the amounts of delay reduction were identical with both terminal-link durations. Another possible explanation follows from a recent study by Grace (2002). He found that acquisition was faster if the terminal links after a transition were shorter but also that acquisition was slower if the terminal links before the transition were shorter. Grace (2002) argued that more highly valued terminal links after a transition produce faster acquisition but that more highly valued terminal links before a transition interfere with acquisition. In the present experiment, the terminal-link durations were identical before and after a transition, so these opposing effects of the preceding and the following schedules may have canceled out, leading to the observed result: no difference in rates of behavioral change with 1-sec and 20-sec terminal links. In summary, there are several possible explanations for why the 20-fold variation in terminal-link duration had no measurable effect on the rates of behavioral change after a transition in the reinforcement percentages.

This study also shows that the magnitude of the change in reinforcement percentages had no systematic effect on the rates at which response percentages approached a new asymptote. This is not to say that the observed behavioral changes were the same in all cases. Figure 2 clearly shows much larger changes in the rich response key percentages after a large change in reinforcement percentages (e.g., from 10% to 90%) than after a small change (e.g., from 10% to 30%). However, when response percentages were measured as a proportion of the total change from the pretransition percentages to the asymptotic percentages after a transition, there was no significant difference across the different reinforcement percentages. As is shown in Figure 2, much of the change in rich response key percentages took place in the first half of the session immediately after a transition, and by the third session after a transition, rich response key percentages showed little or no systematic changes. These results are similar to those from a previous study on choice behavior in transition (Mazur, 1992), which also found that rates of change were not systematically related to the magnitude of the change in reinforcement percentages.

Just as with overall transition rates, an analysis of the effects of individual reinforcers found little difference between the FT 1-sec and the FT 20-sec terminal links, with one exception. Figure 3 shows that for the first several responses after a reinforcer, the differences in response percentages after rich-key and lean-key reinforcers were larger with the shorter terminal links. However, after about the first 5 responses, the effects of the previous reinforcer were about equally large with the

1-sec and the 20-sec terminal links. In both cases, if no additional reinforcers were delivered, the effects of the previous reinforcer were still evident 100 responses later.

Similar aftereffects of reinforcers have been observed in previous studies of concurrent VI schedules (Menlove, 1975) and concurrent-chain schedules (Mazur et al., 2001). However, other studies have found the opposite: a *negative recency effect*, in which there was a lower response percentage on the alternative that had just delivered a reinforcer (Fantino & Royalty, 1987; Killeen, 1970). Fantino and Royalty proposed that negative recency effects are most likely to be seen when the initial links are two independent VI schedules, because immediately after a reinforcer is delivered for one alternative, the probability of reinforcement is temporarily increased for the opposite alternative. Conversely, when a single VI schedule arranges the reinforcers for both alternatives (as in the present experiment), there is no particular advantage to selecting the opposite alternative immediately after reinforcement.

One interesting aspect of the aftereffects of individual reinforcers is that they were just as large in the fifth session after a switch in the reinforcement percentages as in the first session after a switch (Figure 4). One might have supposed that individual reinforcers would have a greater effect on subsequent responding immediately after a transition (when response percentages were changing rapidly at a global level) than several sessions later (when response percentages were relatively stable). Indeed, just such an effect was observed by Mazur et al. (2001) with a concurrent-chain procedure in which the key with the shorter terminal-link delay (e.g., 5 vs. 20-sec) was periodically switched. Mazur et al. found that the aftereffects of individual reinforcers were about three times as large in the first session after the terminal links were switched as they were eight sessions later. Furthermore, their experiment included a few conditions in which reinforcement percentages were switched (rather than the terminal-link delays). In those conditions, there were no decreases in the aftereffects of individual reinforcers across sessions, a finding that was corroborated by the much more extensive set of data of the present experiment. Taken together, these results suggest that immediately after some types of transitions in the reinforcement contingencies (i.e., changes in terminal-link delays), the effects of individual reinforcers are temporarily enhanced, whereas after other types of transitions (i.e., changes in reinforcement percentages), there is no such enhancement. Additional research will be needed to determine why the effects of individual reinforcers are increased after some types of changes in the reinforcement contingencies, but not after others.

In their study on the effects of individual reinforcers when the reinforcement schedules changed seven or more times per session, Davison and Baum (2000) showed that individual reinforcers had effects on choice behavior that lasted even after several additional reinforcers were delivered. Were these clear and persistent

effects of individual reinforcers specific to their procedure, with its rapid switches in reinforcement schedules? Their procedure was quite different from that of most previous experiments using operant conditioning methodology to study choice, in which the same reinforcement schedules are kept in effect for many sessions. In the present experiment, the same reinforcement percentages were used for between five and nine sessions, and the effects of individual reinforcers were similar to those found by Davison and Baum in at least two ways. First, although somewhat different methods of analysis had to be used because of the different experimental designs of the two studies, the effects of specific sequences of two or three reinforcers (Figures 5 and 6) were similar to those observed by Davison and Baum (see their Figure 15). Second, by pooling large amounts of data, the effects of individual left-key or right-key reinforcers could be detected even after as many as seven or eight subsequent reinforcers had been delivered (Figure 7).

These results suggest that, at least to a first approximation, individual reinforcers may have similar effects on choice behavior, regardless of whether the reinforcement contingencies have been changing every few minutes or only about once a week. In both cases, individual reinforcers have measurable effects on subsequent choice behavior, which gradually decreases as additional responses are made and as additional reinforcers are delivered. Many theories about conditioning in general (e.g., Bush & Mosteller, 1955; Estes, 1950; Killeen, 1994) and about choice behavior in particular (e.g., Couvillon & Bitterman, 1985; Devenport & Devenport, 1993; Mazur, 1992) include the core assumption that individual reinforcers have effects that diminish gradually over time, by a process such as an exponentially weighted moving average. These hypothesized effects of individual reinforcers can be observed in ongoing choice behavior, but averaging over large samples of data may be necessary to see them.

REFERENCES

- BAILEY, J. T., & MAZUR, J. E. (1990). Choice behavior in transition: Development of preference for the higher probability of reinforcement. *Journal of the Experimental Analysis of Behavior*, **53**, 409-422.
- BUSH, R. R., & MOSTELLER, F. (1955). *Stochastic models for learning*. New York: Wiley.
- COUVILLON, P. A., & BITTERMAN, M. E. (1985). Analysis of choice in honeybees. *Animal Learning & Behavior*, **13**, 246-252.
- DAVISON, M., & BAUM, W. M. (2000). Choice in a variable environment: Every reinforcer counts. *Journal of the Experimental Analysis of Behavior*, **74**, 1-24.
- DEVENPORT, J. A., & DEVENPORT, L. D. (1993). Time-dependent decisions in dogs (*Canis familiaris*). *Journal of Comparative Psychology*, **107**, 169-173.
- DRAGOI, V., & STADDON, J. E. R. (1999). The dynamics of operant conditioning. *Psychological Review*, **106**, 20-61.
- DREYFUS, L. R. (1991). Local shifts in relative reinforcement rate and time allocation on concurrent schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, **17**, 486-502.
- ESTES, W. K. (1950). Toward a statistical theory of learning. *Psychological Review*, **57**, 94-107.
- FANTINO, E. (1969). Choice and rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, **12**, 723-730.

- FANTINO, E., & DAVISON, M. (1983). Choice: Some quantitative relations. *Journal of the Experimental Analysis of Behavior*, **40**, 1-13.
- FANTINO, E., & ROYALTY, P. (1987). A molecular analysis of choice on concurrent-chain schedules. *Journal of the Experimental Analysis of Behavior*, **48**, 145-159.
- FISCHER, M. E., COUVILLON, P. A., & BITTERMAN, M. E. (1993). Choice in honeybees as a function of the probability of reward. *Animal Learning & Behavior*, **21**, 187-195.
- GRACE, R. (1994). A contextual model of concurrent-chains choice. *Journal of the Experimental Analysis of Behavior*, **61**, 113-129.
- GRACE, R. (2002). Acquisition of preference in concurrent chains: Comparing linear-operator and memory-representational models. *Journal of Experimental Psychology: Animal Behavior Processes*, **28**, 257-276.
- KILLEEN, P. R. (1970). Preference for fixed-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior*, **14**, 127-131.
- KILLEEN, P. R. (1982). Incentive theory: II. Models for choice. *Journal of the Experimental Analysis of Behavior*, **38**, 217-234.
- KILLEEN, P. R. (1994). Mathematical principles of reinforcement. *Behavioral & Brain Sciences*, **17**, 105-172.
- MACEWEN, D. (1972). The effects of terminal-link fixed-interval and variable-interval schedules on responding under concurrent-chained schedules. *Journal of the Experimental Analysis of Behavior*, **18**, 253-261.
- MARK, T. A., & GALLISTEL, C. R. (1994). Kinetics of matching. *Journal of Experimental Psychology: Animal Behavior Processes*, **20**, 79-95.
- MAZUR, J. E. (1992). Choice behavior in transition: Development of preference with ratio and interval schedules. *Journal of Experimental Psychology: Animal Behavior Processes*, **18**, 364-378.
- MAZUR, J. E. (2000). Two- versus three-alternative concurrent-chain schedules: A test of three models. *Journal of Experimental Psychology: Animal Behavior Processes*, **26**, 286-293.
- MAZUR, J. E. (2001). Hyperbolic value addition and general models of animal choice. *Psychological Review*, **108**, 96-112.
- MAZUR, J. E., BLAKE, N., & McMANUS, C. (2001). Transitional choice behavior in concurrent-chain schedules. *Behavioural Processes*, **53**, 171-180.
- MENLOVE, R. L. (1975). Local patterns of responding maintained by concurrent and multiple schedules. *Journal of the Experimental Analysis of Behavior*, **23**, 309-337.
- SAVASTANO, H. I., & FANTINO, E. (1996). Differences in delay, not ratios, control choice in concurrent chains. *Journal of the Experimental Analysis of Behavior*, **66**, 97-116.
- SQUIRES, N., & FANTINO, E. (1971). A model for choice in simple concurrent and concurrent-chains schedules. *Journal of the Experimental Analysis of Behavior*, **15**, 27-38.
- STUBBS, D. A., & PLISKOFF, S. S. (1969). Concurrent responding with fixed relative rate of reinforcement. *Journal of the Experimental Analysis of Behavior*, **12**, 887-895.
- VAUGHAN, W. (1985). Choice: A local analysis. *Journal of the Experimental Analysis of Behavior*, **43**, 383-405.

(Manuscript received February 20, 2002;
revision accepted for publication June 10, 2002.)