TIMING THE SECOND RESPONSE IN TWO-RESPONSE AVOIDANCE

STEPHEN FAIRHURST AND JOHN GIBBON

NEW YORK STATE PSYCHIATRIC INSTITUTE AND COLUMBIA UNIVERSITY

Rats were trained on a free-operant avoidance task requiring two lever presses within R seconds, with the opportunity for each response distinguished by differing stimuli. Response latencies at a variety of response-shock intervals were found to be proportional to the time available for the response. These results are shown to be consonant with a scalar expectancy model of timing behavior.

Key words: avoidance, response-shock interval, timing, response latency, scalar expectancy model, rats

Our recent theorizing about avoidance (Gibbon, 1979) has conceived of the task as (1) detecting an increment in aversiveness associated with the passage of time, and (2) forestalling the delivery of the aversive event, thereby reducing current aversiveness to some background or irreducible context level. Both of these features are presaged in Anger's (1963) account. The idea is that the behavior in an avoidance situation is governed by this relative reduction. This reduction depends in part on the response task and in part on the degree to which subjects can make the required temporal discriminations.

It seems to us that the avoidance task is intrinsically a difficult one if the components of the behavior take any appreciable time themselves. Consider a rat that must perform a sequence of responses, say depressing five different levers in succession to avoid shock. Even were one to successfully train subjects on this task, behavior might well deteriorate rapidly. To maintain the behavior, shocks must occasionally be delivered. That is, we expect extinction to intervene between successive shocks so that subjects gradually perform the chain more slowly and get shocked when they fail to complete it in time. But then the aversive event ought to act as a punisher for the earlier components of the response chain, since it di-

rectly follows them, albeit on some schedule. Subjects would successfully avoid when a response chain reaches completion in time but would be intermittently punished for early components of the chain. If we carry this argument further, one might conceive of early portions of the final component being punished as well, when the chain is nearly, but not quite, completed in time.

A simple retranslation of this chained-avoidance argument holds that this is precisely the situation for any avoidance contingency, except that the subcomponents before the final one are not experimentally specified. A rat, in pressing a lever to avoid shock, must position itself properly, raise its paws or body or head over the lever, and move down upon it in a manner analogous to Skinner's early description of the initial shaping of appetitive behavior (Skinner, 1938). Thus, the intrinsic nature of the avoidance contingency puts subjects into a situation in which early components of the required chain may be frequently punished, especially at the outset of training. In our view, something of this nature is what constrains avoidance acquisition. To the extent that early components are correlated with punishment, responding is suppressed, whereas to the extent that the completed chain is correlated with a reduction in aversiveness, responding is supported.

The notion of a chain with discrete units as components of successful avoidance leads naturally to the attempt to study an explicit avoidance chain in detail. In the experiment reported here, we trained subjects on a two-

This research was supported by National Science Foundation Grants BNS76-01229 and BNS78-23616 to John Gibbon. Reprints may be obtained from Stephen Fairhurst, Department of Psychophysiology, New York State Psychiatric Institute, 722 West 168th Street, New York, New York 10032.

response avoidance task that might be thought of as an FR-2 avoidance contingency, but with multiple external stimuli distinguishing the first from the second response. Note that without multiple distinguishing stimuli, subjects under a free-operant ratio avoidance contingency might learn that flurries of responding are effective but fail to learn the actual FR requirement. The correlation of successive responses in a chain requirement with specific cues permits the subject to distinguish shocks delivered for failing to make the first response from shocks delivered for failing to make the second response. One reason for interest in such a paradigm is the question of where in the chain errors occur. The location of chain errors was examined parametrically over a range of preshock intervals, including a range of performances from maximally to minimally efficient.

A second and more important focus of this work is the role played by the latency of the subcomponents of an avoidance chain. A ratio requirement alters the temporal context within which successive responses are made. Immediately after a successful avoidance, subjects in a free-operant avoidance situation face a delay of, say, R sec before the next shock is due. If a response occurs within this interval, the delay is reset. Immediately after a response, when shock is due R sec hence, the potential shock rate is $1/R$ sec, the lowest rate that the subject can achieve. With the passage of time, the potential shock rate increases to more aversive levels, such that the new shock rate at any given time, t , since the last response is $1/(R-t)$. Our avoidance model (Gibbon, 1971, 1972, 1979) argues that it is this differential—between $1/R$ and $1/(R-t)$ —that engenders responding. Subjects of course have some difficulty in discriminating the two aversiveness levels, but the theory says that they do so stochastically by detecting when the ratio of the two levels exceeds some criterion.

In this experiment we ask how this conception of the task applies to the idea of subcomponents of an avoidance response. In the simplest case, two responses are required for avoidance. The question arises: What is the temporal context for each of the two responses considered singly? Perhaps the differential between the background and the current aversiveness level operates in a similar fashion here as when a single response is required. The mo-

tivation for making the second response, once the first response has been completed, may be thought of in this way—that is, as supported by a reduction from current levels of aversiveness to the level of $1/R$. The motivation for completing the first component is less clear. There is a sense in which the first response of a two-response chain results in a discriminable increase in aversiveness. That is, the first response produces conditions that have been at least occasionally accompanied by shock at short delay. Another way to put this is that completion of the first response associated with an R -sec preshock interval puts subjects into a new avoidance task that has a preshock interval that is shorter by precisely the latency of the first response. Since short preshock intervals are demonstrably harder to master than long preshock intervals, subjects, in a sense, “move from the frying pan into the fire” by completing the first response in two-response avoidance. It is as though the temporal context has substantially worsened precisely to the degree that the first response is well-timed and thus closer to shock.

One way that subjects might master the task without, so to speak, facing the increased aversiveness produced by performing the initial component, is to chain the two responses as closely as possible in time so that they are performed as a unit and the task reverts to standard free-operant avoidance. Less interest would then accrue to the performance since it would be analogous to the single-response task. On the other hand, the task should still be somewhat more difficult than standard avoidance, in that performing two responses takes longer in absolute time. Our previous analysis of avoidance suggests that subjects fail to compensate for a fixed execution time involved in the avoidance latency (Gibbon, 1979). One result that would be expected, then, is that for a long response requirement, no matter how performed, a higher shock rate would be expected with a short R interval than with a long R interval. We give the reasons in some quantitative detail elsewhere (Gibbon, 1979). A qualitative description would note that a 2-sec execution time represents 20% of a 10-sec preshock interval but 2% of a 100-sec shock interval. Thus subjects that respond at about the same proportion of the interval would nearly always complete a 2-sec response when the preshock interval is

Fig. 1. Possible response patterns for two-response avoidance, where each response is associated with a fixed component, M .

100 sec, but might fail to complete a 2-sec response initiated late in a 10-sec interval. Thus, minimally, a ratio requirement in avoidance may be expected to increase the difficulty of the task even if subjects treat the multiple requirement as a "unit."

The first two alternatives shown in Figure 1 depict possible chaining patterns for two responses, where each response is associated with a fixed time, M . In the first case, the first latency (L_1) is "timed" in the sense that a substantial time, say about a third of R , elapses before the first response is initiated. A fixed time, M , then elapses, comprised at least in part of the actual time required to physically execute the response. L_1 is then followed immediately by the second response (L_2) also requiring at least M sec to complete. In the second line, the first response is made immediately after the start of the trial (and hence immediately after the preceding second response) and is followed by a timed second response.

The interesting case arises if the responses are sufficiently distinct and the opportunities to respond are perhaps separated slightly in time, so that subjects are, so to speak, encouraged to view a FR-2 requirement as two distinct tasks, each with its own temporal background, as depicted in the third line in Figure 1.

Our initial attempt to make the responses distinct involved two separate levers with different response topographies required (one vertical and one horizontal). From a relatively large subject population, we were unable to train any rats to perform this task. We subsequently moved to a ratio requirement on a single lever that was withdrawn from the chamber briefly between responses and that

was associated with distinct signals for the first and second response. As described below, this task was difficult but not impossible for subjects to acquire. If subjects treat the two responses as a "unit," then one would expect the Response-1 minus Response-2 latency (or Response-2—Response-1 latency) to be short and insensitive to changes in the R interval. If, on the other hand, subjects treat Response 1 as one task and Response 2 as a second, then both latencies may be expected to be sensitive to preshock intervals.

In the extreme case, the two tasks might be regarded as entirely independent, except that the latency of the first defines the preshock interval of the second. Thus a comparison of the first response latency with a single-response task should reveal no essential differences. We call this possibility "loose chaining." Another alternative, "tight chaining," involves a closer interaction between the two tasks. Because of the requirement of a second response, the first response tends to occur earlier. Subjects may adopt a "conservative" strategy that entails completion of the two-response requirement at about the same point as for a single-response task—even though both responses are "timed" and in that sense independent tasks. These two possibilities are considered in more detail later.

METHOD

Subjects

The subjects were three male hooded rats, screened from a pool of 32 animals. All subjects in the pool were experimentally naive and approximately three months old at the start of screening. All were placed on a standard free-operant avoidance schedule with a shock-shock (S) time of 10 sec and a response-shock (R) time of 100 sec. Thirteen of the original pool mastered this single-response avoidance well enough to be screened with the two-response paradigm described below. Only three animals maintained responding in the two-response situation.

Apparatus

All subjects were conditioned in an SPA 100 experimental chamber with internal dimensions of 20 cm wide by 20 cm high by 23 cm long, inside an SPC 300 sound-attenuating chamber. The front wall of the experi-

mental chamber was equipped with a Coulbourn E 23-05 retractable lever centered 3.5 cm from the left wall and 6.5 cm above the floor, and with a cue light 4.0 cm above the center of the lever. A second cue light was located 3.5 cm from the right wall and 10.5 cm from the floor. The floor consisted of 16 stainless steel bars, each .35 cm in diameter, with their centers 1.5 cm apart. Shock was delivered to the floor through an SP 4008 J scrambled shock source, with an intensity of .8 ma and a duration of .5 sec. A houselight provided dim illumination. A tone and clock used for cues were generated by a Foringer 1166-4-M1 stimulus source.

Procedure: Two-Response Avoidance

In this procedure, two responses were required for avoidance of shock. After delivery of shock, the lever entered the chamber, the *S* clock was started, the cue light above the lever was lit, and a tone was sounded. When the first response was made, with latency L_1 , the lever retracted and returned immediately (the retraction and return cycle took about .5 sec), and the tone was terminated. When the second response (L_2) was made, the lever retracted and returned, the *R* clock was started, the left cue light went off and the right one turned on, and a click and tone were presented. The first response in *R* time (L_1) caused retraction and return of the lever and cessation of the tone, but the clock continued. A second response (L_2) caused the lever to retract and return, reset the *R* clock, and restored the stimuli to their start-of-*R*-time state. Failure to make the two responses within the *R* or *S* interval produced shock (the lever was withdrawn during shock), turned off the *R* clock and started the *S* clock, and returned the stimuli to their *S*-time configuration.

For recording purposes, each *R* or *S* interval was divided into 10 bins, and responses were recorded according to their bin number, both cumulatively in counters and sequentially on paper tape for later computer analysis.

Subjects were given alternate blocks of 10 sessions of a baseline reference condition with $S = 10$ and $R = 100$, and 10 sessions of testing at $S = 10$, $R = 50$; $S = 10$, $R = 10$; and $S = 10$, $R = 5$. Data are presented for the *R*-time behavior in the last five sessions in each block. *S*-time behavior was very similar to *R*-time behavior at $R = 10$.

RESULTS

Completed Avoidance, $L_1 + L_2$

Data in avoidance timing experiments typically show a mean response latency and standard deviation that are linear in the preshock interval, R , with a small intercept (cf. Gibbon, 1979).

In Figure 2 the mean latencies and SDs for the completed avoidance ($L_1 + L_2$) and the shock probabilities, $P(S^*)$, per trial (per *R*-clock initiation) are plotted against R , averaged for three subjects. The two-response data

Fig. 2. Mean response latency, SD, and shock probability vs. preshock interval for a shuttling response (circles), single lever press (squares), and the sum of two lever presses (diamonds), each the average of three subjects.

for the present experiment (filled diamonds) are shown along with the corresponding data from two previous experiments: an average of three subjects on a single lever-press avoidance task (Gibbon, 1979, unfilled squares) and an average of three subjects on a shuttle-box avoidance task (Libby & Church, 1974, unfilled circles). The functions are clearly similar in form but differ slightly in their intercepts.

Linear regressions are shown for each of the mean and *SD* functions. The approximation to linearity is quite good ($R^2 > .985$ for all functions). The mean functions show small positive intercepts and the *SD*s small negative ones.

The shock probabilities (probability that the *R* interval will time out before the avoidance is completed) all show the same declining form but different levels for the two-response and single-response tasks. The declining form may be associated with a fixed component that effectively shortens the interval in which the decision to respond can be made. For completed avoidances, then, the current data are very similar to other avoidance timing experiments with respect to the mean and *SD*, and similar in form for $P(S^*)$.

Avoidance Components, L_1 , L_2

The latency of the first press, L_1 , is taken as the time from the start of the trial to the first press, and the latency of the second press, L_2 , as the time between the first and second press, omitting the time required for the lever to retract and return. The *SD*s for each press are *SD*s of each of the latencies. Shock probabilities are, for the first press, the probability that no response is made, and for the second press, the probability that, given a first response, no second response is made before the interval elapses.

If only one press is timed, then the latency and *SD* for the other press should be small, and more important, constant. That is, the untimed press should be made relatively quickly, and its latency should not vary with the setting of the *R* interval. For example, if the first press is timed and the second untimed, the second press should act like a fixed component, increasing the shock probabilities more at short intervals than at long ones, and should produce the typical declining shock probability function. The timed press, however, should

display substantial latencies linearly related to *R*. Since the latency for the untimed response should be small, the latencies and *SD*s for the timed press should be nearly equal to those for the completed avoidance. If both presses are timed, then both should show latencies and *SD*s that increase with *R*.

In Figure 3, the mean latency, *SD*, and shock probability for the first press (triangles) and for the second press (inverted triangles) are shown averaged over the three subjects, along with the functions for the completed avoidance (diamonds). Linear regressions are shown for the latency and *SD* functions. As in Fig-

Fig. 3. Mean response latency, *SD*, and shock probability vs. *R* interval for the first (triangles) and second (inverted triangles) response, and the sum of the two responses (diamonds), average of three subjects.

ure 2, the latency and *SD* functions are approximately linear, and shock probabilities decline with increasing *R* values.

It is clear that timing is exhibited in both presses. For both presses the mean latencies and *SD*s increase approximately linearly, and the shock probabilities decline with increasing *R*. The approximation to linearity for the latency and *SD* functions is quite good ($R^2 > .974$ for all functions).

The forms of the functions clearly indicate that each press is timed; their levels may tell us what is being timed. In all cases the latency and *SD* functions are lower for the second response than for the first, and the shock probabilities are higher for the second response. A shorter latency, smaller *SD*, and higher shock probability are what would be expected if a shorter interval were being timed. We can readily interpret the second press as being associated with a shorter interval. At the start of the trial, the animal has *R* seconds in which to respond. On the average, the first response will be made at L_1 , the mean latency of the first press. A second response is now possible, and the data show that the second press is timed, rather than being made immediately or

at some constant delay after the first. However, the animal no longer has an interval of length *R* before shock but only an interval of $R - L_1$. If second presses are distributed through this preshock interval in the same way that first presses are distributed through their preshock interval (*R*), the mean latencies and *SD*s when plotted against *R* should be lower and shock probabilities higher than for the first press. If the data for each press are plotted against the effective preshock interval, *R* for the first press, and $R - L_1$ for the second, the two functions should superimpose, if both are timed in the same manner.

In Figure 4, the mean latencies, *SD*s, and shock probabilities for each response are shown for each animal plotted against the effective interval: *R* for the first press and $R - L_1$ for the second press.

The functions retain the property exhibited in Figure 3, near proportionality to the preshock interval, but now the functions for each of the two presses are nearly colinear. For each animal the mean and *SD* for each press is nearly proportional to the time available for that press, with approximately the same slope and intercept. This is evidence that the sub-

Fig. 4. Mean response latency, *SD*, and shock probability vs. preshock interval for the first (triangles) and second (inverted triangles) response for each subject and for the average of three subjects.

jects are timing the effective preshock intervals in the same way. Again the addition of a small fixed component is indicated by a non-zero intercept.

The two shock-probability functions do not superimpose as well as do the means and *SDs*, but they are closer to one another than when plotted against *R*. This, together with the large difference in the level of shock probability between this experiment and comparable single-response tasks (Figure 2), suggests that some other factor associated with the two-response task may be entering into the shock probabilities.

The data confirm that scalar timing provides a good description of the latencies of each response separately. The mean latencies are nearly proportional to the interval, the *SDs* are proportional to the size of the latency and hence proportional to the interval, and the shock probabilities decline as the length of the interval increases. And for each press, the interval being timed is the time during which it is possible to make that press: the entire interval for the first press and the time remaining for the second press.

Latency Distribution

In Figure 5, the two panels on the left show latency distributions averaged over subjects for each of four *R* intervals (100, circles; 50, diamonds; 10, triangles; and 5, squares) plotted against real time. That is, each point represents the relative frequency with which responses occurred in the class interval whose midpoint is represented on the *x*-axis. (Bins with real-time values less than 1 sec were excluded to eliminate the effects of *M*.) The top panel is for the first response latencies and the lower panel for second response latencies. The functions are plotted on a log scale and on this scale appear quite similar in form but displaced to the right with increasing *R* interval. The declining, approximately linear (exponential in real time) form is pronounced for two of the three subjects but less so for the third. However, all subjects showed a marked similarity of form and displacement in time with successively increasing *R* intervals. If these distributions were exactly similar in form and simply displaced to the right by a constant, then they would be expected to super-

Fig. 5. Response latency distributions averaged for three subjects at *R* = 5, 10, 50, 100. First response (top row) and second response (bottom row) vs. real time (left column) and vs. proportion of interval (right column).

impose when plotted against proportions of the interval—relative time—as shown in the right-most two panels. Again the L_1 latency is in the top panel and L_2 in the lower panel. To a first approximation, superimposition is achieved here, with some deviation from strict superposition observed in the first one or two bins. Although the form of these functions for the individual subjects differed slightly (and idiosyncratically) from that shown in the average, the superposition was about as good for individuals as that shown here. Thus subjects distributed their latencies over the remaining time in the interval in a similar manner relative to the size of that interval regardless of the absolute time remaining.

Relationship Between L_1 and L_2

The data presented so far represent averages for many trials and show that, on the average, both responses are timed. Let us now examine how timing takes place within a given trial.

Suppose, for example, that R is 100 sec, and that the animal responds, on the average, when 1/5 of the available time has elapsed. Then, on the average, the first press should occur at 20 sec. The time available for the second press will, on the average, be 80 sec, and the second press will be made, on the average, when 1/5 of 80, or 16 sec of this time has elapsed. But now suppose that on a particular trial the first press is actually made when only 10 sec—1/10 of the interval—has elapsed. There are then several alternatives for the second response. The animal might revert to its average proportion of the interval and

- A. time 1/5 of the average time left, $1/5(80) = 16$ sec, or
- B. time 1/5 of the actual time left, $1/5(90) = 18$ sec. Alternatively, it might maintain the same proportion of the interval as was represented by the (deviant) first press, and
- C. time 1/10 of the average time left, $1/10(80) = 8$ sec, or
- D. time 1/10 of the actual time left, $1/10(90) = 9$ sec.

Similarly, each of these strategies could be applied when the first press is made late. Mathematically (ignoring for now the contribution of M) they are expressed as:

- A. $l_2/(R-L_1) = L_1/R$
- B. $l_2/(R-l_1) = L_1/R$
- C. $l_2/(R-L_1) = l_1/R$
- D. $l_2/(R-l_1) = l_1/R$,

where L_i is the main latency, l_i is the particular latency, $i = 1, 2$. For Alternative A, l_2 is independent of l_1 ; for B, l_2 varies inversely with l_1 ; for C, l_2 varies directly with l_1 ; and for D, l_2 increases to a maximum and then decreases as l_1 increases.

Figure 6 shows the relationship between l_1 and l_2 for the three subjects at $S = 10$, $R = 100$. In the upper panel the mean L_2 for an L_1 in a given bin is shown for each subject (H-23, triangles; H-122, squares; H-134, diamonds) and for the average of three subjects (filled circles). The mean represents individuals quite well. The mean (with its SD) is reproduced in the lower panel, which shows the several theo-

Fig. 6. Correlation of first and second responses. Top: Mean L_2 for a given L_1 for three subjects (triangle, square, diamond) and for the average of three subjects (circles). Bottom: The average L_2 latency for a given L_1 is shown with a filled circle and its SD with a barred line. The dashed lines A, B, C, D show the theoretical value of L_2 under various timing strategies.

retical relationships described above. In both panels, the area below and to the left of the solid diagonal marked "shock" delineates the area in which responses are possible, where the sum of the latencies is less than T . The four dashed lines marked A, B, C, and D show the predictions for L_2 according to the four alternatives: (A) the average proportion of the average time left; (B) the average proportion of the actual time left; (C) the same proportion of the average time left; and (D) the same proportion of the actual time left.

The data are more like Curve D than any of the others but deviate systematically from D toward B. Both D and B assume that the animal is timing the actual time left on each trial, whereas A and C assume that it considers each remaining interval to be equal to the mean. In D the animal is timing the same proportion of the remaining interval as it did the entire interval; in B it reverts to the mean proportion. Thus the data suggest that the animal times the actual remaining intervals on each trial and that when the proportion of the interval at which it made its first response differs from the mean, the proportion at which the second response is made tends to regress toward the mean proportion. This raises the interesting possibility that the variation in these latencies results primarily from variations in the threshold at which a discriminable increment in aversiveness is detected. Variation in Latency L_2 appears to be less dependent on variation in the estimation of the time remaining than on variation in the proportion of the interval that the animal uses for a given trial.

If variation in the estimate of the time remaining were the primary determinant of the variability in latency, then the two latencies, L_1 and L_2 , should be positively correlated. To see this, consider that if, on a given trial, the animal's appreciation of the remaining time were incorrectly long, the first response would be made with a latency longer than the mean and the second response also made with a latency longer than the mean. Of course, this assumes that error in appreciating R is directly correlated with error in appreciating $R - L_1$. In Figure 6, however, it is clear that long L_1 latencies (longer than their mean) are associated with short L_2 latencies (shorter than their mean).

DISCUSSION

The experimental data for the two-response situation indicate that the animal times both responses and that each response is made when a given (similar) proportion of the available time has elapsed.

In scalar expectancy theory, the animal is seen as estimating the interval and responding when a given proportion of this time has elapsed. The variations in its estimates are proportional to their size, so that the SD is a constant portion of the mean, which in turn is proportional to the size of the interval. In effect, a given distribution of estimates is spread across the timing interval, whatever that interval may be, and the probability of a response occurring within any given proportion of the interval, say within the first fifth of the interval, is the same for any interval.

Typically, the mean latencies plotted against the length of the interval are linear with a small positive intercept. This intercept represents, we believe, a small fixed-time component in the latency, which may represent (or include) a motor time—that is, the time it takes the animal to actually execute the response. The SD s often show a small negative intercept, which is also consonant with the idea of a small fixed component (Gibbon, 1979).

Shock probability—the proportion of cases in which the interval elapses before a response is made—typically shows a graded decline with increasing R . If the response latencies were strictly proportional to the interval, this decline would be quite unexpected, since the proportion of trials in which the latency is longer than R should be the same for all intervals, just as the distribution of latencies less than R is the same for all intervals. But if there is a small fixed component in the latency, latencies are effectively increased by that amount, increasing the probability of shock. Suppose, for example, that this fixed component is one second. Then a response begun immediately at the start of the trial is completed one second after the start, a response begun four seconds into the trial is completed at five seconds, and so forth. Any response begun less than a second before the interval elapses cannot be completed before shock intervenes. When the interval is short, say five

seconds, this last second, in which a response cannot be completed, represents a fifth of the trial, so that responses begun in the last 20% of the trial are ineffective and result in shock. When the interval is long, say 100 seconds, this last second represents only 1% of the interval. Thus the presence of a small fixed component will increase the shock probability at short intervals much more than at long intervals and can give rise to a function that declines with increasing R .

We have previously fitted the single lever-press and shuttle-box-task data with the expectancy model (Gibbon, 1972, 1979). In this model there is associated with each trial in an avoidance situation a given amount of aversiveness, A , which spreads across the preshock interval. If the animal's estimate of the length of the R interval is x , then at the start of the trial, the level of aversiveness, $a(0)$ is A/x . At any time, t , in the interval, the total amount of aversiveness A remains the same but is spread across the shorter remaining portion of the interval, $x - t$, so that the level of aversiveness $a(t) = A/(x - t)$ is higher. The animal compares the current level of aversiveness $a(t)$ with the initial level, $a(0)$, forming the expectancy ratio:

$$r(t) = \frac{a(t)}{a(0)} = \frac{A/(x - t)}{A/x} = \frac{x}{(x - t)} = (1 - t/x)^{-1}. \quad (1)$$

Note that the amount of aversiveness, A , cancels, so that the comparison is a relative one—"How much worse is the situation now than before?"—rather than an absolute one—"How bad is it now?"

For the single-response case, the response will be made when the expectancy ratio $r(t)$ crosses some threshold value B , i.e., when the level of aversiveness is B times its initial value.

The expectancy ratio $r(t)$, the ratio of current aversiveness to baseline aversiveness, plotted against t rises hyperbolically from an initial value of 1.0, as shown in Figure 7 (top panel). When the expectancy ratio reaches the Threshold B , the response is made. The time at which this Threshold B is crossed defines the response latency, l . At $t = l$, the expectancy ratio is B times the initial value of 1.0, or B . That is,

$$r(l) = B = (1 - l/x)^{-1} \quad (2)$$

and hence

$$l = x(1 - 1/B). \quad (3)$$

Potentially there are two sources of variance in Equation 3. The estimate of the preshock interval, x , from a variety of other evidence, surely contains some variance, and that variance displays the scalar property (Church & Gibbon, 1982; Gibbon, 1972; Gibbon & Church, 1981). The Threshold B may also be expected to vary, as would be required to ac-

Fig. 7. Scalar expectancy models for FR-2 avoidance. Top: Single response. The expectancy ratio $R(T)$ rises from 1.0 to Threshold B , where the response is made. Middle: L_2 = new avoidance task. A new expectancy ratio, $R(T - L_1)$ starts after L_1 is made and rises to Threshold B , where L_2 is made. Bottom: L_2 = component of $L_1 + L_2$. $R(T)$ crosses Threshold $\sqrt{B}$ where L_1 is made, then continues to rise to Threshold B , where L_2 is made.

commodate the infrastructure of the data observed in Figure 6.

It is readily shown that the mean latency, L , with either or both sources of variation should be

$$L = (1 - l/B)R, \quad (4)$$

assuming the estimate, x , of R is accurate in the mean. We refer the reader to our previous discussion (Gibbon, 1979) for a quantitative examination of the role of the fixed time M in the analogues of Equation 4. For an appreciation of the qualitative relationships to be expected, we will for the present deal only with the mean values, L , B , and R , and consider the contribution of the fixed component to be negligible.

Equation 4 may be expected to hold for the first response in a two-response task. The first response should occur at a proportion of the preshock interval, R , defined by the first threshold, say, B_1 . A second response is also timed, but timed relative to the preshock interval $R - L_1$. We wish now to distinguish two ways in which the first and second responses might be related.

L₂ As a New Avoidance Task

If, as we argue, the second response times the remaining preshock interval, then one conception of this task is that it is entirely separate from the first response except that the first response defines the preshock interval. That is, subjects may be regarded as again sampling from an internal representation of the preshock interval adjusted for L_1 and forming an expectancy ratio that crosses the threshold when

$$r(L_2) = B_2 = [1 - L_2/(R - L_1)]^{-1}, \quad (5)$$

assuming an unbiased estimate of $R - L_1$. Thus on the average,

$$L_2 = (1 - 1/B_2)(R - L_1). \quad (6)$$

The data show that the two responses are made at equal proportions of the time available for them; that is, $L_2/(R - L_1) = L_1/R$. Substituting this empirical relationship into the above equation and solving for B_2 , we find

$$B_2 = B_1 \equiv B. \quad (7)$$

Thus, if having crossed the Threshold B_1 and responded once, the animal resets his expectancy ratio to 1.0 and begins timing a new

interval equal to the time remaining before shock, it will respond for a second time when the new expectancy ratio crosses that same threshold. The first response having been made when the situation became B_1 times more aversive than at the start of the trial, the second response will be made when it becomes B_1 times more aversive than at the start of the second opportunity. Equivalently, starting with the assumption that the two thresholds B_1 and B_2 are equal, the colinear relationship $L_2/(R - L_1) = L_1/R$ follows. This situation is depicted in Figure 7 (center panel). The first latency, L_1 , occurs when the expectancy ratio first reaches B , and the second latency, L_2 , occurs when the expectancy ratio again reaches B for the new, shorter, preshock interval.

L₂ As a Component of L₁ + L₂

A second conception of the second response latency would regard it as a component of the total completed avoidance latency in the following sense: Imagine that subjects time the complete avoidance, $L_1 + L_2$, inserting l_1 along the way. They would thus follow Equation 4 with $L = L_1 + L_2$. Then if the subcomponents L_1 and L_2 of the task are differentiated, they must be performed in a manner that conforms to the overall requirement of Equation 4. On this view, the second response must occur (i.e., the avoidance be completed) when the expectancy ratio, calculated as the aversiveness ratio from the beginning of the trial, crosses B . This is shown in Figure 7 (bottom panel). There, the expectancy ratio continues to rise until it reaches the level B , at which time the second response will be made. This is indicated by the threshold for L_2 . But the first component must have been completed previously, and the colinearity in the data requires that it be timed in the same manner as L_2 . That is, the empirical relationship $L_2/(R - L_1) = L_1/R$ forces the first response to occur when the expectancy ratio reaches $B_1 = \sqrt{B}$. The second response will then occur when the aversiveness is $\sqrt{B}$ times $\sqrt{B}$, or B worse than at the outset of the trial.

The two conceptualizations may be thought of as "loose" or independent chaining—two tasks—versus "tight" or dependent chaining—one task with two components. The distinction between them can be characterized more precisely by looking at the sources of within-trial as opposed to between-trial variability.

Between successive trials, variability can be introduced both by variations in the estimates of the preshock interval and by variation in the threshold for initiating responding. The sources of variability within a trial, however, differ in the two versions of the task.

In the independent or loosely-chained version, the second response represents a wholly new task that might well be expected to involve both a new estimate of the length of the preshock interval and a change in threshold. In the dependent or tightly-chained alternative, represented in the bottom panel of Figure 7, variability within a trial is restricted. Here a single estimate (x) of the time to shock is made at the onset of the trial. This estimate is then updated as the trial elapses and when the aversiveness associated with the updated time exceeds the first threshold, the response L_1 is made. There is no necessity for a second estimate of the remaining time to be made on this version, though one might readily conceive of a change in the component subthreshold B_1 and B_2 . Thus variability within the trial might arise from two sources, change in estimation of the interval and change in threshold, in the loosely-chained version, but from only one source, change in threshold, in the tightly-chained version.

There are two lines of evidence that suggest that the tightly-chained version of the task is closer to the truth. The first is the observed negative correlation between L_1 and L_2 . Although this correlation has not been fully analyzed in the sense that the constraints imposed by truncation at shock delivery force some negative correlation, nevertheless the data suggest that subjects 1) detect in a fairly direct way the actual, as opposed to the average, remaining time to shock, and 2) vary somewhat in their threshold from Response 1 to Response 2. If the independent task idea is taken seriously, then subjects should be able to re-estimate the actual time left depending on their first response latency. One way for subjects to accomplish this would be to store in memory the distribution of L_1 -to-shock times experienced during training. Although this seems unlikely on its face, there might well be a second estimate of the overall preshock interval R and adjustment for L_1 .

Variation in the threshold from Response 1 to Response 2, on the other hand, would accommodate some negative correlation between

Latency 1 and Latency 2, according to the second, tightly-chained version of the theory. The reason is that if the first threshold happened to be an outlier, it would be expected to be followed by a choice of threshold for the second latency that is closer to the mean, so that some negative correlation is built in via regression to the mean for the second threshold sample.

A second line of evidence comes from an examination of the absolute values of B and of the fixed-component M required for these data. We saw earlier that the latency functions and their SD s for the total latency, $L_1 + L_2$, corresponded quite well to the latencies and SD s obtained from a single lever-press task and were shorter than latencies obtained from a shuttling task. This correspondence between two-response total latency and single-response latency again suggests the tightly-chained conception with two subcomponents rather than the loosely-chained version with two independent components. To analyze this possibility further, we estimated B and M for the data in Figure 2 for the shuttle and single lever as well as the sum of the two components required in this experiment. The fitting procedure allowed the fixed component to differ for the three sets of data but required B , which defines the slope, to be the same. The results are shown in Figure 8, with the B value for the three sets found to be 1.74 and the M value for the two-response data to be 1.4. ω^2 values are, for the latencies, .979, .952, and .995 for the shuttle, single-lever, and two-response tasks respectively, and for the SD s, .943, .792, and .995, respectively. The fits to the shock-probability functions are less good, with the two-response-task prediction in particular being systematically lower than the data.

In the right-hand column, the values of M and B just obtained are applied to the two subcomponents of the two-response task, assuming the tightly-chained model in which the threshold for each subcomponent is $\sqrt{B}$ and the fixed time associated with each is half of the total M ($M/2$). The predictions of the latency and SD functions are quite good, ω^2 values for the latencies being .976 and .981 for L_1 and L_2 respectively, and for the SD s, .978 and .963, respectively. As would be expected from the shock-probability data in the left-hand column, the shock-probability functions for each response separately are also not well

Fig. 8. Mean latency, SD, and shock probability fitted by scalar expectancy theory. Left column: Shuttle response (circles), single lever press (squares), and sum of two lever presses vs. R interval, all with $B = 1.74$. Right column: First (triangles) and second (inverted triangles) of two lever presses vs. preshock interval with threshold $= \sqrt{B} = \sqrt{1.74}$, and $M_1 = M_2 = M/2$.

fit. The form of the function is about right, but the levels of the data are much higher than is predicted from theory using a small fixed component. This again suggests, as noted earlier, that another factor enters into the determination of shock probabilities in the two-response task—perhaps a breakdown of timing on some proportion of the trials (cf. Gibbon, 1971).

Quite independently of the mechanism for producing higher shock rates in this task, it seems to us that these data and our analysis force the timing mechanism to contain at least the following general features: (1) The preshock interval for the first response is R sec and for the second response approximately R minus the first response latency. (2) Timing of these two preshock intervals involves an approximately constant average threshold or proportion of the interval, regardless of the absolute size of the interval. (3) Preshock in-

tervals, both first and second, are estimated in a scalar fashion so that variance increases as the square of the mean. (4) The two responses are probably best conceived as subcomponents of a single task rather than as two independent tasks, even though Points 1 and 2 might suggest such independence. (5) Shock rates at all values of the preshock interval for the two-response task are uniformly higher than for the single-response task but roughly proportionally so. This increase in shock rates of the two-response task over its single-response counterpart is not adequately explained by the assumption of a fixed component in the latency associated with the execution of the response. Such a fixed component appears to be involved in the latencies and SDs and does predict the form of the shock-rate function for a single-response task. However, the fixed component involved in the latencies and SDs is much too small to produce the high shock rates observed in the two-response task.

REFERENCES

- Anger, D. The role of temporal discriminations in the reinforcement of Sidman avoidance behavior. *Journal of the Experimental Analysis of Behavior*, 1963, 6, 477-506.
- Church, R. M., & Gibbon, J. Temporal generalization. *Journal of Experimental Psychology: Animal Behavior Processes*, 1982, 8, 165-186.
- Gibbon, J. Scalar timing and semi-Markov chains in free-operant avoidance. *Journal of Mathematical Psychology*, 1971, 8, 109-138.
- Gibbon, J. Timing and discrimination of shock density in avoidance. *Psychological Review*, 1972, 79, 68-92.
- Gibbon, J. Timing the stimulus and the response in aversive control. In M. D. Zeiler & P. Harzem (Eds.), *Advances in analysis of behaviour* (Vol. 1). Reinforcement and organization of behaviour. Chichester, Eng. and New York: Wiley, 1979.
- Gibbon, J., & Church, R. M. Time left: Linear versus logarithmic subjective time. *Journal of Experimental Psychology: Animal Behavior Processes*, 1981, 7, 87-108.
- Libby, M. E., & Church, R. M. Timing of avoidance responses by rats. *Journal of the Experimental Analysis of Behavior*, 1974, 22, 513-517.
- Skinner, B. F. *The behavior of organisms*. New York: Appleton-Century, 1938.

Received October 13, 1981

Final acceptance October 15, 1982