Timing processes in the reinforcement-omission effect

ROBERT C. MELLON, THERESA M. LEAK, and STEPHEN FAIRHURST
New York State Psychiatric Institute, New York, New York

and

JOHN GIBBON
*New York State Psychiatric Institute, New York, New York
and Columbia University, New York, New York*

A comparison was made of the effects of long-term exposure to fixed-interval reinforcement with unsignaled and with signaled reinforcement omissions. When successive fixed-interval cycles ended in reinforcement, subjects showed a clear "scalloped" response pattern. When reinforcement was omitted, but a brief signal was given in lieu of reinforcement, responding in the next cycle started earlier than it did after reinforcement. When reinforcement was omitted without exteroceptive cues, response rates peaked near the time of reinforcement and then declined to a flat but substantial level. The classic *reinforcement-omission effect* was observed, in that total responding was greatest after uncued omission, somewhat less after cued omission, and least after reinforcement. However, response rate plotted as a function of time showed that the uncued-omission condition had a very different function from that of the cued-omission or reinforcement conditions. A failed-discrimination account of the omission effect might accommodate the three functions if the discrimination is considered to be a temporal one. The temporal-discrimination account argues that high-rate responding reflects the accumulation of subjective time proximal to the memory of the time of reinforcement. The accumulation resets completely with food reinforcement, incompletely with cues in lieu of reinforcement, and not at all in the absence of cues.

In a fixed-interval schedule of reinforcement, rats and pigeons typically show a low response rate early in the interreinforcement interval and accelerated responding as the time of reinforcement draws near. Staddon and Innis (1966, 1969) showed that when brief time-out periods (i.e., lights out) are occasionally substituted for reinforcement, higher rates of responding are observed in the subsequent interval, due largely to a decrease in time to begin responding at the start of the interval. This *omission effect* is in some respects similar to the *frustration effect* (Amsel & Roussel, 1952), whereby rats trained to pick up food from the goalboxes in a double runway show an increased running speed to the second goalbox when reinforcement is omitted from the first. In the runway preparation, more vigorous responding following omission has been explained as reflecting a change in motivation experienced when an expected reward does not appear (Amsel, 1958). In the fixed-interval literature, the omission effect has been understood in terms of a failure of discriminative control after omission. Specifically, post-

reinforcement pausing is under the discriminative control of reinforcement itself, but is only partially controlled by the cues that are presented in lieu of reinforcement, and the short latency and high rate in the interval subsequent to omission reflect a greater functional continuity with the previous interval (Staddon, 1967).

A common feature of the frustration and discrimination accounts of reinforcement omission is that the subjects are assumed to respond earlier and with more vigor because they anticipate reinforcement, which then fails to appear. However, the frustration account holds that nonreinforcement itself induces a higher rate, while the discrimination account holds that nonreinforcement fails to fully inhibit the high rate associated with the end of the previous nonreinforced interval. Thus, both accounts appeal to a change in responsiveness early in the postomission interval.

Support for the discrimination account is taken from two findings. First, when time-out duration is varied, the time to first response after omission is directly related to the duration of the preceding blackout (Staddon & Innis, 1969). Presumably, longer time-outs are more likely to be attended to, disrupting the functional continuity of the pre-omission and postomission discriminative stimuli. However, the differential effects of time-out durations on post-time-out pausing are short-lived in rats; after 10 sessions, all time-out durations produced comparably shorter pausing. Moreover, the degree of pausing lengthened somewhat with continued exposure to the time-out procedure.

This research was supported by National Institutes of Health Grant MH41649 to the Research Foundation for Mental Hygiene (John Gibbon, Principle Investigator). R.C.M. is now at the University of Maryland. Requests for reprints may be addressed to J. Gibbon, N. Y. S. Psychiatric Institute, Department of Biopsychology, Unit 50, 722 West 168th Street, New York, NY 10032.

—Accepted by previous editor, Vincent M. LoLordo

The second finding is that total responding is higher when reinforcement is omitted but entirely uncued (i.e., 2-min fixed interval interspersed with the 4-min fixed intervals) than it is when there is an omission signaled by a time-out. According to Staddon and Innis (1969), this implies that the time-out has a mild inhibitory aftereffect. This observation also runs counter to a frustration account, since the more closely an omission resembles reinforcement, the more frustrating that omission should be, and, hence, the *greater* the corresponding frustration effect.

This aspect of the frustration and discrimination accounts of the effects of reinforcement omission in fixed-interval schedules was tested directly in a study reported by Kello (1972). In a free-operant situation, it is possible to withhold all or part of the stimulus events that normally accompany food presentation, including blackout, the onset of the magazine light, and the banging of the food hopper. If the frustration account is correct, frustration-based responding should increase with the number of food cues that accompany food omission (Amsel, 1958; Kello, 1972). If, on the other hand, the discrimination account is correct, post-cue *pausing* should increase with the number of food cues accompanying food omission. Since an uncued omission least resembles reinforcement, responding in this condition should be greatest. Thus, the two accounts make opposite predictions about the effects of partially cued reinforcement omissions.

In the Kello (1972) experiment, pigeons' keypecks were reinforced on a fixed-interval 2-min schedule for seven training sessions, followed by reinforcement-omission training for another seven sessions. The subjects were exposed to 60 fixed intervals daily, with reinforcement randomly omitted from half of the odd-numbered intervals (i.e., from 25% of all intervals). Omission trials were of three types, also randomly selected: (1) unsignaled; (2) signaled by blackout; and (3) signaled by blackout and magazine light. Total responding was lowest in intervals preceded by reinforcement, higher after reinforcement was omitted but cued with a blackout and magazine light, and higher still when omission was cued only by a blackout. In this experiment, only two birds were used, and one bird showed its highest responding after uncued omissions, while the other showed the same response rate for uncued omission and omission cued by blackout. An examination of cumulative records revealed that differences in the total number of responses were due largely to differences in the length of the postreinforcement (or post-nonreinforcement) pauses. The length of pausing before the first peck was consistently greatest following reinforcement, growing shorter as reinforcement cues were degraded. Pausing was negligible after uncued omission.

The discrimination account argued that cued and uncued omission simply reflected different degrees of similarity to reinforcement. However, in cued omissions, as in reinforcement, there is an external cue indicating the precise time at which reinforcement may be delivered, whereas in uncued omission, the only "cue" accessible to the animals consists of their temporal estimates of the expected time of reinforcement. When uncued omissions

are interspersed with reinforcement, the situation is very similar to a peak procedure in which reinforcement is occasionally withheld (Catania, 1970; Roberts, 1981). With sufficient training on a peak procedure, animals show a peak in responding around the time of expected reinforcement, and a decline thereafter. With little omission training, such as in the Kello (1972) and Staddon & Innis (1966, 1969) experiments, this temporal discrimination is unlikely to have developed to the point at which a decline in responding beyond the expected time of reinforcement might be observed. As Kello's cumulative records indicate, there is no discernible pause, but, rather, animals simply respond straight through the omission. We will show that the increased responding seen after omissions reflects some imprecision in timing the fixed-interval duration, and that the "lack of discriminative control" described by Staddon and Innis (1966, 1969) is, in fact, a lack of good temporal discrimination of the time of reinforcement.

METHOD

Subjects and Apparatus

The subjects were four experimentally naive White Carneau pigeons (*Columba livia*) obtained from the Palmetto Pigeon Plant, Sumter, South Carolina. They were housed individually with free access to water, and were maintained at 80% of their free-feeding weight.

The experimental chamber was a standard Lehigh Valley Electronics pigeon conditioning chamber. The subject space was 30 x 34 x 34 cm (length x width x height). An aluminum wall of the chamber had three 2.5-cm-diam response keys, each mounted 25 cm above a mesh floor. Only the center key was operative. It could be lit white or red, and was located directly above the 5 x 5 cm food magazine, which was 10 cm above the floor. A 5-cm speaker was mounted 7 cm to the right of the magazine, and could emit a 0.25-sec audible "beep." A 15-W houselight was mounted on the same wall, 32 cm above the floor. An externally mounted fan provided ventilation and masking noise. The chamber was housed in a light- and sound-attenuated room.

Procedure

The subjects served five days per week throughout the study, and were given supplemental feedings as needed, to maintain prescribed 80% ad-lib weights. The subjects were trained to peck a white key for food reinforcement (3.5-sec access to pigeon checkers). When food was presented, the houselight was extinguished, and the magazine was illuminated. When all of the subjects were pecking the white key (within 1-5 sessions), Phase I began.

Phase I: Uncued omission. The subjects were placed in a dark chamber. Onset of the houselight and red keylight signaled the start of the first fixed-interval cycle. Each cycle was 16 sec long, and was associated with a probability of reinforcement of .5. If a given cycle was scheduled to end in reinforcement, the first keypeck after the 16-sec interval had elapsed produced food. The next cycle followed immediately, with the houselight and red keylight on. Cycles scheduled for nonreinforcement (omission) ended at 16 sec without regard to responding (noncontingent), and the next cycle began without a signal (i.e., the houselight and red keylight remained illuminated). A trial consisted of n , 16-sec cycles. Trials ended with a probability of .042 after each cycle, and were separated by a variable-time intertrial interval (ITI) averaging 50 sec (with a range of 2.13-173.54 sec), during which the keylight was off and disabled and the houselight was on. (The reason for the trial structure with an ITI defined by the dark keylight is irrelevant to the present report, and reflects our interest in the analogy be-

tween this procedure and the peak procedure.) For the present purposes, only cycles after the first reinforcement in a trial were analyzed. The maximum number of cycles per trial was 32, with a mean of 12. The mean number of trials per session was 8, so that the mean number of cycles per session was 96. Thus, on average, the subjects received 48 reinforced cycles and 48 omission cycles per day. Such a design means that half of the time, reinforcement was given 16 sec after the previous reinforcement, that one quarter of the time it was given 32 sec after the previous reinforcement, that one eighth of the time, it occurred 48 sec from the previous reinforcement, and so on. The birds received 86 sessions before moving on to Phase II of the experiment.

Phase II: Cued omission. Except for events at the end of successive cycles, the basic procedure was the same as in Phase I. An additional reinforcement cue—a 0.25-sec audible “beep”—was sounded at the start of reinforcement. Omission cycles still ended on a response-noncontingent basis, but omission was signaled with part of the complement of stimuli normally associated with reinforcement. Specifically, the 0.25-sec “beep” was presented and the food hopper swung briefly into feeding position, just long enough to make its characteristic “bang,” but not long enough for the bird to obtain food. The magazine light was also briefly flashed. Training in Phase II continued for a total of 34 sessions.

Phase III: Cued and uncued omission. All events were as in Phase II, except that on omission cycles, the probability was .5 that the omission was cued (“beep” plus “bang” and magazine light), and .5 that the omission was uncued, as in Phase I. Thus, for a given cycle, the probability of food reinforcement was .5, the probability of cued omission was .25, and the probability of uncued omission was .25. Training in Phase III continued for a total of 75 sessions.

RESULTS

Phase I: Uncued Omission

Figure 1 shows response rates from Days 6–10 of training in Phase I. Note that these data are from a training period roughly equal, in terms of cycles, to that of Kello’s (1972) experiment. Data of individual subjects, as well as pooled data from all subjects, are presented. The left-hand section of each panel (showing heavy-line functions) depicts average responding as a function of time since reinforcement. The origin represents the start of postreinforcement cycles. At 16 sec, either reinforcement is given or there is an uncued omission (shown, respectively, by heavy-line and thin-line functions in the middle section of each panel). The bar graphs on the right-hand side of each panel depict average *total* response rate during the second cycle shown (i.e., overall response rate during 16–32 sec. The dark bar represents responding in cycles following a reinforcement, and corresponds to the heavy line in the response-rate function. The light bar represents responding in cycles following an uncued omission, and corresponds to the thin line. Total responding in a cycle is the typical index of the omission effect, and the classic omission effect is clearly seen for each bird. Responding following uncued omission is nearly double that following reinforcement. A two-way ANOVA (birds $\times$ reinforcement vs. omission) was performed on the data from each of the five sessions entering into the average in the histogram bars. The .05 level of significance was used here (and is used throughout this paper). There was a very large effect of omission versus

reinforcement, and a somewhat smaller—but still highly significant—effect of subjects. There was also a significant interaction between subjects and conditions, which was due to a somewhat higher differential for Bird 5426 associated with reinforcement versus uncued omissions.

The functions in the average graph depicting responding as a function of time demonstrate that overall response rate is greater following an uncued omission because there is no discernible pause in responding when there is no external cue to signal the end of one cycle and the start of the next. When reinforcement is given at 16 sec, however, there is a break, together with a pause before responding accelerates toward 32 sec, resulting in lower overall responding.

Figure 2 shows response rates for individual subjects during the *last* 10 sessions of training in Phase I, together with the pool of all subjects’ data. As before, the origin represents the start of postreinforcement cycles, and heavy-line functions show what happened when the subsequent cycles were reinforced. The thin-line functions depict responding in cycles that began with an uncued omission. The bar graph again depicts total average responding during 16–32 sec, the dark bar representing responding following reinforcement, and the light bar representing responding following uncued omission.

By the end of this phase of training, the omission effect, as measured by total responding during a cycle, remained essentially unchanged. Responding following uncued omission was nearly double that following reinforcement, and again, a two-way ANOVA showed these differences to be reliable. However, the form of the response function can be seen to have changed substantially. The subjects showed a decline in rate after the uncued omission, which then leveled off to flat responding. This is in contrast to the completely flat responding in post-uncued-omission intervals observed in the early sessions of the present phase, and by Kello (1972). It is likely that increased sharpness of the timing functions reflects training well beyond that in previous studies (Kello ran seven sessions of FI training followed by only seven sessions of omission training). Further, the shorter cycle used here (16 sec as opposed to 2 min) may have facilitated accurate timing.

The peak near 16 sec and the subsequent decline indicate that the time of the omitted reinforcement was discriminated. Note that especially when several omissions in a row occur by chance, the procedure is similar to the peak procedure (e.g. Gibbon & Church, 1990; Roberts, 1981), in which reinforcement normally scheduled at a fixed interval is sometimes omitted. The light-line functions show that responding declines after reaching a peak near 16 sec, but in this procedure, the declining response rate does not reach zero, presumably because subjects expect food occasionally at harmonics of the 16-sec interval (e.g., at 32 and 48 sec after a previous reinforcement).

It should also be noted that postreinforcement response rates reach their peak and begin to decline a few sec before the scheduled end of the interval in reinforced cycles. This is a common finding with extended fixed-interval training (e.g., Dews, 1970), and is sometimes

Figure 1. Average responses per sec and overall response rate (bar graphs) for individual subjects from Days 6–10 of training in Phase I, and the pool of all subjects' data over the same period. Heavy-line functions indicate postreinforcement rates, and dark bars indicate total responding during 16–32 sec for the postreinforcement cycles. Lighter functions and bars show rates when the cycle began with an uncued omission.

seen in the peak procedure. It may indicate "hopper tending," an appetitive behavior incompatible with keypecking, near the time of food. However, the decline after omission continues beyond 16 sec, indicating that the animal "knows" that the first time of possible reinforcement (i.e., 16 sec) has passed.

Phase II: Cued Omission

In this phase of the experiment, reinforcement omissions were cued by brief stimuli that normally accompanied food reinforcement. To mimic the uncued-omission

procedure, the signals were not contingent upon responding. The data are shown in Figure 3. Again, they are averages from the last 10 days of training for individual subjects and the pool of all subjects, and the heavy-line functions and dark bars indicate the average response rates in the cycles that began with reinforcement. The dashed-line functions and hatched bars indicate average response rates when reinforcement was omitted and cued. When reinforcement omissions were signaled, all subjects showed a clear break in responding, and then a substantial omission effect. Responding started consid-

Figure 2. Average responses per sec for individual subjects from the last 10 days of training in Phase I, and the pool of all subjects' data over this same period. Heavy-line functions and dark bars indicate postreinforcement responses per sec, and lighter functions and bars show rates when subsequent cycles went unreinforced.

erably earlier after cued omissions than after reinforcement, and the bar graphs show that cued omissions evoked greater total amounts of keypecking in the next cycle, although they did not result in a higher peak rate. A two-way ANOVA (birds $\times$ reinforcement vs. cued omission) showed a large main effect of omission and a large subject effect. There was also a significant interaction, due to the reduced cued-omission effect for Bird 4596 (post hoc analysis showed that the omission effect did not quite reach significance for this bird). Note also that the difference in total response rate between cycles that began with reinforcement and cycles that began with un-

cued omission (Phase I) is greater than the difference between cycles that began with reinforcement and cycles that began with cued omission (Phase II). This is the omission effect seen by Kello (1972) and Staddon and Innis (1966, 1969). This difference is examined in the same setting in Phase III, below.

Phase III: Cued and Uncued Omission

The purpose of this phase was to examine the omission effect in the context of cued- and uncued-omission cycles. By the end of this phase, the subjects had been exposed to 195 sessions (approximately 18,700 cycles)

Figure 3. Average responses per sec for individual subjects from the last 10 days of training in Phase II, and the pool of all subjects' data over this same period. Heavy-line functions and dark bars indicate postreinforcement rates. Dashed functions and hatched bars show rates when subsequent cycles were unreinforced, but cued.

of omission training, and were thus given a good deal of opportunity to habituate to a schedule filled with failed expectations. The data in Figure 4, taken from the last 10 days of training in this phase, indicate that the omission effect was persistent, although somewhat diminished, and cue dependent. The heavy-line functions and dark bars indicate average rates when the previous cycle ended in food. The dashed-line functions and hatched bars are from cases in which a postreinforcement cycle was followed by a cued omission, and the lighter, solid-line functions and bars show response rates when a postreinforcement cycle was followed by an uncued omission.

The histograms depicting overall rate of responding during 16–32 sec show that for three of the four birds, responding is greatest after uncued omission, less after cued omission, and least after reinforcement—as Kello (1972) found. A two-way ANOVA (subjects [birds] $\times$ reinforcement condition [reinforcement, cued omission, or uncued omission]) was performed on these data. The analysis again showed a large subject effect, a large effect due to reinforcement condition, and a significant interaction. Post hoc analysis showed that the significant interaction was due to the low rate of responding of Bird 4596 after uncued omission.

Figure 4. Average responses per sec for individual subjects from the last 10 days of training in Phase III, and the pool of all subjects' data over this same period. Heavy-line functions and dark bars indicate postreinforcement rates. Dashed functions and hatched bars show rates when subsequent trials were unreinforced and cued. Lighter solid functions and bars show rates when subsequent trials were unreinforced and uncued.

The decline in response rate after uncued omissions is sharper than it is in Phase I, especially for Bird 4596. Evidently, continued training on this procedure (or perhaps the additional training with cued omissions) results in a fairly sharp peak, usually 1 or 2 sec before the 16-sec delivery point, followed by a sharp decline to a low, steady rate of responding. After cued omissions, responding still begins earlier than it does after reinforcement, but the effect is diminished somewhat from Phase II. Note also that the peak after cued omissions generally occurs slightly earlier than it does after reinforcement.

DISCUSSION

While a failed-discrimination account appears consonant with data for cued omissions, it is harder to understand its role in uncued omissions, in which there is no pause following reinforcement omission, but rather a decline to a low, roughly constant rate during the subsequent interval. In the cued case, the external stimulus allows the animal to judge the precise time of the omitted reinforcement, and hence there is a clear break in responding at that point. It does not seem reasonable to group these two kinds of omission effects together as re-

sulting from a similar mechanism. One might possibly argue that from a "failed-discrimination" point of view, discrimination has failed partially in the cued case and completely in the uncued case. However, the response-rate functions show that when the omission is uncued, the animal is temporally processing the time of expected reinforcement, but, as in any timing task (e.g., a peak procedure), there is a degree of error in temporal judgments, which results in the peak in responding near 16 sec, followed by a smooth decline. In the next section, we show that these two very different function forms are both consonant with a timing account.

Timing: Scalar Expectancy Theory (SET)

Imagine subjective time accumulating in a counter/accumulator and response rate reflecting a shift to a high rate when elapsed time approaches the time of prior reinforcement experiences. Reinforcement cues reset the accumulator completely, but a subset of these (i.e., omission cues) may leave some residual accumulation. According to this view, the reason that response rates increase earlier after cued omission is not because the omitted event is still anticipated, but because the next reinforcement is anticipated sooner. In the absence of all omission cues (i.e., in the uncued case), the timer is not reset at all, but continues to accumulate throughout the next interval. Thus, the decline in responding indicates that the accumulation has exceeded that associated with reinforcement at the end of the cycle. Subsequent, higher, accumulation values may be associated with prior reinforcement at 32 sec (or 48 sec, and so on) after a preceding reinforcer. We will develop this idea further when the data from Phase III are fit to the timing model.

Notice that a failure of reset with a diminished set of reinforcement cues would predict not only that responding after cued omissions should start earlier, but also that it should reach its peak and decline earlier. To analyze this effect, we selected cycles of one or two *uncued* omissions in a row that were preceded either by reinforcement or by a cued omission. These data are shown in Figure 5. The solid-line functions show response rates when one or two uncued omissions in a row followed food reinforcement; the dashed-line functions indicate response rates when successive uncued omissions follow a *cued* omission (that in turn had been preceded by a food reinforcement, not shown here). The following four features of these data are apparent in all individual subjects: (1) responding starts earlier in the interval after a cued omission than after reinforcement, and this is the hallmark of an omission effect; (2) for at least three of the four birds, and in the pooled data, the peak after cued omission is slightly earlier than the peak after food, and it also declines earlier; (3) responding does not rise quite as high after a cued omission as it rises after a food reinforcement; and (4) the decline of both functions is to a common level (i.e., after an uncued omission, responding declines to about the same overall rate [approximately one response per sec], regardless of whether the omission was preceded by food or by cued omission); if

(3) reflected differential motivation, however, we would expect that differential to also have affected the height of the flat portion of the functions at 32 and 48 sec. We will later argue that the difference in peak heights mainly reflects the relative precision with which food is anticipated after a complete reset, as opposed to a partial reset, of the timing system.

Scalar Expectancy

The bottom panel of Figure 6 shows a schematic scalar timing model for responding in the baseline uncued-omission condition (Phase I). Memory distributions for the three possible reinforcement times are depicted on the ordinate. Because the response functions peak slightly before 16 sec, reflecting what we believe is a tendency to approach the hopper just before reinforcement time, we have adjusted the memory distributions to reflect this tendency. Thus, the peaks are located at 13.5, 29.5, and 45.5 sec. The diagonal represents mean accumulated subjective time after a preceding reinforcement. The model assumes that animals begin responding at a high rate when accumulated time is within a response window centered around the peak of the memory distributions. High-rate responding occurs when subjects are within one standard deviation (*SD*) above or below the peak memory time. The scalar property is reflected in the size of the response windows, which increase proportionally with increases in the remembered time. Thus the relative time window associated with the first harmonic (32 sec) is twice as broad as that associated with the fundamental (16 sec). Similarly, the window around the second harmonic (48 sec) is three times as broad as that associated with 16 sec. Notice that the *SD* window for 16 sec does not overlap with the others, while the windows for the longer times do overlap. This is because the *SD* in memory for the reinforcement times increases proportionally with each harmonic, but the distance between harmonics remains the same (i.e., 16 sec).

A computer fit to the baseline data is shown in the top panel of Figure 6. The smooth heavy-line function through the data points represents the scalar-expectancy account, assuming that the subjects are timing independently three different postfood times. The computer algorithm performed an exhaustive search for the best-fitting function under the following constraints: (1) The short reinforcement time, nominally 16 sec, was allowed to vary in order to accommodate the leftward shift of the peak evident in the data. The best-fitting function found this to be 13.5 sec, reflecting 2.5 sec of hopper tending prior to reinforcement, as in the schematic diagram above. The second and third times of reinforcement were similarly shifted by 2.5 sec (i.e., to 29.5 and 45.5 sec, respectively). (2) The fit determined a single coefficient of variation, $\gamma = \sigma/\mu = .232$, for the memory distributions associated with the three reinforcement times (the scalar property). (3) The absolute level of responding in the window near 13.5 sec was permitted to differ from response rates near 29.5 and 45.5 sec. The fitting algorithm found the rate near 13.5 sec to be higher (3.1 responses per

Figure 5. Peak-rate analysis of data from Phase III. Solid functions show response rates when one or two uncued omissions in a row followed food reinforcement. Dashed functions show rates when one or two uncued omissions in a row followed a cued omission (that itself had been preceded by food [not shown]).

sec) than the rate near the longer times (1.0 responses per sec). (4) The fitting algorithm picked the best relative time window for responding by determining a threshold on the relative proximity to the remembered times of reinforcement (e.g., as described in Gibbon & Church, 1990). This value was found to correspond to approximately one *SD* of the memory distributions, as in the schematic diagram.

Both the data and the fit show a clear peak near 16 sec, but a high, flat response rate near the harmonics. The scalar property causes the *SD* to increase proportionally with the value of the harmonics, thus producing the high,

flat response rate. Leak and Gibbon (1995) have shown similar high, flat tails in timing functions for multiple intervals when the separation between reinforcement times is small (less than about 4:1). Note that here the separation is two and three times the shortest interval.

Motivation or Timing

We now wish to consider variants of the timing account that might produce the omission effect. In SET, a distinction is made between motivational changes that may alter responsiveness by changing the threshold (i.e., the

Figure 6. Top: Data and scalar-timing-model fit for responding during successive uncued omissions after reinforcement. The heavy-line function through the data is the fit for anticipation of 13.5, 29.5, and 45.5 sec with a common threshold at about 25% of the relative discrepancy between current time and target time. Bottom: Schematic representation of the accumulation of subjective time (ordinate) with real time (abscissa). The memory distributions on the ordinate represent independent memories for the remembered times of reinforcement with the scalar property on variance (coefficient of variation = .232). The hatched response windows reflect the threshold on relative discrepancy that resulted in one *SD* above and below the mean for each distribution. Note the overlap between the response windows for the longer harmonics.

size of the response window) and effects that are due to timing (i.e., the pacemaker and count-accumulation process). The former are under motivational control, while the latter are affected by timing parameters.

One way in which enhanced responsiveness might occur is via an emotional effect, whereby the threshold is less conservative around the target time. In the top panel of Figure 7, we show the uncued-omission data from both the postcued omission condition (open circles) and the postreinforcement condition (filled circles, as in Figure 6). If enhanced responding were due to a motivational bias inducing a less conservative threshold, responding should increase on both sides of the peak time, but no shift would be expected in the location of the peak. The data, however, show a leftward shift, as well as a broadening of the curve.

An alternative motivational-arousal view might argue that the timing system simply runs faster after cued omissions. In these circumstances, one would expect some leftward shift in the peak after cued omissions, as opposed to after reinforcement, but the shift should be associated with a function that is *more*, rather than *less*, peaked. This is because the rate of the pacemaker governs the degree of precision, the resolution of the system estimating reinforcement times. A faster pacemaker results in greater precision. In the data in Figure 7, it is clear that the resolution has decreased rather than increased.

Moreover, it may be shown that an increase in pacemaker rate should produce a smaller leftward shift in the rise to the peak, and a larger leftward shift in the decline from the peak. If anything, however, the reverse is seen in the data after cued omissions.

Figure 7. Top: Postreinforcement uncued-omission data (filled data points and heavy-line function) and postcued-omission uncued-omission data (open points and light-line function). The postcued-omission data were fit with the same threshold and target-time parameters found for the postreinforcement uncued-omission data, but with incomplete reset with some variability, and a lower response rate near the shortest interval. Bottom: Schematic representation of the accumulation of subjective time with real time. The memory distributions are the same as those in Figure 6. The distribution associated with the starting point for accumulation induced by variation in the degree of reset associated with cued omission has been added. The heavy-line diagonal beginning slightly above zero is set at the mean starting position required to induce the 2.0-sec leftward shift in the peak. The dashed lines set at ± 1 *SD* of this mean illustrate the effect of this variance on the response windows.

Another feature, and one that was apparent in Phase II as well, is that generally the postcued-omission peak is lower than the postreinforcement peak. Thus, while the total amount of responding is somewhat greater in a postcued-omission interval, it is greater by virtue of starting earlier, not by virtue of showing a higher peak in anticipation of food. This finding argues against a frustration account, which would postulate more vigorous responding, late as well as early, in a postomission interval.

Failed Discrimination as Incomplete Reset

We propose that the failed-discrimination view described by Staddon and Innis (1966, 1969) is actually a failure of the timing system to reset completely with cued omission as opposed to reinforcement. After a cued omission, the accumulator begins with some nonzero value, and so reaches its criterion time somewhat earlier. Given the same threshold used for omission after reinforcement, we expect a symmetric leftward movement in the rise toward reinforcement and in the decline after reinforcement. The thin-line function in Figure 7 through the open circles is a fit to the postcued-omission data allowing an incomplete reset with some variability, but with no change in the threshold or remembered target times. High-rate responding for the shortest time interval was somewhat reduced in the postcued-omission fit (2.7 vs. 3.1 responses per sec), but later response rates were the same. The subjects were assumed to have an average positive count in their accumulation of subjective time, but there was some variability from cued omission to cued omission in terms of how much of a residue had been left. It is this variability that induces the somewhat less precise, broader peak function. Its maximum reflects a 2.0-sec reduction in peak time, corresponding to a positive average accumulation remaining when the next interval after a cued omission is begun.

In the bottom panel of Figure 7, a schematic representation of the incomplete-reset account is shown. The memory distributions on the ordinate are the same as those for the postreinforcement case (Figure 6); however, additional variability has been added in the accumulation process, reflecting variable starting positions (highly peaked distribution near the origin). The heavy-line accumulation diagonal has been set with a positive y -intercept at the mean residue of subjective time after incomplete reset. This results in the leftward shift of 2.0 sec (to 11.5 sec) seen in the postcued-omission data. The dashed lines represent starting positions $\pm 1SD$ around this mean. They illustrate how variability in the degree of reset effectively broadens the response windows around the remembered reinforcement times (hatched bands). The postcued-omission-peak functions fit this account quite closely. While we cannot argue that these data are definitive, they clearly suggest that a failure of reset of the timing mechanism is at the root of enhanced responding in the omission effect.

We have argued that this interpretation is consistent with the discrimination-failure hypothesis for the omission effect only if one identifies discrimination failure with failure to reset the timing mechanism before timing the next interval. It is less clear, however, how the uncued-omission-peak function might fit with the discrimination hypothesis. If partial discrimination failure is argued to result in partial reset failure, complete discrimination failure should appear when there is no reset, as in the uncued-omission procedure—yet the peaks that the subjects evince after uncued omissions look very similar to peak-procedure response functions (e.g., Gibbon & Church, 1990), centered near the target time. The difference here is that there is also a high right tail of additional responding, consonant with expectation of food at later times that are appreciated much less precisely. Thus, it appears that the cued and uncued situations are different, with cues forcing an incomplete or partial reset of the timing process, while in the uncued-omission procedure, the animal simply continues timing with its internal clock to estimate the first target time and subsequent harmonics.

REFERENCES

- AMSEL, A. (1958). The role of frustrative nonreward in noncontinuous reward situations. *Psychological Bulletin*, *55*, 102-119.
- AMSEL, A., & ROUSSEL, J. (1952). Motivational properties of frustration: I. Effect on a running response of the addition of frustration to the motivational complex. *Journal of Experimental Psychology*, *43*, 363-368.
- CATANIA, A. C. (1970). Reinforcement schedules and psychophysical judgements: A study of some temporal properties of behavior. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 1-42). New York: Appleton-Century-Crofts.
- DEWS, P. B. (1970). The theory of fixed-interval responding. In W. N. Schoenfeld (Ed.), *The theory of reinforcement schedules* (pp. 43-61). New York: Appleton-Century-Crofts.
- GIBBON, J., & CHURCH, R. (1990). Representation of time. *Cognition*, *37*, 23-54.
- KELLO, J. E. (1972). The reinforcement-omission effect on fixed-interval schedules: Frustration or inhibition? *Learning & Motivation*, *3*, 138-147.
- LEAK, T., & GIBBON, J. (1995). Simultaneous timing of multiple intervals: Implications of the scalar property. *Journal of Experimental Psychology: Animal Behavior Processes*, *21*, 3-19.
- ROBERTS, S. (1981). Isolation of an internal clock. *Journal of Experimental Psychology: Animal Behavior Processes*, *7*, 242-268.
- STADDON, J. E. R. (1967). Attention and temporal discrimination: Factors controlling responding under a cyclic-interval schedule. *Journal of the Experimental Analysis of Behavior*, *10*, 349-359.
- STADDON, J. E. R., & INNIS, N. K. (1966). An effect analogous to "frustration" on interval reinforcement schedules. *Psychonomic Science*, *4*, 287-288.
- STADDON, J. E. R., & INNIS, N. K. (1969). Reinforcement omission on fixed-interval schedules. *Journal of the Experimental Analysis of Behavior*, *12*, 689-700.