Published in final edited form as:

J Exp Anal Behav. 2013 March ; 99(2): 159–178. doi:10.1002/jeab.12.

Pigeons' Choices with Token Stimuli in Concurrent Variable-Interval Schedules

James E. Mazur and Dawn R. Biondi
Southern Connecticut State University

Abstract

Twelve pigeons responded on concurrent variable-interval schedules that delivered token stimuli (stimulus lights for some pigeons, and white circles on the response keys for others). During exchange periods, each token could be exchanged for food on a fixed-ratio 1 schedule. Across conditions, the exchange requirements (number of tokens that had to be earned before they could be exchanged for food) varied between one and four for the two response keys. The main findings were that the pigeons' response percentages varied as a function of the number of tokens earned at any given moment, and they were determined by both the delays to food and by the number of food deliveries in the exchange periods. In some conditions, tokens had to be earned but were not visible during the variable-interval schedules for one or both keys. When one key had visible tokens and the other did not, the pigeons showed a preference for the key without visible tokens. A model based on the matching law and a hyperbolic delay-discounting equation could account for the main patterns of choice responding, and for how response percentages changed as successive tokens were earned. The results are consistent with the view that the token stimuli served as discriminative stimuli that signaled the current delays to food.

Keywords

token stimuli; concurrent variable-interval schedules; generalized matching law; hyperbolic discounting; keypeck; pigeons

Token stimuli exert substantial control over human behavior in everyday life. Whether the tokens are money, airline bonus miles, college course credits, or supermarket coupons, people spend large amounts of time and effort earning and collecting them, and making sure they have received their fair share of them. Many experiments have shown that the delivery of tokens can affect the behavior of nonhuman organisms as well. Early work with chimpanzees demonstrated that as long as tokens could be exchanged for food, water, or some other primary reinforcer, they could function as conditioned reinforcers—they could be used to maintain existing behaviors or to train completely new behaviors (Cowles, 1937; Wolfe, 1936). Kelleher (1956) showed that when responding on various reinforcement schedules led to tokens (poker chips that could be exchanged for food), chimpanzees' behavior patterns on these schedules were similar to those obtained on schedules with food as the reinforcer. For example, when the poker chips were delivered on fixed-ratio (FR) schedules, the animals exhibited postreinforcement pauses followed by rapid responding (Kelleher, 1958). When tokens were delivered on fixed-interval (FI) schedules, accelerating response patterns were observed, just as when food is delivered in FI schedules (Kelleher, 1957).

Token stimuli have been studied with other species as well. Malagodi (1967) used marbles as tokens with rats, and he found that their response patterns on FR and variable-interval (VI) schedules of token reinforcement were similar to those obtained with food reinforcement. In a series of experiments, Hackenberg and his colleagues have shown that stimulus lights can be used as tokens for pigeons (e.g., Bullock & Hackenberg, 2006; Raiff, Bullock, & Hackenberg, 2008). The tokens were light-emitting diodes (LED) arranged in a horizontal row above the response keys. One LED was lit each time a token was earned during *token-production periods*; later, during *exchange periods*, each LED could be exchanged for one food presentation (and the LEDs were turned off, one at a time, as each was exchanged for food). As in the studies with chimpanzees and rats, the response patterns of the pigeons on schedules of token reinforcement resembled those observed on reinforcement schedules with food reinforcers.

Findings such as these have provided strong evidence that token stimuli can serve as conditioned reinforcers, in the sense that they can be used to shape, strengthen, and maintain behaviors in much the same way that primary reinforcers do. However, it has often been noted that typical conditioned reinforcers may serve other functions besides strengthening behaviors. Williams (1994) discussed some of these functions, including marking and bridging. The term *marking* refers to the notion that a conditioned reinforcer can provide immediate feedback for a particular response, thereby increasing the salience of that response. *Bridging* suggests that a conditioned reinforcer can fill the gap between a response and the delivery of a primary reinforcer, which may help the animal associate the response and the primary reinforcer. In a recent review, Hackenberg (2009) described evidence that token stimuli can serve other functions as well. For instance, because tokens are associated with primary reinforcers such as food or water, they sometimes appear to act as Pavlovian conditioned stimuli, eliciting conditioned responses associated with these primary reinforcers (e.g., rats chewing on ball bearings that could be exchanged for food; Boakes, Poli, Lockwood, & Goodall, 1978). In addition, there is considerable evidence that tokens can serve as discriminative stimuli that indicate when primary reinforcers will or will not be delivered (e.g., Kelleher, 1958; Yankelevitz, Bullock, & Hackenberg, 2008).

A few studies have used tokens in choice situations, and these studies have helped to distinguish between the reinforcing and discriminative properties of tokens. Jackson and Hackenberg (1996) used a discrete-trial procedure to give pigeons a choice between one LED presented immediately and three LEDs presented after a delay (and each LED was eventually exchanged for one food delivery). In different experiments, they varied both the delays between responses and token deliveries and the delays between token deliveries and food. They found that the pigeons tended to choose one immediate token over three delayed tokens (similar to the “impulsive” choices found when pigeons choose between immediate and delayed food reinforcers), but only if the tokens could be immediately exchanged for food. If the delay between the choice response and food was the same for both alternatives, the pigeons chose the three-token alternative, even though the three tokens were delayed. This suggested that the tokens were serving as discriminative stimuli that indicated the number of food deliveries, but that the pigeons’ choices were controlled mainly by the delays to food, not by the delays to the tokens.

Hackenberg and Vaidya (2003) used a similar procedure to examine three different types of delays: (1) the delay from a choice response to the token delivery, (2) the delay from a choice response to the start of the exchange period, and (3) the delay from a choice response to food. They found that pigeons’ choices were affected by the latter two delays, but there was no evidence that delays to token deliveries had any effects on choice. In an experiment in which pigeons earned tokens by responding on FR schedules, Foster and Hackenberg (2004) found that if the unit price (i.e., the number of responses required per token) was held

constant, pigeons showed a strong preference for whichever alternative had the shorter delay to food. For example, if one alternative required 25 responses to earn one token that could be then exchanged for 2 s of food, and if the other alternative required 50 responses to earn two tokens that could then be exchanged for 4 s of food, the pigeons almost always chose the one-token option. In summary, all of these studies suggested that what affected the pigeons' choices was not how quickly the tokens were delivered, but how quickly they could be exchanged for food.

These previous studies on choice with token stimuli all used discrete-trial procedures, in which only a single brief response was needed to make each choice. In contrast, the present experiments examined the effects of token stimuli with pigeons responding on concurrent VI schedules. Although tokens or points have been used as reinforcers in concurrent VI schedules with human participants (e.g., Bradshaw, Szabadi, & Bevan, 1979; Fukui & Kimura, 1987; Ruddle, Bradshaw, Szabadi, & Bevan, 1979), we are not aware of any previous experiments using tokens in concurrent VI schedules with nonhumans.

One advantage of using concurrent VI schedules is that responding can occur continuously when the schedules are operating, so choice proportions can be measured from moment to moment as tokens are presented or removed. Another advantage is that the characteristics of these schedules have been extensively studied with traditional primary reinforcers. Numerous experiments have shown that behavior on concurrent VI schedules can be well described by the generalized matching law (Baum, 1974, 1979):

$$\frac{B_1}{B_2} = b \left(\frac{R_1}{R_2} \right)^s, \quad (1)$$

where B_1 and B_2 are the response rates of two different behaviors, R_1 and R_2 are the reinforcement rates for those two behaviors, b is a measure of bias for one alternative or the other, and the exponent s reflects sensitivity to the differences in reinforcement rates for the two alternatives. In all conditions of the present experiments, the scheduled rates of food deliveries were 60 reinforcers per hour for each response key. Therefore, in the absence of other factors, Equation 1 would predict equal rates of responding on both keys. However, the VI schedules did not deliver food reinforcers directly; rather, the VI schedules delivered only tokens, and a certain number of tokens had to be earned before they could be exchanged for food. For some pigeons, the tokens were stimulus lights, and for others they were small white circles on the response keys. Similar to the procedures used by Hackenberg and colleagues, experimental sessions were divided into token-production periods, in which the tokens were delivered according to the concurrent VI schedules, and exchange periods, in which the tokens were exchanged for food deliveries.

In Experiment 1, the number of tokens required to enter the exchange period varied between one and four in different conditions. For example, in Condition L3-R1, three tokens had to be earned on the left key before they could be exchanged for three food deliveries, but only one token had to be earned on the right key before it could be exchanged for one food delivery. To determine the effects of the tokens on choice behavior, we examined (1) the overall distributions of responding in conditions with different exchange requirements, and (2) how response proportions changed within a session depending on how many tokens had been earned for each key. Based on previous research with token stimuli, we expected that the delay to the exchange period would have a substantial effect on choice. For example, in Condition L3-R1, we hypothesized that choice percentages on the left key would be low when three more tokens had to be earned before the exchange period, but that left responses would be greater when only one more left token had to be earned. To separate the effects of the token stimuli from those of the exchange requirement, Experiment 1 also included two

conditions with the same VI schedules and exchange requirements as in other conditions, but no token stimuli were presented during the token-production periods. Once the pigeons reached the exchange period, however, the token stimuli appeared at once, and as in the other conditions each token could be exchanged for one food delivery.

Experiment 2 was conducted to examine the roles of the tokens in another way. In some conditions, the exchange requirement was two tokens for each key, and in other conditions it was three tokens for each key. However, for one key, each token appeared as soon as it was earned during the token-production period, but for the other key the tokens did not appear until the exchange period began. This arrangement allowed us to determine whether the passage of time since the last exchange period might exert stimulus control over the pigeons' choices, even in the absence of visible tokens. If the passage of time exerted only weak discriminative control as compared to the visible tokens, this should be evident through a comparison of responding on the keys with and without the visible tokens.

EXPERIMENT 1

Method

Subjects—The subjects were 12 male White Carneau pigeons (Palmetto Pigeon Plant, Sumter, South Carolina) maintained at approximately 80% of their free-feeding weights. All had previously participated in a variety of other experiments on choice, including adjusting-ratio and adjusting-delay procedures, and so no pretraining was necessary. However, none of the pigeons had previous experience with token stimuli.

Apparatus—Three experimental chambers were used, each 30 cm long, 30 cm wide, and 31 cm high. The chambers had three response keys, each 2 cm in diameter, mounted in the front wall, 24 cm above the floor and 8 cm apart, center to center. A force of approximately 0.15 N was required to operate each key. Behind each key was a 12-stimulus projector (Med Associates, St. Albans, VT) that could project different colors or shapes onto the key. 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. A shielded 2-W white houselight was mounted at the top center of the rear wall of each chamber.

In one chamber, a horizontal row of eight 2-W stimulus lights stretching across the top of the front wall served as the tokens. The lights were mounted in a 30-cm long rectangular panel that extended 4 cm from the wall and 3 cm from the ceiling. The lights were 1.2 cm in diameter and 2.8 cm apart, center to center. The four lights on the left had red translucent caps, and the four on the right had green translucent caps. In the second chamber, the token stimuli were four 0.2-cm white circles that could be projected onto each response key in a horizontal row. The number of white circles projected on a key could vary from 0 to 4, indicating the number of tokens that had been accumulated for that key. They could be projected on a background that was red, green, or black. The third chamber was the same as the second, except that the token stimuli were four 0.2-cm white circles projected onto each response key in a vertical row rather than a horizontal row. These three variations in the token stimuli were included to increase the generality of the results, and to reduce the possibility that the pigeons' performance might be influenced by some specific details related to the visual appearance or spatial arrangement of the tokens.

Each chamber was enclosed in a sound-attenuating box containing a ventilation fan. All stimuli were controlled and responses recorded by a personal computer using MED-PC software.

Procedure—The pigeons were divided into three groups of 4. The pigeons in Group 1 were tested in the chamber in which the tokens were the red and green lights, The pigeons in Group 2 were tested in the chamber with the horizontal rows of white circles on the response keys, and those in Group 3 were tested in the chamber with the vertical rows of white circles on the response keys. Experimental sessions were usually conducted 6 days a week. Groups 1 and 2 received the same 11 experimental conditions in the same order. The pigeons in Group 3 received only six of these conditions (see Table 1).

Group 1. All sessions lasted for 60 min, and alternated between token-production periods, in which tokens were earned on two concurrent VI schedules, and exchange periods, in which the tokens were exchanged for food deliveries. During the token-production periods, the left key was lit red, the right key was lit green, and the white houselights were on. Pecks on the red or green keys produced token deliveries on concurrent VI 60-s VI 60-s schedules. The VI schedules used dependent scheduling; that is, a single VI 30-s schedule assigned tokens to either the left or right key with equal probability, with the restriction that 8 of every 16 tokens were assigned to each key. Once a token was assigned to a key, the VI schedule stopped and no further tokens could be assigned to either key until the token was collected. There was also a 2-s changeover delay.

When the first token was earned for a red-key response, the following events occurred. The red and green keylights and the houselight were turned off. The leftmost red token light flashed on and off for 1 s (0.1 s on, 0.1 s off), and then it remained lit. Similarly, when the first token was earned for a green-key response, the red and green keylights and the houselight were turned off, and the rightmost green token light flashed on and off for 1 s (0.1 s on, 0.1 s off), and then it remained lit. If additional tokens were required before the exchange could occur, the red and green keylights and the houselight were again lit, and the token-production period resumed; otherwise, the exchange period began, as described below. As each additional token was earned, another token light was lit (from left to right for the red token lights, and from right to left for the green token lights). The delivery of each new token was signaled in the same way as the first—the keylights and houselight were turned off, and the new token light flashed on and off for 1 s, then remained lit.

The number of tokens that had to be earned on a particular key before they could be exchanged for food will be called the *exchange requirement*. The condition names in Table 1 refer to the exchange requirements for the two keys. For example, in Condition L3-R1, the exchange requirement was three tokens for the left key, but only one token for the right key. As Table 1 shows, in Conditions 1-7, the exchange requirement was one token for one of the keys, and it was one, two, three, or four tokens for the other key. The exchange requirement was two tokens for each key in Condition 8, and it was three tokens for each key in Condition 9.

When the exchange requirement was met for the left (red) key, the two side keys were turned off, the center key was lit red, and a red exchange period began. After a single peck on the center key, the keylight was turned off, one of the red token lights was turned off, and food was presented for 3 s. Then, if there were additional red token lights still lit, the center key was again lit red, and each peck on the center key produced another food delivery and another red token light was turned off. Once all the red tokens were exchanged for food deliveries, the two side keys were lit and a new token-production period began. If any green token lights were lit when the red exchange period began, they remained on throughout the exchange period, and were carried over to the next token-production period (where they counted toward the next green exchange requirement).

Similarly, when the exchange requirement was met for the right (green) key, the two side keys were turned off, the center key was lit green, and a green exchange period began. The green exchange procedure followed the same sequence as the red exchange procedure: Each peck on the green center key turned off one of the green token lights and produced a 3-s food presentation, and this cycle was repeated until all the green token lights were turned off. Any red token lights that were lit remained on and were carried over to the next token-production period.

Conditions 10 and 11 used the same procedures as the other conditions, with one important difference: Although the tokens were earned according to the same concurrent VI schedules, they were not visible to the pigeons during the token-production period, but only during the exchange period. For example, in Condition 10 (labeled L3N-R3N), the exchange requirement was three tokens for each key, but there was no stimulus change when the first two tokens were earned on a key. The red and green keylights remained on throughout the token-production period until three tokens had been earned on one key, then the three token lights were lit at once above that key, and the exchange period began. The exchange procedure was the same as in the other conditions: The side keys were dark, the center key was either red or green, and each peck on the center key turned off one of the token lights and led to 3-s food delivery. The purpose of Conditions 10 and 11 was to examine the effects of the tokens by comparing these conditions to Conditions 9 and 6, respectively, which had identical contingencies for food deliveries, but where the tokens were visible during the token-production periods.

The first condition of the experiment lasted for a minimum of 20 sessions, and all other conditions lasted for a minimum of 15 sessions. For each session the mean left-key response percentage was calculated. After the minimum number of sessions, a condition was terminated for each pigeon individually when the following criteria were met: (a) Neither the highest nor the lowest single-session response percentage could occur in the last six sessions of a condition. (b) The mean response percentage across the last six sessions could not be the highest or the lowest six-session mean of the condition. (c) The mean response percentage of the last six sessions could not differ from the mean of the preceding six sessions by more than 5%.

Group 2. For this group, the procedures during both the token-production periods and the exchange periods were similar to those of Group 1, but the stimulus conditions were somewhat different because the token stimuli were white circles presented on the response keys rather than lights above the keys. During the token-production periods, the left key had a green background color and the right key has a red background color. When the first token was earned for a left response, the following events occurred. The red background color on the right key and the white houselight were turned off. A 2-cm white circle appeared on the left side of the left key, and this circle flashed on and off for 1 s (0.1 s on, 0.1 s off), then remained on. Similarly, when the first token was earned for a right-key response, the green background color on the left key and the houselight were turned off, and a 2-cm white circle appeared on the left side of the right key, flashed on and off for 1 s (0.1 s on, 0.1 s off), then remained on. If additional tokens were required to meet the exchange requirement, the red or green background color on the opposite key and the houselight were again lit, and the token-production period resumed; otherwise, the exchange period began, as described below. As each additional token was earned, another white circle was lit (from left to right for both keys). The delivery of each new token was signaled in the same way as the first—the background color on the opposite key and the houselight were turned off, and the new white circle flashed on and off for 1 s, then remained on.

When the exchange requirement was met for the left key, all lights in the chamber were turned off except for the accumulated token circles on the left key (which were now on a black background) and the left-key exchange period began. After a single peck on the left key, the token circles were turned off, and food was presented for 3 s. Then, if there were additional tokens to be exchanged for food on the left key, the remaining token circles reappeared on the key (one fewer token circle than before the food delivery), and each peck on the left key produced another food delivery, and another token was turned off. Once all the left-key tokens were exchanged for food deliveries, the two side keys were again lit (red and green, respectively) and a new token-production period began. If any token circles had accumulated on the right key before the left exchange period began, these token circles reappeared on the right key, where they counted toward the next right-key exchange requirement.

When the exchange requirement was met for the right key, all lights in the chamber were turned off except for the accumulated token circles on the right key (which were now on a black background) and the right-key exchange period began. The sequence of events was the same as in the left-key exchange period—each peck on the right key led to a 3-s food delivery and one token circle was turned off until all the right-key tokens had been exchanged for food.

As Table 1 shows, Group 2 received the same 11 conditions as Group 1. (One pigeon, P23, died before Condition 10, so there are only data from the first 9 conditions for this pigeon.) In Conditions 10 and 11, no token circles were visible during the token-production periods. Although the tokens were not visible, the same requirements for earning tokens were in effect as in Conditions 9 and 6, respectively. When the exchange requirement was met for one key, the corresponding number of token circles appeared on that key, and an exchange period began. The procedures and stimuli during the exchange periods were the same as in the other conditions.

Group 3. The procedure for this group was exactly like that for Group 2 except that the token circles appeared in a vertical row rather than a horizontal row on the left and right keys, appearing from bottom to top on the keys as successive tokens were earned. As Table 2 shows, this group received only six of the conditions.

Results and Discussion

All analyses were based on the data from the six sessions of each condition that satisfied the stability criteria. For each of the three groups of pigeons, Figure 1 shows the overall percentage of left responses for the conditions in which one or both of the exchange requirements was one token (the first seven conditions for Groups 1 and 2, and the first four conditions for Group 3). Although there was considerable variability both among subjects and across conditions, Figure 1 shows a tendency for the pigeons to make more responses on whichever key had the smaller exchange requirement. When the left key had an exchange requirement of four tokens (Condition L4-R1), the pigeons made an average of 41% of their responses on the left key. In the opposite condition, where the right key had an exchange requirement of four tokens (Condition L1-R4), the percentage of left responses was 61%. These results bear some similarity to those obtained with FR schedules by Foster and Hackenberg (2004), who found that when the unit price was the same for the two alternatives, pigeons showed a preference for the key with the shorter delay to the exchange period. In the present experiment, that key was, of course, the one with the smaller exchange requirement. However, the modest trend in the overall response percentages shown in Figure 1 provides no information about the systematic changes that occurred as successive tokens were earned in each condition.

To examine how the pigeons' choices were affected by the number of token stimuli present at any given moment, response percentages were calculated separately for each possible configuration of tokens in each condition. Figure 2 shows the results for individual pigeons from all the conditions that had a one-token exchange requirement for one or both keys. The means for all 12 pigeons are presented in the top panel of Figure 3. For all cases where the left and right exchange requirements were counterbalanced across conditions, the results were averaged across the two conditions. For example, Groups 1 and 2 received both Conditions L3-R1 and L1-R3, and the results from these two conditions were averaged (and labeled Ex 3-1 in Figures 2 and 3). These figures show the choice percentages for the key with the larger exchange requirement, plotted as a function of the number of tokens that had been earned at any moment. These percentages are based on the number of responses made since the previous token was received. For example, in the Ex 4-1 conditions, there were only four possible token stimulus configurations: zero, one, two, or three token stimuli could be on for the key with the larger exchange requirement. Note that for the key with the one-token exchange requirement, no token stimuli were ever on during the token-production phase, because the exchange period began as soon as one token was earned.

The clearest trend in Figures 2 and 3 is that response percentages for the key with the larger exchange requirement increased steadily as more tokens were earned for that key. As an example, the response percentages for Pigeon 21 (top center panel in Figure 2) in the Ex 4-1 conditions was 31% when no tokens had been earned, but rose to 73% when three tokens had been earned. Figure 2 shows that this pattern of increasing response percentages as more tokens were earned for the key with the larger exchange requirement was found for all 12 pigeons in the Ex 4-1, Ex 3-1 and Ex 2-1 conditions. In most cases, response percentages for the key with the larger exchange requirement were less than 50% when no tokens had been earned, but they gradually rose above 50% as more tokens were earned.

This pattern can be explained as a result of a trade-off between reinforcer delays and amounts. At times when three or four tokens still had to be earned before the exchange period, this meant that the average delay to the exchange period was three or four times larger than for the key with a one-token exchange requirement. This longer delay evidently outweighed the fact that there would be three or four reinforcer deliveries when the exchange period was reached. However, when only one more token had to be earned (on either key) before an exchange period would begin, the pigeons showed a preference for the key that would deliver three or four reinforcers rather than one. Therefore, these shifts in preference are similar to the preference reversals that have been found in many other studies when animals must choose between smaller but more immediate and larger but more delayed reinforcers (e.g., Green, Fisher, Perlow, & Sherman, 1981; Green & Snyderman, 1980; Snyderman, 1983).

A comparison of the results among the Ex 4-1, Ex 3-1, and Ex 2-1 conditions showed less consistency, especially for the individual subjects. Despite this variability, a trend that can be seen in the group means in Figure 3 is worth mentioning. To make this point more clearly, the bottom panel in Figure 3 plots the same means in a different way: The x-axis shows the number tokens still needed before the next exchange period. The advantage of plotting the data in this way is that it allows a direct comparison across conditions at times when the pigeons were equally far from the next exchange period. When the same number of tokens still had to be earned, response percentages tended to be highest in the Ex 4-1 conditions, and progressively lower in the Ex 3-1 and Ex 2-1 conditions. One might reasonably expect a higher response percentage when four tokens were delivered in the exchange periods rather than three or two tokens. Indeed, the lack of a clearer difference among the Ex 4-1, Ex 3-1 and Ex 2-1 conditions might seem a bit surprising, and it suggests that the number of tokens delivered in the exchange period had less control over responding

than did the delay to the start of the exchange period (as measured by the number of tokens still to be earned).

The results from the conditions with exchange requirements of two or three tokens for each key are more complex because there were more possible combinations of tokens earned. For instance, in Condition L3-R3, which had an exchange requirement of three tokens for each key, there were nine possible token stimulus configurations during the token-production periods: At different times, zero, one, or two tokens might be visible for either key. For each pigeon, Figure 4 shows the left-key choice percentages for all nine token configurations. The filled circles are from times when no right tokens had been earned, and for all 12 pigeons, left response percentages increased as more tokens were earned on the left key. Similar increases in left-key responding were observed when one token (open circles) or two tokens (triangles) had been earned on the right key. However, left response percentages were also systematically related to the number of right-key tokens that had been earned, as can be seen from the fact that the filled circles are consistently above the open circles, which are in turn consistently above the triangles. Figure 4 shows that response percentages were jointly determined by the number of tokens earned on each key. Stated simply, response percentages tended to be higher for whichever key was closer to meeting the exchange requirement at any given moment.

Similar results were found in Condition L2-R2, except that there were only four possible token configurations: At different times, zero or one token might be visible for either key. Figure 5 shows the results from this condition, plotted in the same way as in Figure 4. Here too, response percentages were jointly determined by the number of tokens earned on each key: Left response percentages were higher when one more token had been earned on the left key, and they were also higher when no tokens had been earned on the right key.

These analyses appear to show that choice percentages were strongly controlled by the number of tokens earned at any moment, but it is also possible that responding was controlled by the time since the last exchange period. For example, because four tokens had to be earned on the left key in Condition L4-R1, a left exchange period occurred about every 4 min, on average. It is possible that responding on the left key was low immediately after a left exchange period and then increased as time passed, regardless of the number of left tokens that were earned. The two final conditions of this experiment were conducted to separate the effects of the token stimuli from the effects of time by eliminating the display of tokens during the concurrent VI schedules. Figure 6 shows the results from the 7 pigeons that received Condition L4N-R1N, plotted as a function of the number of tokens earned (although the tokens were not visible to the pigeons until the exchange period). For most of the pigeons, the left response percentage was lower when no tokens had been earned on the left key, which was immediately after a left-key exchange period had occurred. This implies that the left-key exchange period (which included four food deliveries) served as a stimulus for a decrease in responding on the left key, in much the same way as food deliveries lead to postreinforcement pauses on simple FR and FI schedules. However, there were no further changes in the response percentages as one, two or three tokens were earned. This pattern contrasts with the steady increase in choice percentages when the tokens were visible, as shown in the data for Condition L4-R1 in Figures 2 and 3, and it shows that these increases were due to the visibility of the tokens and not the passage of time since the last exchange period.

Figure 7 shows the left response percentages for the 7 pigeons that received Condition L3N-R3N, plotted in the same way as the results from Condition L3-N3 in Figure 4. The difference between these two figures is striking. There were obvious changes in the choice percentages as each token was earned (and visible) in Condition L3-R3, but no systematic

changes when the tokens were earned (but not visible) in Condition L3N-R3N. These results again demonstrate that choice percentages were controlled by the presence of the token stimuli and not by the passage of time.

In summary, Experiment 1 demonstrated the role that the tokens played as discriminative stimuli by showing how choice percentages changed as each token was earned, and how most of the differences were eliminated when the tokens were not visible for either key in the last two conditions. Experiment 2 examined choice behavior when the tokens were visible for one key but not for the other key during the token-production periods. Based on the idea that the token stimuli are conditioned reinforcers, one might predict that the pigeons would show a preference for whichever key provided visible tokens. If, however, the tokens are simply discriminative stimuli for upcoming food, it is not obvious what effect their presence versus absence might have on choice. There might be no clear preference for either alternative, or possibly even a preference for the key without visible tokens. Experiment 2 was therefore conducted to see what type of preference, if any, would be found when one alternative had visible tokens and the other did not, and to try to determine the reasons for any observed preferences.

EXPERIMENT 2

Method

Subjects and Apparatus—The 8 pigeons from Groups 1 and 3 of Experiment 1 were used, and they were tested in the same chambers.

Procedure—For both groups, the general procedure was the same as in Experiment 1, except that during the token-production periods, the token stimuli were visible for one of the two keys but not for the other key. As in Experiment 1, the token stimuli for both keys were visible during the exchange periods. Table 2 shows the order of conditions for both groups. In some conditions, the exchange requirement was three tokens for each key, and in other conditions it was two tokens for each key. For example, in Condition L3-R3N, the exchange requirement was three tokens for each key, and “N” indicates that the tokens for the right key were not visible during the token-production periods. The key for which the tokens were not visible alternated between right to left in succession conditions. Each condition was conducted for a minimum of 15 sessions, and the criteria for terminating a condition were the same as in Experiment 1.

Results and Discussion

All analyses were based on the data from the six sessions of each condition that satisfied the stability criteria. For both groups, Figure 8 shows the mean left response percentages in each of the six conditions. As the side without the tokens switched during successive conditions, the pigeons' response percentages shifted accordingly: They displayed a preference for whichever key had no tokens visible during the token-production periods. This result was found with both the three-token and two-token exchange requirements. The results from individual pigeons showed this same pattern: For every pigeon and every pair of conditions, the left-key response percentages were lower when tokens were visible for the left key than when they were visible for the right key. A repeated-measures analysis of variance (ANOVA) found that there was significantly more responding on whichever key had no tokens visible, $F(1,7) = 35.02, p < .001$.

The reasons for this preference for the alternative with no tokens visible can be understood by examining how left response percentages changed as each token was earned. Figure 9 shows the group means of these choice percentages for the three-token and two-token

conditions of this experiment, and compares them to the corresponding conditions from Experiment 1 (Conditions L3-R3 and L2-R2), in which the tokens were visible for both keys. One conclusion that can be drawn from Figure 9 is that, as in Experiment 1, choice percentages were strongly controlled by the number of tokens visible, and the passage of time had little or no effect. The top row of panels (from the conditions of Experiment 1) shows that choice percentages varied as each additional token appeared, for both keys. The middle row in Figure 9 shows that when the right-key tokens were not visible, the number of right-key tokens earned had little or no effect. (Whatever small effects remained must have been due to the passage of time since the last right-key exchange period.) Conversely, the bottom row in Figure 9 shows that when the left-key tokens were not visible, the number of left-key tokens earned had little or no effect.

The reason that the pigeons' overall choice percentages favored the key without tokens can be explained as follows. Consider the results in the left middle panel of Figure 9, which are from the conditions in which only the left-key tokens were visible. Because no tokens were ever visible for the right key, the absence of these right-key tokens provided no indication of how much time remained before the next exchange period. However, the absence of left-key tokens did provide such information—three more tokens had to be earned on the left key before the next left exchange period. As a result, the left response percentage was only about 15% when no left-key tokens were visible. When one left token was visible, this indicated that two left tokens still had to be earned, and the left response percentage was 36%. When two left-key tokens were visible, this indicated that only one left-key token still had to be earned, and the choice percentage finally shifted to the left key (63%). Because the left response percentage only rose above 50% when two left-key tokens were visible, the overall left response percentage in this condition averaged 38%.

The opposite pattern can be seen in the bottom left panel in Figure 9. No left-key tokens were visible in this condition, but the pigeons exhibited a preference for the left key both when no right-key tokens were visible (filled circles) and when one right-key token was visible (open circles). Only when two right-key tokens were visible (triangles) did preference shift to the right key, and the overall left response percentage in this condition was 60%. In short, the overall response percentages that are shown in Figure 8 should not be construed as evidence that the pigeons had a stable and enduring preference for whichever alternative did not have visible tokens. Rather, these overall percentages are simply aggregates across some times when there was more responding on the side without visible tokens, and other times when there was more responding on the side with visible tokens. The same reasoning can be applied to the results from the conditions with two-token exchange requirements, as shown in the right panels of Figure 9. This analysis of the choice patterns with different numbers of visible tokens shows that whether the pigeons displayed a preference for or against the key with visible tokens depended on how much time remained before the next exchange period. It was therefore the temporal proximity to food, not the visibility of the tokens that controlled their preferences.

GENERAL DISCUSSION

Experiments by Hackenberg and his colleagues have found that stimulus lights can serve as effective tokens for pigeons. The stimulus lights exerted strong control over behavior both on single schedules (e.g., Bullock & Hackenberg, 2006; Raiff, Bullock, & Hackenberg, 2008) and in choice situations (e.g., Hackenberg & Vaidya, 2003; Jackson & Hackenberg, 1996). The present experiments extended this research by examining the effects of tokens on pigeons' behavior in a concurrent VI choice situation. Three different types of token stimuli were used for the three groups of pigeons—a row of stimulus lights near the ceiling of the chamber for Group 1, a horizontal row of white circles on the response keys for Group 2,

and a vertical row of white circles on the response keys for Group 3. All three types of token stimuli produced effective stimulus control, and there were no systematic differences in performance among the three groups. The effectiveness of each of these different types of tokens as discriminative stimuli can be seen by comparing the results from conditions with and without tokens visible (e.g., Figure 4 vs. Figure 7).

As described in the Introduction, previous research has suggested that token stimuli, and conditioned reinforcers in general, can serve a variety of functions. Token stimuli can serve as replacements for primary reinforcers, maintaining existing behavior patterns or establishing new ones (e.g., Kelleher, 1956; Malagodi, 1967; Wolfe, 1936). They can also have marking, bridging, or eliciting functions (Hackenberg, 2009). However, in the two experiments reported here, the main finding was that the tokens served primarily as discriminative stimuli. The pigeons' choice responses were strongly controlled by the number of tokens that were visible at any particular time.

It should be emphasized, however, that it was not simply the absolute number of tokens present but the number relative to the exchange requirements that determined the pigeons' choice percentages. For instance, suppose that one left token and no right tokens were currently visible: How did the pigeons distribute their choices? In Condition L2-R1, the mean left-key choice percentage with this configuration of tokens was 61%. In Condition L4-R1, the mean choice percentage with this same configuration of tokens was 29%. This difference in behavior reflects the fact that a single visible token signaled two different situations in these two conditions. In Condition L2-R1, it indicated that only one more token had to be earned on the left key before an exchange period, but in Condition L4-R1, three more tokens had to be earned on the left key. Therefore, the same token configuration signaled two different delays to food in these two conditions, and the pigeons' choice percentages reflected this difference. To put it simply, choice percentages were not controlled by the absolute number of tokens earned and visible at any moment, but rather by the current delay to food, as measured by the number of tokens still needed before the next exchange period.

We conducted a series of calculations to determine whether two basic principles of choice—the matching law and delay discounting—could together account for the patterns of behavior observed in these two experiments. Our starting point was the basic matching relation, except that we assumed that the response ratios were equal, not to the reinforcement ratios, but to the ratio of the values of the two alternatives, V_1 and V_2 , as has been suggested by others (e.g., Killeen, 1972):

$$\frac{B_1}{B_2} = \frac{V_1}{V_2}. \quad (2)$$

To calculate V_1 and V_2 , we used the following hyperbolic-decay equation (Mazur, 1984):

$$V = \sum_{i=1}^n P_i \left(\frac{A}{(1+K D_i)} \right), \quad (3)$$

where V is the value of a reinforcer that could be delivered after any one of n possible delays, A represents the reinforcer amount, and P_i is the probability that there will be a delay of D_i seconds. K is a free parameter that reflects the rate at which a reinforcer's value decreases with increasing delay. Equation 3 states that the value of a reinforcer that is delivered after a variable delay can be calculated by averaging the values derived from all the different possible delays, each weighted by its probability of occurrence. For simplicity, in these calculations the parameter K was set equal to 1. The units used to measure A are

arbitrary because they cancel out when calculating response percentages, but it was assumed, again for simplicity, that A was proportional to the number of food deliveries in an exchange period (i.e., for exchange periods delivering one, two, three, and four tokens, values of 100, 200, 300, and 400 were used for A in Equation 3).

Condition L1-R2, with exchange requirements of one token for the left key and two tokens for the right key, will be used to illustrate how V_1 and V_2 were calculated. With the VI 60-s schedules used in these experiments, each second there was a $1/60$ chance that a token would become available, to be collected with the next response. The probability of earning a token on the left key after 1 s is therefore $1/60$. From simple rules of probability, it follows that the probability of earning a token after 2 s is $1/60 * (1 - 1/60)$, and more generally, that the probability of earning a token after t s is $1/60 * (1 - 1/60)^{(t-1)}$. Individual values of P_i and D_i were used in Equation 3 to calculate the total value of the left key—a VI 60-s schedule with a one-token exchange requirement.¹ The calculations also included an additional delay of 2 s to account for the 1-s flashing of the token stimulus plus 1 s to approximate the time needed to make a single peck on the exchange schedule. With these assumptions, and with $A_1 = 100$ for a one-token exchange requirement, Equation 3 yields a value of 4.15 for the left key. Now suppose that one token is currently visible for the right key, indicating that one more token needs to be earned before the exchange period. In this situation, the delays and probabilities are the same as for the left key, but $A_2 = 200$ because two tokens will be delivered for the right key, so the value of the right key is currently 8.30 (double that of the left key). Therefore, Equation 2 predicts a 1:2 response ratio, which corresponds to a left-key response percentage of 33%.

Next, suppose that no token lights are lit in Condition L1-R2, which indicates that two tokens still need to be earned on the right key. Because this stimulus configuration signals a greater expected delay to food for the right key, Equation 3 was used to calculate a different value of V_2 . The same methods were used as described above, except that P_i is now the probability that *two* tokens will be obtained after a total delay of D_i seconds. In other words, this approach assumes that the values of the schedules are based on the delays to food (not to the tokens). For this case, Equation 3 yields a value of 2.88 for V_2 . The value of V_1 is the same as before, so Equation 2 predicts a response ratio $4.15/2.88$, which corresponds to a left-key response percentage of 59%. In summary, this model predicts that the left-key response percentage will shift from 59% when no tokens are present to 33% when one right-key token is present.

Using these techniques, separate values of V_1 and V_2 , and left-key response percentages were calculated for each configuration of tokens that could occur in each different condition. Finally, for the conditions in which the tokens were not visible for one key, we assumed that the pigeons did not know how many tokens had been earned at any time, so a single value of V was calculated from the average of the values with different numbers of earned tokens. For example, in Condition L2N-R2, the values of V_2 were 2.88 with no right tokens visible and 8.30 with one right token visible (as explained in the preceding example), but because no tokens were ever visible for the left key, V_1 was kept constant at 5.59 (the average of 2.88 and 8.30).

Figures 10 and 11 show the results of these calculations. Figure 10 shows the predictions for the conditions with an exchange ratio of one token for one of the keys, plotted in the same

¹To obtain estimates of V for the different schedules, the probability of earning a token was calculated for every 1-s interval for the first 60 s of each VI schedule, and for every 10-s interval thereafter, up to a maximum of 1000 s (at which point the cumulative probability of a token delivery was virtually 100%). These approximations were deemed accurate enough for the present purposes, because no pigeon would experience all the theoretically possible token delays in a finite number of sessions, and because the purpose of the calculations was to illustrate the general predictions of the equations, not to obtain precise quantitative fits to the data.

way as the data shown in the top panel of Figure 3. Although the predictions are not always similar to the data in terms of exact response percentages, they do capture the main trends shown in Figure 3—the increasing response percentages as more tokens were earned in all conditions, and the differences among the conditions with different exchange ratios.

Figure 11 shows the predictions for the conditions with exchange requirements of two or three tokens on both keys, plotted in the same way as the data in Figure 9. For the conditions with tokens visible for both keys (top row), the effects of the different numbers of earned tokens shown in Figure 9 are nicely approximated by the predictions. For the conditions in which tokens were not visible (middle and bottom rows in Figures 9 and 11), the predictions are, of course, identical regardless of the number of tokens earned. In the actual data, there were some small differences based on the number of tokens earned (even though they were not visible), suggesting that the pigeons' choices were controlled to at least a small degree by the passage of time since the last exchange period. However, the predictions of the effects of the visible tokens were similar to the actual results. For example, the bottom left panel of Figure 11 shows that when the exchange requirements were three tokens and the left tokens were not visible, the model predicts a preference for the left key when zero or one right tokens were visible, but a preference for the right key when two right tokens were visible. These predictions were consistent with the pigeons' actual choices, as were the predictions from the other conditions with tokens visible for only one of the keys. As a result, the model also predicts that, when averaged across an entire session, the overall response percentages in these conditions will always favor the side without the visible tokens. This is what was found in Experiment 2, both with two-token and three-token exchange requirements.

The schedules for earning tokens in these experiments were similar in some ways to chain schedules, in which a sequence of two or more schedules, each with a unique discriminative stimulus, must be completed, and food is delivered only after the final schedule is completed. With chain schedules, it is typically found that responding comes under discriminative control of the stimuli associated with each schedule in the chain, and that responding is slowest in the early links, and increased in the links closer to primary reinforcement (e.g., Gollub, 1958; Kelleher & Golub, 1962). As we have shown, the token stimuli in these experiments exhibited stimulus control over choice behavior in similar way. However, this does not imply that token stimuli are identical in all ways to discriminative stimuli in chain schedules. As discussed in the introduction, many previous studies have shown that token stimuli can serve several other functions (e.g., marking, bridging, and serving as conditioned reinforcers).

It should be noted that these predictions were obtained using a model with no free parameters. Closer quantitative fits to the data could almost surely be obtained by varying K as a free parameter or by adding free parameters for bias and sensitivity in Equation 2. Our purpose, however, was not to develop a complete quantitative account of the data, but simply to demonstrate that the main findings of these experiments could be described using an approach that is based on three basic assumptions: (1) that choice in concurrent VI schedules can be described by a matching relation, (2) that delayed reinforcers decrease in value according to a hyperbolic function, and (3) that the main function of the tokens was to serve as discriminative stimuli that signaled the expected delays to food at different times. The first two assumptions have been supported by the results from many previous experiments. To the extent that the simulations succeeded in capturing the main features of the pigeons' choice behavior, we conclude that this provides support for the third assumption.

Acknowledgments

This research was supported by Grant R01MH38357 from the National Institute of Mental Health. The content of this article is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of Mental Health or the National Institutes of Health. We thank Michael Lejeune, Kimberly Rakiec, and Lisa Wanek for their help in conducting this research.

References

- Baum WM. On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*. 1974; 22:231–242. doi:10.1901/jeab.1974.22-231. [PubMed: 16811782]
- Baum WM. Matching, undermatching, and overmatching in studies of choice. *Journal of the Experimental Analysis of Behavior*. 1979; 32:269–281. doi:10.1901/jeab.1979.32-29. [PubMed: 501274]
- Boakes RA, Poli MM, Lockwood MJ, Goodall GG. A study of misbehavior: Token reinforcement in the rat. *Journal of the Experimental Analysis of Behavior*. 1978; 29:115–134. doi:10.1901/jeab.1978.29-115. [PubMed: 16812033]
- Bradshaw CM, Szabadi EE, Bevan PP. The effect of punishment on free-operant choice behavior in humans. *Journal of the Experimental Analysis of Behavior*. 1979; 31:71–81. doi:10.1901/jeab.1979.31-71. [PubMed: 16812124]
- Bullock CE, Hackenberg TD. Second-order schedules of token reinforcement with pigeons: Implications for unit price. *Journal of the Experimental Analysis of Behavior*. 2006; 85:95–106. doi:10.1901/jeab.2006.116-04. [PubMed: 16602378]
- Cowles JT. Food-tokens as incentives for learning by chimpanzees. *Comparative Psychological Monographs*. 1937; 12:1–96.
- Foster TA, Hackenberg TD. Unit price and choice in a token reinforcement context. *Journal of the Experimental Analysis of Behavior*. 2004; 81:5–25. doi:10.1901/jeab.2004.81-5. [PubMed: 15113130]
- Fukui I, Kimura H. An experimental analysis on the relation between reward value and choice behavior. *Japanese Journal of Psychology*. 1987; 58:109–112. doi:10.4992/jjpsy.58.109. [PubMed: 3430880]
- Gollub, LR. The chaining of fixed-interval schedules (Unpublished doctoral dissertation). Harvard University; Cambridge, MA: 1958.
- Green L, Fisher EB, Perlow S, Sherman L. Preference reversal and self control: Choice as a function of reward amount and delay. *Behaviour Analysis Letters*. 1981; 1:43–51.
- Green L, Snyderman M. Choice between rewards differing in amount and delay: Toward a choice model of self control. *Journal of the Experimental Analysis of Behavior*. 1980; 34:135–147. doi:10.1901/jeab.1980.34-135. [PubMed: 16812184]
- Hackenberg TR. Token reinforcement: A review and analysis. *Journal of the Experimental Analysis of Behavior*. 2009; 91:257–286. doi:10.1901/jeab.2009.91-257. [PubMed: 19794838]
- Hackenberg TD, Vaidya M. Determinants of pigeons' choices in token-based self-control procedures. *Journal of the Experimental Analysis of Behavior*. 2003; 79:207–218. doi:10.1901/jeab.2003.79-207. [PubMed: 12822687]
- Jackson K, Hackenberg TD. Token reinforcement, choice, and self-control in pigeons. *Journal of the Experimental Analysis of Behavior*. 1996; 66:29–49. doi:10.1901/jeab.1996.66-29. [PubMed: 8755699]
- Kelleher RT. Intermittent conditioned reinforcement in chimpanzees. *Science*. 1956; 124:679–680. doi:10.1126/science.124.3224.679. [PubMed: 17800891]
- Kelleher RT. Conditioned reinforcement in chimpanzees. *Journal of Comparative and Physiological Psychology*. 1957; 50:571–575. doi:10.1037/h0044327. [PubMed: 13502499]
- Kelleher RT. Fixed-ratio schedules of conditioned reinforcement with chimpanzees. *Journal of the Experimental Analysis of Behavior*. 1958; 1:281–289. doi:10.1901/jeab.1958.1-281. [PubMed: 16811227]

- Kelleher RT, Gollub LR. A review of positive conditioned reinforcement. *Journal of the Experimental Analysis of Behavior*. 1962; 5:543–597. doi: 10.1901/jeab.1962.5-s543. [PubMed: 14031747]
- Killeen P. The matching law. *Journal of the Experimental Analysis of Behavior*. 1972; 17:489–495. doi:10.1901/jeab.1972.17-489. [PubMed: 16811604]
- Malagodi EF. Acquisition of the token-reward habit in the rat. *Psychological Reports*. 1967; 20:1335–1342. [PubMed: 5625463]
- Mazur JE. Tests of an equivalence rule for fixed and variable reinforcer delays. *Journal of Experimental Psychology: Animal Behavior Processes*. 1984; 10:426–436. doi: 10.1037/0097-7403.10.4.426.
- Raiff BR, Bullock CE, Hackenberg TD. Response-cost punishment with pigeons: Further evidence of response suppression via token loss. *Learning & Behavior*. 2008; 36:29–41. doi:10.3758/LB.36.1.29. [PubMed: 18318424]
- Ruddle HH, Bradshaw CM, Szabadi EE, Bevan PP. Behaviour of humans in concurrent schedules programmed on spatially separated operanda. *The Quarterly Journal of Experimental Psychology*. 1979; 31:509–517. doi:10.1080/14640747908400742.
- Snyderman M. Delay and amount of reward in a concurrent chain. *Journal of the Experimental Analysis of Behavior*. 1983; 39:437–447. doi:10.1901/jeab.1983.39-437. [PubMed: 16812328]
- Williams BA. Conditioned reinforcement: Neglected or outmoded explanatory construct? *Psychonomic Bulletin & Review*. 1994; 1:457–475. doi:10.3758/BF03210950. [PubMed: 24203554]
- Wolfe JB. Effectiveness of token rewards for chimpanzees. *Comparative Psychological Monographs*. 1936; 12:1–72.
- Yankelevitz RL, Bullock CE, Hackenberg TD. Reinforcer accumulation in a token-reinforcement context with pigeons. *Journal of the Experimental Analysis of Behavior*. 2008; 90:283–299. doi: 10.1901/jeab.2008.90-283. [PubMed: 19070337]

Fig. 1.

For each group in Experiment 1, the mean percentages of responses on the left key are shown for the conditions in which at least one of the two keys had a one-token exchange requirement.

Fig. 2.

For each pigeon in Experiment 1, the percentages of responses on the key with the larger exchange requirement are shown for all conditions in which the other key had a one-token exchange requirement. These percentages are plotted as a function of the number of tokens that had been earned for the larger exchange schedule.

Fig. 3.

For Experiment 1, group means of the response percentages on the key with the larger exchange requirement are shown for all conditions in which the other key had a one-token exchange requirement. The top panel shows these means in the same form as Figure 2, as a function of the number of tokens earned. The bottom panel shows the same data plotted as a function of the number of tokens that still had to be earned for the larger exchange schedule.

Fig. 4. For Condition L3-R3 in Experiment 1, the percentages of responses on the left key are shown for each pigeon, plotted as a function of the number of left and right tokens that had been earned at any moment.

Fig. 5. For Condition L2-R2 in Experiment 1, the percentages of responses on the left key are shown for each pigeon, plotted as a function of the number of left and right tokens that had been earned at any moment.

Fig. 6. Left-key response percentages are shown from Condition L4N-R1N of Experiment 1, in which no token stimuli were visible during the token-production periods, but four tokens had to be earned on the left key (and only one on the right key) before entering an exchange period.

Fig. 7.

Left-key response percentages are shown from Condition L3N-R3N of Experiment 1, in which no token stimuli were visible during the token-production periods, but three tokens had to be earned on either the left key or the right key before entering an exchange period.

Fig. 8. Mean percentages of left-key responses in each condition of Experiment 2. “N” denotes the side with no tokens visible during the token-production periods.

Fig. 9.

For the 8 pigeons in Experiment 2, mean left response percentages are shown as a function of the number of tokens earned when all tokens were visible (top row), when right tokens were not visible (middle row), and when left tokens were not visible (bottom row). The results in the top row are from Experiment 1, and they include only the data from the 8 pigeons that participated in Experiment 2.

Fig. 10.

Predictions based on Equations 2 and 3, as explained in the text. The results are plotted in the same form as the data from the pigeons in the top panel of Figure 3.

Fig.11. Predictions based on Equations 2 and 3, as explained in the text. The results are plotted in the same form as the data from the pigeons in Figure 9.

Table 1

Order of Conditions in Experiment 1.

Condition	Groups 1 and 2	Group 3
L1-R1	1	1
L2-R1	2	2
L1-R3	3	3
L3-R1	4	-
L1-R4	5	-
L4-R1	6	4
L1-R2	7	-
L2-R2	8	5
L3-R3	9	6
L3N-R3N	10	-
L4N-L1N	11	-

Note: Condition names refer to the exchange requirements for the left (L) and right (R) keys. N indicates that tokens earned were not visible during the token-production periods.

Table 2

Order of Conditions in Experiment 2.

Condition	Group 1	Group 3
L3-R3N	1	1
L3N-R3	2	2
L3-R3N	3	-
L3N-R3	4	-
L2-R2N	5	3
L2N-R2	6	4

Note: Condition names refer to the exchange requirements for the left (L) and right (R) keys. N indicates that tokens earned were not visible during the token-production periods. Group 2 did not participate in Experiment 2.