JOURNAL OF THE EXPERIMENTAL ANALYSIS OF BEHAVIOR 2000, **73,** 241–260 NUMBER 3 (MAY)

# *HUMAN PERFORMANCE ON NEGATIVE SLOPE SCHEDULES OF POINTS EXCHANGEABLE FOR MONEY: A FAILURE OF MOLAR MAXIMIZATION*

ERIC A. JACOBS AND TIMOTHY D. HACKENBERG

UNIVERSITY OF FLORIDA

Panel pressing was generated and maintained in 5 adult humans by schedules of points exchangeable for money. Following exposure to a variable-interval 30-s schedule and to a linear variable-interval 30-s schedule (which permitted points to accumulate in an unseen "store" in the absence of responding), subjects were exposed to a series of conditions with a point-subtraction contingency arranged conjointly with the linear variable-interval schedule. Specifically, points were added to the store according to the linear-variable interval 30-s schedule and were subtracted from the store according to a ratio schedule. Ratio value varied across conditions and was determined individually for each subject such that the subtraction contingency would result in an approximately 50% reduction in the rate of point delivery. Conditions that included the subtraction contingency were termed *negative slope* schedules because the feedback functions were negatively sloped across all response rates greater than the inverse of the variable-interval schedule, in this case, two per minute. Overall response rates varied inversely with the subtraction ratio, indicating sensitivity to the negative slope conditions, but were in excess of that required by accounts based on strict maximization of overall reinforcement rate. Performance was also not well described by a matching-based account. Detailed analyses of response patterning revealed a consistent two-state pattern in which bursts of high-rate responding alternated with periods of prolonged pausing, perhaps reflecting the joint influence of local and overall reinforcement rates.

*Key words:* maximization, matching, feedback functions, variable-interval schedules, negative slope schedules, panel press, adult humans

A feedback function is a quantitative description of the relation between reinforcement rate and response rate (Baum, 1973). With a feedback function it is possible to specify in quantitative detail how a particular reinforcement rate depends on a particular response rate. Every schedule has its own unique set of feedback functions, determined jointly by the schedule and the pattern of responding. The feedback functions for many common schedules have a positive slope, owing to a positive correlation between response rate and reinforcement rate. But feedback functions can also be negatively sloped. On such negative slope schedules, reinforcement rate and response rate are inversely related.

Feedback functions for negative slope schedules have some theoretically interesting features. In addition to reaching maximum value at relatively low response rates, negative-slope feedback functions normally are sharply peaked: Reinforcement rate reaches its maximum across a narrow range of response rates, then drops precipitously. This provides a clear reference point against which to assess sensitivity to molar reinforcement variables, such as overall reinforcement density, implicit in molar maximizing models. Experiments with rats (Ettinger, Reid, & Staddon, 1987; Reed & Schachtman, 1991) and pigeons (Vaughan & Miller, 1984) provide little evidence to support the predictions of molar maximization: Response rates deviate substantially from the apex of the feedback function, where overall reinforcement rate is at its maximum.

The present experiment extended the use

---

This research was supported in part by Grant MH 50249 from the National Institutes of Health to the University of Florida. The experiment was part of a doctoral dissertation submitted by the first author to the Graduate School at the University of Florida, supported in part by the Maurice C. and Frances A. Holmes Scholarship. We are indebted to Marc Branch, Brian Iwata, Mark Lewis, Hank Pennypacker, and Donald Stehouwer for helpful comments on an earlier version of the manuscript, and to Cynthia Pietras and Manish Vaidya for their assistance in conducting the research. Portions of these data were presented at the Southeastern Association for Behavior Analysis conference, Charleston, 1995, and at the Association for Behavior Analysis conference, Washington, D.C., 1995.

Reprint requests should be addressed to Eric A. Jacobs at the Department of Psychology, Southern Illinois University, Carbondale, Illinois 62901-6502 (E-mail: eajacobs@siu.edu) or Timothy D. Hackenberg at the Department of Psychology, University of Florida, Gainesville, Florida 32611-2250 (E-mail: hack@psych.ufl.edu).

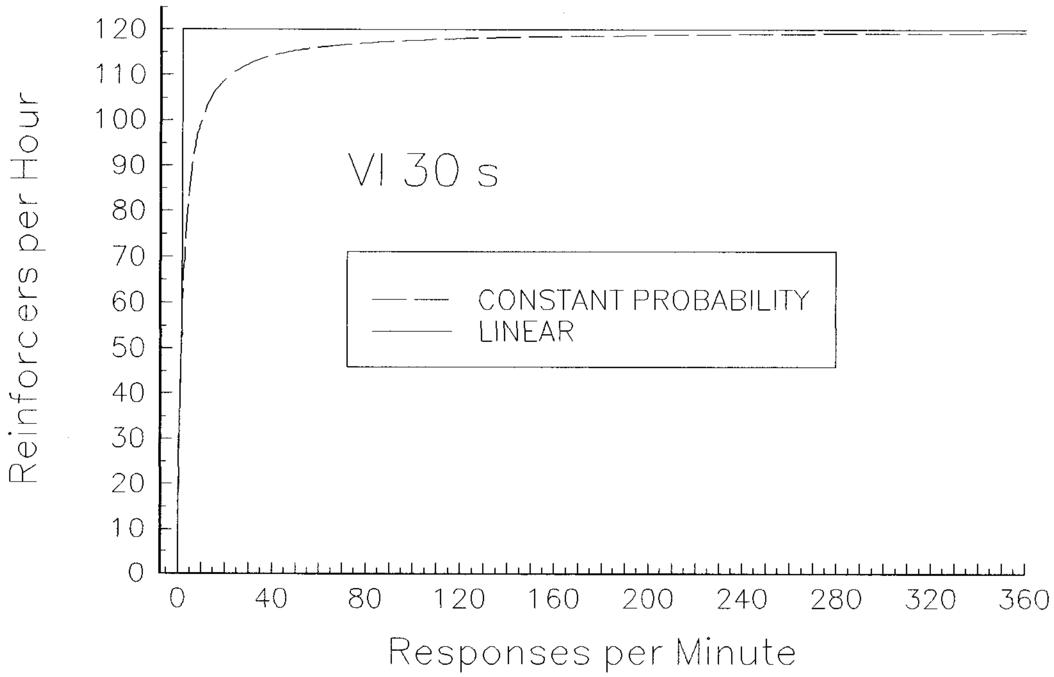

Fig. 1. Molar feedback functions for linear and constant-probability VI 30-s schedules. See text for additional details.

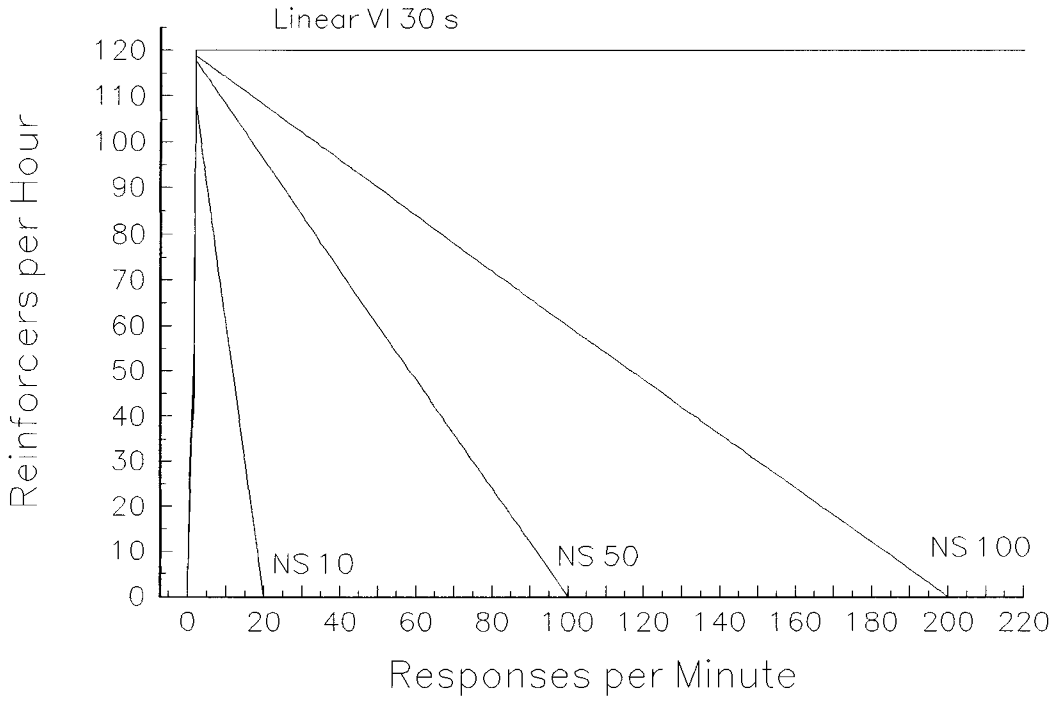

Fig. 2. Molar feedback functions for linear variable-interval (VI) 30-s and negative slope (NS) schedules with subtraction ratios of 10, 50, and 100.

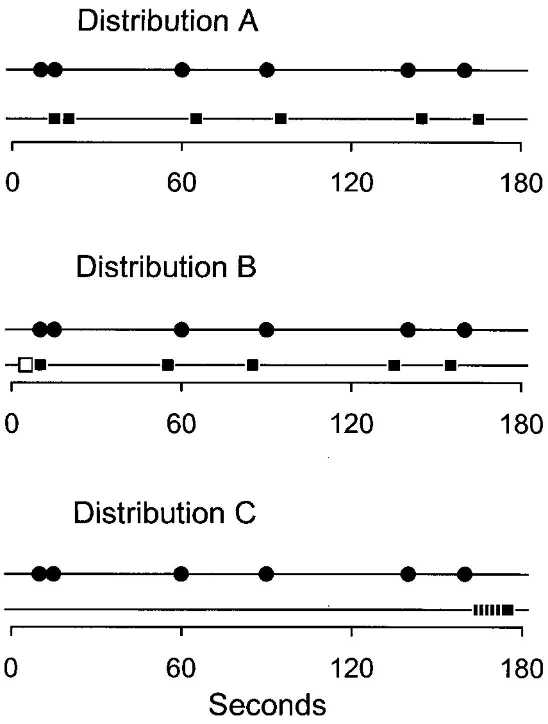

Fig. 3. Event diagrams illustrating possible interactions between the LVI schedule sequence and three response sequences. Filled circles indicate the conclusion of a schedule interval, filled squares indicate reinforced responses, and open squares indicate unreinforced responses.

of negative slope schedules to human subjects. Panel pressing was established and maintained in 5 adult human subjects by schedules of points exchangeable for money. Subjects were exposed to a range of schedules with various feedback-function characteristics, including many with negatively sloped functions in which reinforcement rate varied inversely with response rate.

The extension to human subjects was seen as significant in the context of a growing body of research aimed at examining the sensitivity of human behavior to temporally distributed reinforcement variables. Although there is a vast literature examining human behavior in relation to maximization principles, the majority of studies comprising this literature employ methods that depart significantly from those used in experiments with nonhuman subjects. Most of these experiments, for example, use hypothetical (imagined) rather than real consequences such as money (Kahneman & Tversky, 1979; Rachlin, Raineri, & Cross, 1991). When real consequences have been used, exposure to the contingencies has typically been very brief, with variables manipulated across subjects (Herrnstein, Loewenstein, Prelec, & Vaughan, 1993; Kudadjie-Gyamfi & Rachlin, 1996). These procedural features make comparisons to findings obtained with nonhuman subjects difficult at best. A major aim of the present research therefore was to examine human behavior under conditions in which subjects received prolonged exposure to contingencies that were manipulated on a within-subject basis. This should not only provide a strong test of molar maximizing but should also facilitate comparisons to similar research conducted with nonhuman subjects.

Before discussing how maximization predictions are brought to bear on negative slope schedules, it is important to consider in some detail how such schedules are arranged. A feedback function with a negative slope is constructed from two distinct schedules, one that arranges positive reinforcers and another that cancels them. Reinforcers are scheduled according to a linear variable-interval (LVI) schedule. LVI schedules differ from traditional variable-interval (VI) schedules in that upon completion of an interval a reinforcer is added to a "store" of available reinforcers, and timing of the ensuing interval begins immediately. Reinforcement follows every response when the store is positive, and one is subtracted from the store following each reinforcement. So, for example, if there were three reinforcers in the store, the next three responses would be reinforced and the store would be depleted, thus requiring another interval to expire before reinforcement was again available.

Figure 1 shows the feedback functions for both an LVI 30-s schedule and a conventional VI 30-s schedule. Whereas the VI function asymptotically approaches the programmed reinforcement rate across the entire range of response rates, the LVI function is constant across all but the lowest response rates. LVI schedules are thus more forgiving of deviations in response rate than are traditional VI schedules, in that the maximum reinforcement rate is obtained for all response rates greater than the inverse of the schedule value (two responses per minute, in this case).

An inverse relation between overall re-

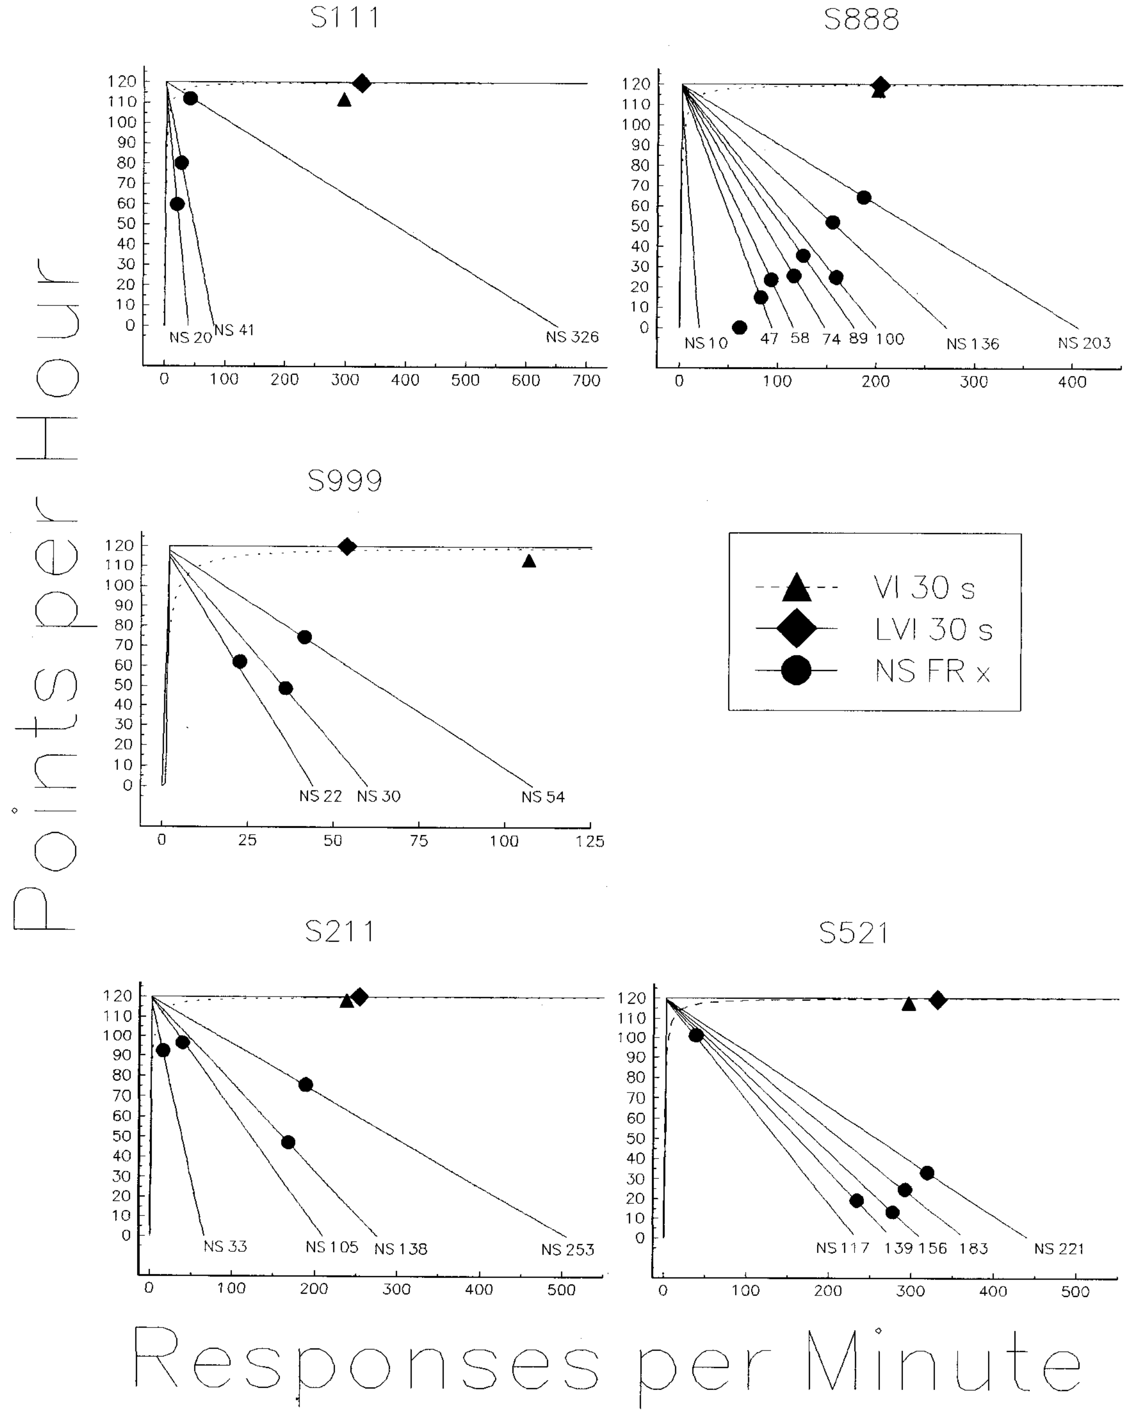

Fig. 4. Mean rates of point delivery versus mean response rates from the last five sessions of the VI 30-s (triangles), LVI 30-s (diamonds), and NS FR *n* (circles) conditions for each subject. The data are superimposed on the molar feedback functions for each condition. The uppermost curve in each graph is the molar feedback function for the LVI condition, and the dashed curve below it is the molar feedback function for the VI condition. The remaining curves are the feedback functions for the NS conditions. Numbers near the base of each function indicate the subtraction ratios of the NS conditions. Note individually scaled abscissae for each subject.

Table 1

The sequence of conditions and the number of sessions conducted under each for all subjects under variable-interval (VI), linear variable-interval (LVI), conjoint (CONJT), and negative slope (NS) schedules of point delivery. The numbers of sessions are in parentheses, and the size of the fixed ratio of the NS schedule is shown without parentheses.

| Subject | | | | |
|---|---|---|---|---|
| S111 | S888 | S999 | S211 | S521 |
| VI (36) | VI (8) | VI (10) | VI (5) | VI (11) |
| LVI (11) | LVI (5) | LVI (13) | CONJT (8) | LVI (7) |
| NS 326 (35) | NS 203 (6) | NS 54 (8) | VI (5) | NS 221 (8) |
| NS 41 (10) | NS 136 (5) | NS 30 (9) | LVI (5) | NS 183 (5) |
| NS 20 (13) | NS 100 (6) | NS 22 (10) | NS 253 (7) | NS 156 (8) |
| | NS 88 (7) | | NS 138 (6) | NS 139 (17) |
| | NS 74 (5) | | NS 105 (28) | NS 117 (25) |
| | NS 58 (6) | | NS 33 (29) | |
| | NS 47 (7) | | | |
| | NS 10 (9) | | | |

sponse rate and overall reinforcement rate can then be established by conjointly subtracting reinforcers from the store according to a second schedule, normally a ratio schedule in which a reinforcer is subtracted following some number of responses since the previous subtraction. In negative slope schedules, the store can be driven into a negative state following periods of relatively high-rate responding, thus requiring periods of relatively low-rate responding before the value of the store is again positive and reinforcement can occur.

Figure 2 shows the theoretical feedback functions that result from an LVI 30-s schedule with reinforcers conjointly subtracted every 100, 50, or 10 responses, that is, according to fixed-ratio (FR) 100, FR 50, and FR 10 subtraction schedules, respectively. The functions are negatively sloped across all response rates that are greater than the inverse of the LVI value (1/30 s, or two responses per minute). Beyond this point, higher response rates result in a linear decrease in reinforcement rates. Note the differences in the slopes of the functions in Figure 2. The rate at which reinforcement rate decreases as a function of response rate is inversely related to the FR

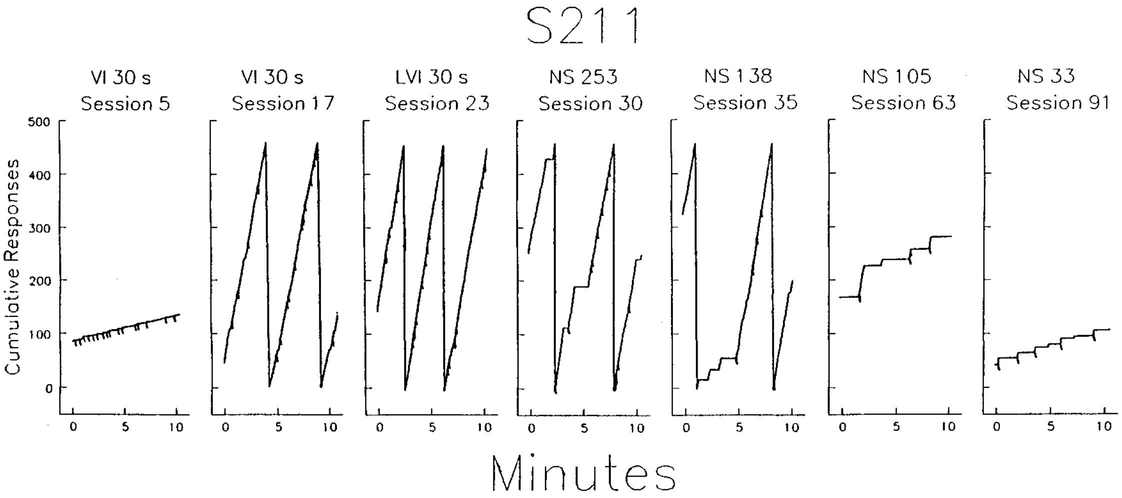

Fig. 5. Portions of representative cumulative response records from sessions conducted under VI, LVI, and NS conditions for Subject 211. Each panel is the segment of the cumulative record from the second 10-min block of the session.

value; reinforcement rate decreases sharply over a narrower range of response rates when reinforcement is canceled every 10 responses than when it is canceled every 50 or every 100 responses.

It is in this context that researchers (Ettinger et al., 1987; Vaughan & Miller, 1984) have viewed negative slope schedules as a strong test of global maximization. By this reasoning, if responding is sensitive to the overall rate of reinforcement, it should fall near the apex of the feedback function—in the region where response rate is approximately 1 divided by the LVI schedule value (two per minute in Figure 2). At this rate, all available reinforcers would be collected in a timely fashion, and the impact of the subtraction ratio on overall reinforcement rate would be minimized. Deviations from this theoretical maximum are then viewed as failures of maximization.

But does this constitute a fair test of maximization? The utility of feedback functions to assess sensitivity to molar reinforcement variables rests on certain assumptions (e.g., Baum, 1992). First, feedback functions assume that variation in molecular characteristics of behavior will not markedly affect overall reinforcement rate. Second, the function relating reinforcement rate to response rate is assumed to be continuous over the entire range of response rates; for every response rate, the function must specify a definite reinforcement rate.

One may question whether the feedback functions presented here and in previous studies of negative slope schedules meet these criteria. Specifically, variability in local (molecular) aspects of responding may produce variability in reinforcement rate in the portion of the curve surrounding the apex of the negative-slope feedback function. Thus, more than one definite reinforcement rate may be possible when the response rate approaches the maximum programmed reinforcement rate. Because the apex of the negative-slope feedback function is a critical point of reference for evaluating sensitivity to overall reinforcement rate, any appreciable deviation in the height or curvature of these functions may render them less useful as tests of overall rate maximization.

The LVI molar feedback functions presented here were based on the method proposed by Vaughan (1982). The LVI feedback functions consist of two discrete sections. At response rates less than the programmed reinforcement rate, the schedule will arrange and store reinforcers at a higher rate than they are collected, so the store will be continually in a positive state and every response will be reinforced (i.e., FR 1). At response rates equal to or greater than the programmed reinforcement rate, reinforcement rate will be constrained by the temporal requirements of the schedule, and the obtained reinforcement rate will equal the programmed rate.

At the intersection of the two sections of the LVI feedback function, where response rate equals the programmed reinforcement rate, obtaining the nominal programmed rate of reinforcement depends upon responses occurring only when the store is positive. Thus, obtained overall reinforcement rate could vary as a function of local characteristics of responding, such as the distribution of interresponse times (IRTs). For illustration, consider the event records of three different response distributions portrayed in Figure 3. Each record shows six responses distributed throughout the 3-min period. Thus, the overall response rate for each distribution was 2.0 per minute, and the maximum attainable overall reinforcement rate was 2.0 per minute. For a response to produce reinforcement, however, the store must be positive.

Distribution A shows responses occurring with the same temporal spacing as the schedule-interval distribution, except the response sequence commences later in the 3-min period than the schedule-interval sequence. All responses occur shortly after an interval has timed out and stored a reinforcer. Hence, all responses are followed by reinforcement, and the value of the store is zero at the conclusion of the 3-min period. Such a pattern would be

→

Fig. 6. Scatter plots depicting IRTs preceding point delivery across sessions for Subjects 211 (top panel) and 888 (bottom panel). Conditions labeled VI, LVI, CONJT, and NS correspond to variable-interval, linear variable-interval, conjoint, and negative slope conditions, respectively.

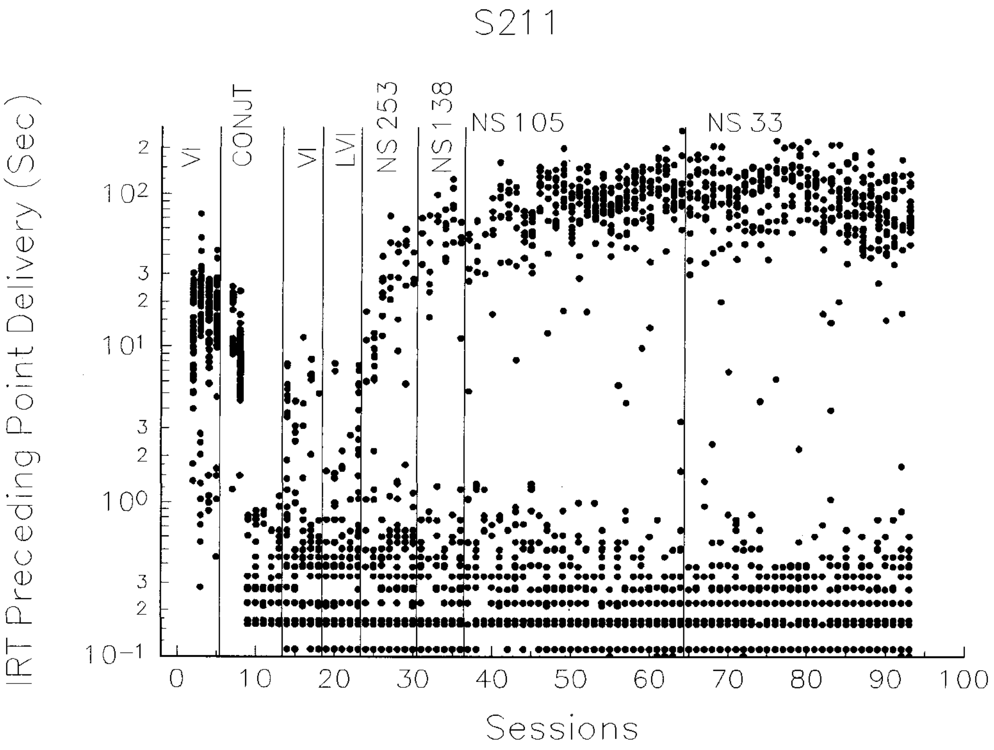
S211
IRT Preceding Point Delivery (Sec)
Sessions
VI
CONJT
VI
LVI
NS 253
NS 138
NS 105
NS 33
2
10²
3
2
10¹
3
2
10⁰
3
2
10⁻¹
0
10
20
30
40
50
60
70
80
90
100

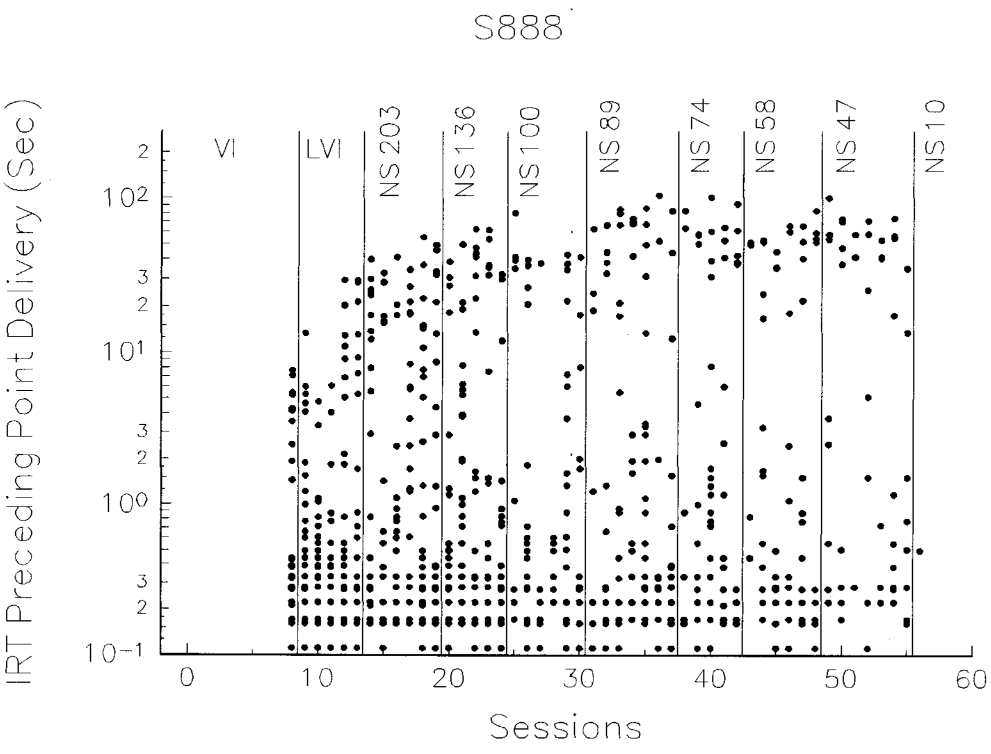
S888
IRT Preceding Point Delivery (Sec)
Sessions
VI
LVI
NS 203
NS 136
NS 100
NS 89
NS 74
NS 58
NS 47
NS 10
2
10²
3
2
10¹
3
2
10⁰
3
2
10⁻¹
0
10
20
30
40
50
60

very unlikely unless reinforcer availability were signaled in some fashion.

In Distribution B, the temporal spacing of the responses again matches that of the schedule-interval distribution, but the response sequence commences earlier than the schedule sequence. As a result of the leftward shift of the response distribution, the first response occurs before the conclusion of the first interval and, therefore, goes unreinforced. The resulting reinforcement rate is 1.67 per minute, and one reinforcer remains in the store at the conclusion of the 3-min period. Thus, slight changes in the temporal distribution of responses can produce changes in overall reinforcement rate under the LVI schedules at the intersection of the two sections of the feedback function.

For the steady-state data to fall along the negative slope functions as drawn, there must be a one-to-one ratio between responses emitted and reinforcers delivered; that is, all responses must be reinforced. Any deviations driven by molecular aspects of responding will produce an overall reinforcement rate that does not lie on the function. But the presence of distinct patterning does not necessarily produce deviations from the theoretical function. In Distribution C, all six responses occur in quick succession toward the close of the 3-min period. Because all responses occur after all of the schedule intervals have timed out, each response is reinforced, yielding a function consistent with that shown in Figure 2.

Because the reinforcers on negative slope schedules are arranged by LVI schedules, the feedback functions depicted for negative slope schedules in Figure 2 may be subject to the same limitations. Like the LVI functions, the negative-slope feedback functions consist of two discrete sections. Inclusion of the subtraction contingency shifts the intersection of the function of the two sections of the feedback function (the apex of the function) slightly to the left, to a lower response rate along the FR 1 based part of the function.

To assess the accuracy of the feedback functions presented here and in prior research, we conducted a series of simulations across a range of response rates in the critical region of the curve where response rate equals the maximum programmed reinforcement rate. The details of these simulations are described in Appendix A, but the results can be summarized briefly. Despite some local variability in point delivery, the vast majority (over 99%) of responses were in fact reinforced. These few unreinforced responses had small but negligible effects on overall reinforcement rate, and did not significantly alter the characteristics of the molar feedback functions. The feedback functions for LVI and negative slope schedules portrayed in Figures 1 and 2 appear to be reasonably accurate and therefore are valid tests of overall rate maximization.

It is possible that greater deviations would occur with response distributions other than those used in our simulations—distributions more closely matched to actual performance on LVI and negative slope schedules. Unfortunately, little is known about the local characteristics of behavior under these schedules. Prior research with LVI and negative slope schedules has focused nearly exclusively on overall response rate; realistic assumptions about local response patterning are therefore difficult to make. Thus, in addition to assessing the adequacy of molar maximization accounts of human performance, the present research sought a more thorough characterization of response patterning on LVI and negative slope schedules. To that end, daily cumulative response records and IRT measures were collected and analyzed in relation to the contingencies. This should not only provide important information on the adequacy of the feedback-function approach to assessing predictions of molar maximization but may also provide insights into the mechanisms that govern behavior on these interesting but seldom-used contingency arrangements.

→

Fig. 7. Scatter plots depicting IRTs preceding point delivery across sessions for Subjects 521 (top panel) and 999 (bottom panel). Conditions labeled VI, LVI, and NS correspond to variable-interval, linear variable-interval, and negative slope conditions, respectively.

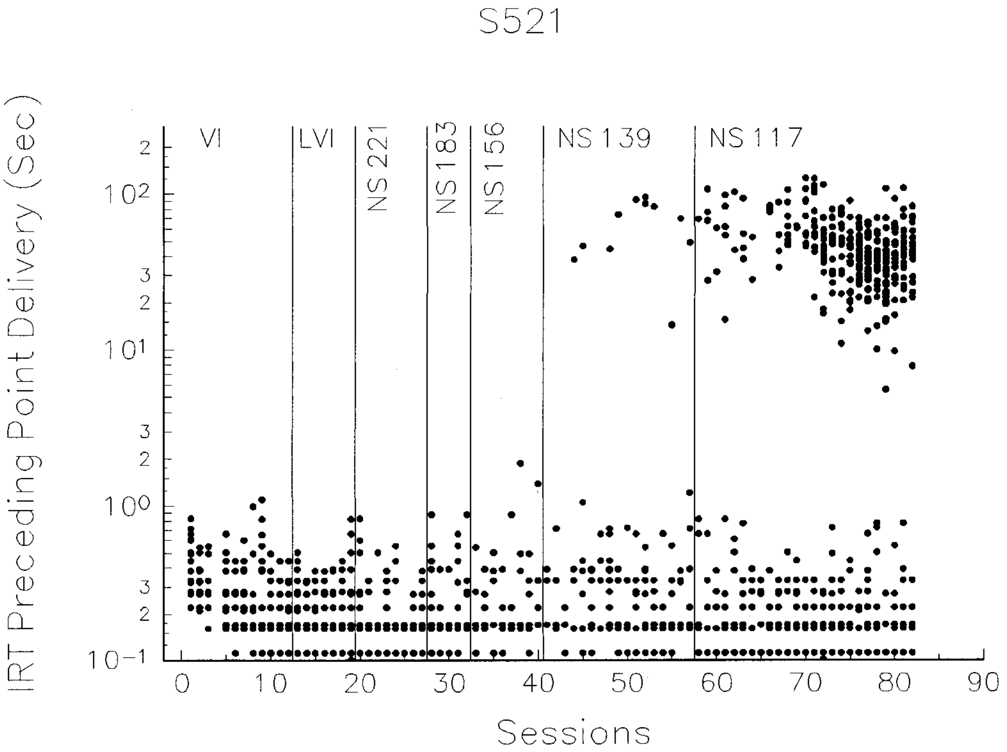
S521
IRT Preceding Point Delivery (Sec)
Sessions
VI
LVI
NS 221
NS 183
NS 156
NS 139
NS 117
2
10²
3
2
10¹
3
2
10⁰
3
2
10⁻¹
0
10
20
30
40
50
60
70
80
90

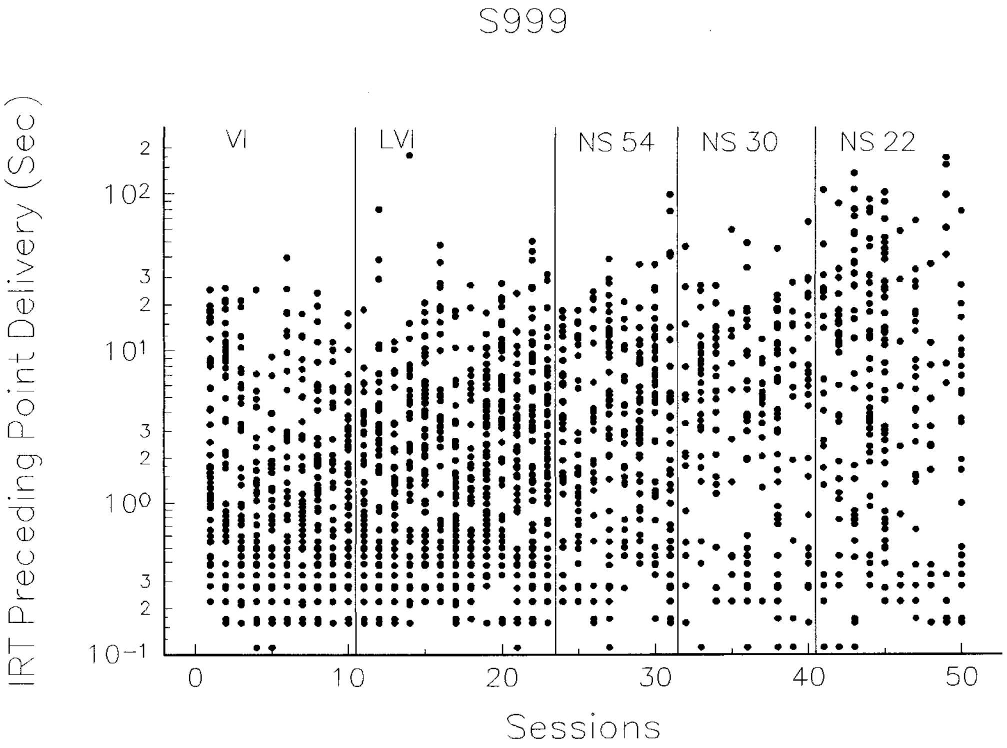
S999
IRT Preceding Point Delivery (Sec)
Sessions
VI
LVI
NS 54
NS 30
NS 22
2
10²
3
2
10¹
3
2
10⁰
3
2
10⁻¹
0
10
20
30
40
50

## METHOD

### *Subjects*

Two men and three women participated in exchange for money. The subjects were recruited via a classified advertisement in a campus newspaper. None of the subjects had previously or was currently enrolled in courses in behavior analysis or learning theory. Each subject was informed prior to the first session that he or she would earn $1.50 for each completed session in addition to bonus earnings of $0.05 for each point. To encourage full participation, collection of bonus earnings was contingent upon completion of the study. The subjects were also instructed not to bring personal items (e.g., food, tools, smoking materials, portable radios, or timepieces) into the work space and that violation of this rule would result in dismissal from the study. Overall earnings (including bonuses) ranged from $3.00 to $7.16 per hour (median = $6.00 per hour).

Six additional subjects were recruited for the study but did not participate beyond the initial condition. Five of these subjects withdrew from the study, and another was dismissed due to insufficient baseline response rates. The data from these subjects will not be considered here.

### *Apparatus*

The operandum consisted of the right side of a dual response panel (12.5 cm by 7.5 cm), illuminated with red light and mounted centrally on a BRS/Foringer human operant panel (HTC-603). A computer monitor was seated on top of the panel and a keyboard was located in front of the panel. Data were recorded and contingencies controlled by a program compiled in Quickbasic® running on a Compuadd® Model 286 computer. Cumulative response records were collected using a Gerbrands cumulative recorder.

### *Procedure*

The following instructions were read to the subject prior to the first session and were displayed on the computer monitor prior to the start of each session.

> Please read carefully. You have already earned $1.50. In order to receive bonus earnings, you will have to earn points by pressing the red panel. Each point you earn is worth $0.05. So, for example, if you earn 80 points then your bonus would be $4.00. You may leave the room at any time in the event of an emergency. Thank you for your participation.

Sessions were divided into 10-min blocks and began when the subject typed his or her subject number on the keyboard and pressed the enter key. Each block was separated by a rest period, during which time the manipulandum was darkened and no programmed contingencies were in effect. Rest periods were terminated by a response on the computer keyboard. Sessions were comprised of five such blocks for Subject 111 and three blocks for the other 4 subjects. Sessions were usually conducted on weekdays at approximately the same time of day.

Table 1 contains the sequence of conditions and the number of sessions conducted under each for the 5 subjects who completed the experiment. The number of sessions per subject ranged from 50 to 105 (median = 81). Conditions were changed when mean response rates and within-session patterns of responding were deemed stable on visual inspection of graphical representations of the data. For Subjects 111, 888, and 999, the final condition was concluded arbitrarily. For Subjects 111 and 999, the final condition ended due to scheduling conflicts. Subject 888 left town during a semester break and did not return to the experiment.

The following are descriptions of each of the four types of experimental conditions.

*Variable-interval (VI) 30 s.* Panel pressing was maintained on a VI 30-s schedule of point delivery. The VI schedule consisted of 50 intervals generated using the method described by Fleshler and Hoffman (1962). The intervals were randomized and were then presented sequentially both within and across sessions. Upon completion of an interval, the next response would result in point delivery. Point delivery was signaled by a brief tone and by incrementing a counter that was continuously displayed on the computer monitor throughout the block. The next interval in the sequence began timing immediately after point delivery.

*Linear variable-interval (LVI) 30 s.* Conditions were the same as in the VI 30-s condition except that, upon completion of an interval, a point was stored in a bank of available reinforcers and the next interval began timing immediately. If a press occurred and the value of the store was positive, one point was delivered and one was subtracted from the store. In the absence of responding, the number of available reinforcers in the store continued to escalate with the completion of each successive interval. Subjects were not informed of the value of the store, which carried across blocks and sessions, nor of a change in conditions.

*Negative slope (NS) FR* n. In this condition, points were added to the store according to an LVI 30-s schedule as described above, but every *n*th response subtracted one point from the store. If points were subtracted more frequently than they were set up by the LVI schedule, the store could be driven into a negative state. Both the value of the store and the position within the FR were carried across sessions. The subtraction ratios for individual subjects are listed in Table 1. These values were based upon the mean response rates from the last five sessions in the preceding condition, and were selected such that the overall rate of point delivery would be reduced by 50% if the steady-state response rates of the preceding condition prevailed. In the first NS condition for Subject 111, for example, the subtraction ratio value was 326, equal to the mean number of responses per minute from the last five sessions of the LVI 30-s condition. If this subject continued to respond at 326 responses per minute in the NS 326 condition, then one point per minute would be subtracted from the store according to the FR, while the LVI 30-s schedule was simultaneously adding two points per minute to the store. The resulting net reinforcement rate would be one point per minute, or approximately half the steady-state rate of point delivery from the LVI condition.

The subtraction ratios of Subjects 888 and 521 were determined slightly differently. The initial subtraction ratio was selected such that the overall rate of point delivery would be 0.5 per minute instead of 1.0 per minute. Both of these subjects showed little sensitivity to the NS schedules after several conditions. To produce decreases in response rates in the NS 58 and NS 47 conditions for Subject 888 and in the NS 139 and NS 117 conditions for Subject 521, subtraction ratios that reduced the rate of point delivery to approximately 0.0 per minute were used. In the final NS condition (NS 10) for Subject 888, the targeted rate of point delivery was −6.3 per minute.

*Conjoint VI 30 s (IRT < 2 s) random ratio (IRT < 2 s)*. For Subject 211, the rate of responding following the initial five-session exposure to the VI condition was insufficient to progress to the remaining conditions. To rectify this, the subject was exposed to a complex schedule of point delivery in which IRTs less than 2 s were reinforced according to random-ratio (RR) and VI 30-s schedules of point delivery, conjointly. On average, points were delivered following IRTs less than 2 s every 30 s by the VI schedule *and* after every *n*th such IRT by the RR schedule, where *n* equals the RR schedule value. The third 10-min block of the session served as a probe during which the RR schedule was suspended and points were scheduled solely by the VI 30-s schedule with the IRT requirement. The RR schedule value was initially set at 30, but was decreased across the first three sessions of the condition until response rates increased, indicating schedule control by a response-based contingency. The ratio requirement was then increased gradually over four sessions. This manipulation established high rates of responding when the ratio requirement was relatively low and eventually sustained those rates with reinforcement densities similar to that produced by efficient VI 30-s performance when the ratio requirement was relatively large. At that point, this schedule was discontinued and the VI condition was reinstated. To maintain overall session earnings similar to the other conditions, sessions were comprised of three blocks that terminated after 10 min or 20 point deliveries, whichever came first.

## RESULTS

Figure 4 shows the mean rate of point delivery versus the mean response rate from the last five sessions under all conditions for each subject. Rate of point delivery increased following the transition to the LVI from the VI conditions for all subjects, as indicated by a slight vertical displacement of the diamonds relative to the positions of the triangles. This result is in accord with the programmed changes in the contingencies across the two conditions, and occurred whether response

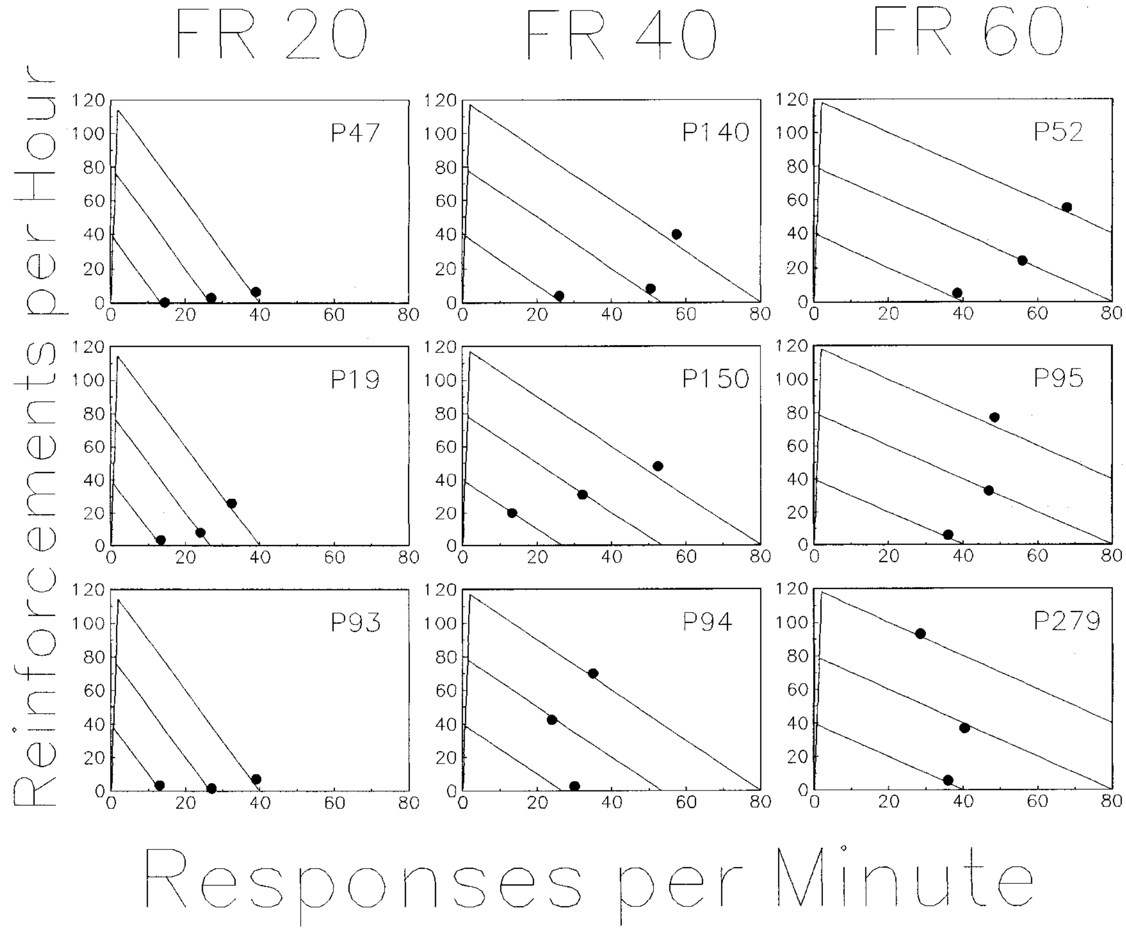

Fig. 8. Reinforcement rates versus response rates of pigeons under linear VI schedules with FR subtraction ratios of 20 (left panels), 40 (middle panels), and 60 (right panels). The data are superimposed on the molar feedback functions for the respective conditions. From Experiment 1 of Vaughan and Miller (1984, p. 341). Adapted by permission of the Society for the Experimental Analysis of Behavior.

rates increased (Subjects 111, 211, and 521), decreased (Subject 999), or remained nearly the same (Subject 888).

The remaining curves in Figure 4 are the molar feedback functions for the NS conditions. If reinforcement rate was consistently maximized, the points would converge at the apexes of the feedback functions, that is, at two responses per minute. Beyond this point, higher rates of responding result in decreasing rates of reinforcement until the rate of point delivery reaches a minimum of zero per minute. The response rates of all subjects across all NS conditions exceeded this value, producing deviations from the maximum obtainable reinforcement rate.

Sensitivity to changes in the contingencies can be assessed by comparing the obtained steady-state values to those encountered earlier in the condition. An appreciation of the situation confronting the subjects following a transition from one condition to the next can be achieved by drawing a vertical line from one of the obtained points in Figure 4 to the abscissa. The point at which this line intersects the feedback function of the subsequent condition will indicate the rate of reinforcement encountered immediately following the transition, assuming similar response rates are obtained. Any steady-state data that are plotted up and to the left of this intersection on the feedback function indicate a reduction in response rate and a corresponding increase in reinforcement rate. For example, the largest absolute change in response rate occurred with Subject 111 during exposure to the first NS condition (NS 326). A subtraction ratio of 326 was selected such that an overall reinforcement rate of 60 points per hour would be obtained if there were no changes in overall response rate. The rate decreased across the NS 326 condition from 326 responses per minute at the conclusion of the LVI condition to 41 responses per minute at the conclusion of the NS condition. The decrease in response rate resulted in a steady-state reinforcement rate that was similar to that obtained in the VI condition. The obtained point is plotted up and to the left of the targeted intersection (e.g., 326 responses per minute, 60 points per hour), indicating sensitivity to differences in the contingencies across the two conditions.

In Figure 4, complete insensitivity to changes in the FR requirement for reinforcement omission would be evidenced by a vertical arrangement of the points from the two conditions. In other words, the response rate would remain the same as the reinforcement rate changed across conditions according to the different FR values. Across subjects and NS conditions, the only instance of complete insensitivity to changes in the FR subtraction requirement occurred under NS 100 for Subject 888. The overall rate of responding under NS 100 was similar to that obtained under NS 136, resulting in a considerable decrease in overall reinforcement rate across the two conditions.

Overall response rate decreased relative to that of the previous condition in the remaining 22 NS conditions across subjects. Although these decreases indicate sensitivity to the contingencies, the changes in response rate seldom produced increases in the steady-state rates of point delivery across conditions. The steady-state rates of point delivery were lower than those of the previous conditions in 18 of 23 NS conditions across subjects. The five exceptions to this were Conditions NS 89 for Subject 888, NS 22 for Subject 999, NS 105 for Subject 211, and NS 139 and NS 117 for Subject 521. Steady-state rate of point delivery increased across consecutive NS conditions only once (Conditions NS 117 and NS 139 for Subject 521).

The data point obtained from Subject 888 under NS 10 deserves comment because it does not fall on the positive region of the molar feedback function for that condition. This subject was exposed to these contingencies for nine sessions and, despite receiving only one point in the first session of the condition, the response rate remained sufficiently high that the rate of point omission exceeded the rate of point availability. By the end of these nine sessions the value of the point store was −1,424. For the value of the store to again become positive, this subject would have had to cease responding entirely for 24 consecutive sessions. The subject left town at this point for a semester break and never returned to the experiment.

With the exception of Subject 999, the within-session response rates were similar across subjects. Portions of representative cumulative records from sessions conducted under VI, LVI, and NS conditions are shown in Figure 5 for Subject 211, selected to represent the other subjects. The leftmost segment (Panel 1) is from Session 5 during the first exposure to the VI condition. The slope of the line is shallow and roughly constant across the 10-min block, indicating a low but steady rate of responding. The segment in Panel 2 is from Session 17 during the second exposure to the VI condition (following the conjoint condition). Note that the slope of the line in this segment is steeper than that of the curve on the left, indicating a higher rate of responding. The next segment is from Session 23 during the LVI condition. The slope of this curve is similar to that from the preceding VI condition, indicating that there was little change in response rates across the two conditions.

Terminal performance under the NS schedules was characterized by alternating periods of high-rate responding and pausing. The development of this two-state response pattern can be seen in the remaining panels, taken from sessions conducted under NS conditions. In Panel 4 (NS 253, Session 30), occasional pauses interrupt periods of high-rate responding. In the subsequent condition (NS 138, Session 35), the high-rate patterns were similar, but pausing was more frequent. Panel 6 (NS 105, Session 63) shows response patterning that is characteristic of terminal NS schedule performance: Pauses alternate with periods of high-rate responding. The high-rate response runs frequently began with a series of point deliveries as the store was depleted and typically continued for many responses thereafter. Occasionally, point delivery occurred in the latter portion of these response runs. This run–pause pattern is even more pronounced and consistent in Panel 7 (NS 33, Session 91). Note that the pauses are more frequent and the runs are shorter.

The within-session response patterns of Subject 999 differed from those of the other subjects. In the VI condition, the overall rate of responding was lower than those of the other subjects, and there was an aperiodic alternation between accelerating and decelerating response rates, producing a wavy cumulative response curve. These characteristics persisted throughout the experiment. In the LVI condition, the overall rates decreased, and momentary fluctuations in response rate became even more pronounced. Under successive NS conditions, the overall rate of responding decreased and pauses became more frequent, but the momentary changes in response rate described above persisted. Only in the final NS condition was there any evidence of the possible development of a run–pause pattern of responding, but even here, this pattern was less pronounced than for the other subjects. Instead, sporadic pausing was often preceded by decelerating response rates and was followed by accelerating response rates.

The run–pause pattern of responding under NS conditions that was apparent in the cumulative records suggested that the IRT preceding reinforcement may be a useful measure to track across conditions. Figures 6 and 7 are scatter plots of the IRTs preceding point delivery across sessions for all subjects except 111. (Subject 111, the 1st subject studied in the experiment, was the subject whose data suggested the present measure.) Consistent with the two-state run–pause patterns of the cumulative records, by the conclusion of the NS conditions the distributions of IRTs preceding point delivery were bimodal for 3 of these 4 subjects (888, 211, and 521). For Subject 999, although there was an increase in the relative frequency of longer IRTs across the NS conditions, the development of a bimodal distribution was not as pronounced.

## DISCUSSION

The general pattern of results suggests that the behavior of all 5 subjects was sensitive to the contingencies of the negative slope conditions. This was evidenced by changes in the steady-state rates of responding and point delivery, in the within-session response patterns depicted in cumulative response records, and in the distributions of IRTs preceding point delivery. Despite clear-cut sensitivity to the contingencies, only rarely did reinforcement rates approach the maximum obtainable. The data from the negative slope conditions in Figure 4 do not lie near the apexes of the feedback functions for the respective schedules. Responding was therefore not in accord with strict versions of molar maximization based on maximizing overall rate of point delivery or minimizing the number of responses per point.

Earlier we entertained the possibility that negative slope schedules are not valid tests of maximization because local characteristics of responding can produce deviations from the theoretical functions. The simulations described earlier suggested that such deviations may be negligible. Although local response patterning can produce deviations from the theoretical functions, the deviations are usually quite small. Still one might argue that the simulations were based on unrealistic assumptions about responding under NS schedules. To test this possibility, we conducted several additional simulations with response characteristics derived from the run–pause pattern observed in our results. The details of these simulations can be found in Appendix B. To summarize briefly, as with the simulations based on a constant response probability, the simulations based on the run–pause pattern produced small but insignificant deviations from the theoretical functions. It seems reasonable to conclude from the simulation data that the feedback functions depicted in Figure 3 are not appreciably altered by local response patterning, even when such patterning is based on characteristics of actual responding.

Like the humans in the present study, the performance of nonhumans on NS schedules is not well characterized by accounts based on maximization of overall reinforcement rate. Figure 8 shows data from Vaughan and Miller's (1984) Experiment 1 with pigeons as subjects. In this experiment, each of three groups of pigeons was exposed to a different FR subtraction ratio on three different LVI schedules. The FR subtraction ratio was 20

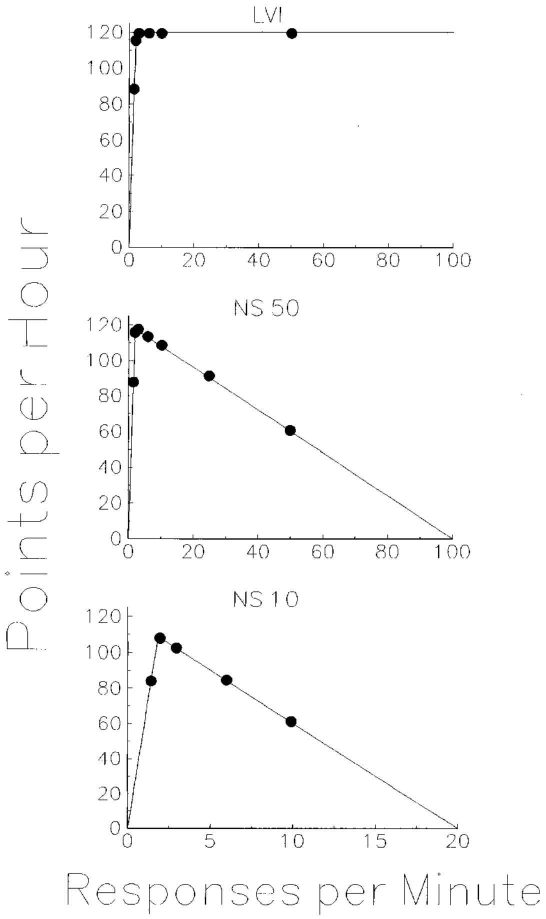

Fig. 9. Mean rate of point delivery versus mean response rate from the 20 simulations conducted at the LVI (top panel), NS FR 50 (middle panel), and NS FR 10 (bottom panel) conditions. The data are superimposed on the molar feedback functions for each condition. Note individually scaled axes.

for one group, 40 for a second group, and 60 for a third group. Steady-state response rates from the NS conditions are superimposed upon the molar feedback functions for the respective schedules. The response rates differed substantially from the inverse of the LVI values, and the reinforcement rates frequently approached zero per minute. Despite substantial experience with these contingencies, low response rates were not obtained, and overall rate of food delivery was not maximized. These results are also consistent with those obtained with rats as subjects (Ettinger et al., 1987; Reed & Schachtman, 1991).

Comparing our steady-state data (Figure 4) to Figure 8 suggests that the performances of humans are somewhat similar to those of pigeons. The relevant data for comparison purposes are the filled circles plotted on the uppermost curves in each of the nine panels of Figure 8. These are the data from conditions in which the LVI schedule value was 30 s at subtraction ratios of 20, 40, and 60, from left to right, respectively. Although interpretation is limited by differences in experimental design (Vaughan & Miller, 1984, manipulated subtraction ratio between subjects, whereas the present study manipulated it within subjects), the data in Figure 8 suggest that rates of responding and reinforcement were inversely related to the subtraction ratio value. This was frequently observed with our human subjects as well (Figure 4). Thus, although it was inconsistent with strict versions of molar maximization, responding both in our humans and in Vaughan and Miller's pigeons was sensitive to NS contingencies.

Vaughan and Miller's (1984) data are at least broadly consistent with models based on a hyperbolic function relating response rate to reinforcement rate, such as the single-alternative form of the matching law (Herrnstein, 1970) and linear system theory (McDowell & Wood, 1984). The filled circles in Figure 8 show the obtained response rates versus the obtained reinforcement rates from NS conditions. Although none of the pigeons were exposed to enough conditions to discern a hyperbolic relation, response rates varied directly with reinforcement rates for 6 of the 9 birds, an outcome that is consistent (although not uniquely so) with these accounts. By contrast, the response rates of humans in the present study (with the possible exception of Subject 888) were not well characterized by a hyperbolic function.

Such deviations from the hyperbolic form are in contrast with the results of several studies conducted with human subjects by Bradshaw, Szabadi, and Bevan (1976, 1977, 1979). Two of these studies (1977, 1979), in particular, deserve mention because, as in the NS conditions of the present study, higher response rates resulted in lower reinforcement rates. In these experiments, button pressing was maintained by points exchangeable for money. Point deliveries were arranged by a multiple schedule comprised of five different VI component schedules. Point delivery was signaled by a 100-ms illumination of a green

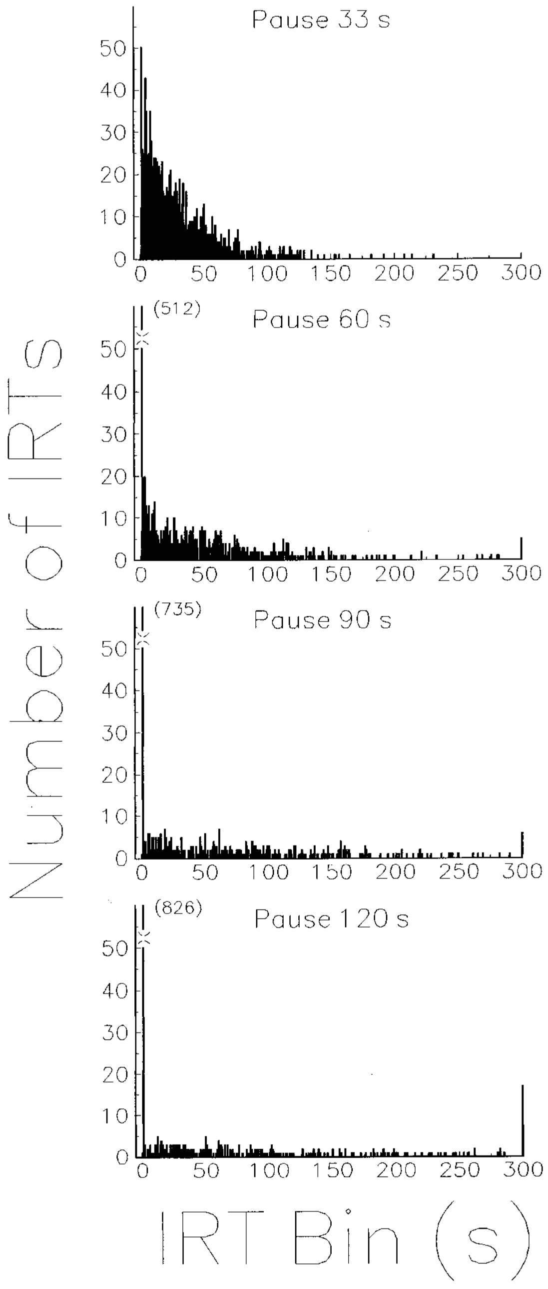

Fig. 10. IRT histograms from simulation sessions in which the programmed mean pause length was 33, 60, 90, and 120 s (top to bottom). Note logarithmic axes.

light and the addition of one point to a counter. Under some conditions, a variable-ratio 34 schedule of point loss was in effect conjointly with the VI schedules. Every 34 responses, on average, produced point loss, signaled by a 100-ms illumination of a red lamp and the deduction of one point from the counter. As in the NS conditions of the present study, response rates in the punishment conditions of the Bradshaw et al. studies (1977, 1979) were lower than baseline rates. Unlike the results of the present experiment, however, response rates in the punishment conditions varied as a function of reinforcement rate in a manner that was consistent with the hyperbolic function underlying both the matching law and linear system theory.

What accounts for the differences between our findings and those of Bradshaw et al. (1977, 1979)? An important difference between studies lies in the discriminability of the contingencies. In Bradshaw's experiments, point losses were signaled, whereas in our experiment they were not. Moreover, unlike our experiment, the schedule-correlated stimuli in Bradshaw's experiments bore an ordinal relation to schedule value: The light correlated with the richest schedule was on the extreme right and that correlated with the leanest schedule was on the extreme left, with a graded series of schedule values in between. In a systematic replication of procedures from Bradshaw's laboratory involving concurrent reinforcement schedules, Horne and Lowe (1993) found that matching was frequently obtained when implicit or explicit ordinal cues were used as discriminative stimuli, but significant deviations from matching occurred when different geometrical shapes were projected onto a response panel as discriminative stimuli. Together, these findings join with others (Madden & Perone, 1999) in showing that matching is more likely when discrimination of schedule contingencies is enhanced. Perhaps if explicit contingency-related stimuli had been used, the present results would be in better accord with the predictions of matching (or with maximizing, for that matter; Heyman & Tanz, 1995). Nonetheless, the present results provide much-needed data on the extent and range of humans' sensitivity to reinforcement variables that occur in the absence of clear contingency-related stimuli.

Apart from questions of matching or maximizing, the present results provide important data on the local characteristics of performance on negative slope schedules. For Subjects 111, 888, 211, and 521, constant, moderate rates of responding established in the VI and LVI conditions gave way under NS conditions to a consistent two-state pattern with periods of pausing alternating with bursts of high-rate responding. The onset of responding was usually met with a series of point deliveries until the store was depleted. Responding often continued after the store was depleted, and additional points were occasionally delivered. This pattern differs considerably from that typical of VI performance, which is characterized by responding at a constant rate with little pausing. The changes in the temporal distribution of responses depicted in the cumulative records were also reflected in changes in the distributions of IRTs preceding point delivery. The distributions of 3 of the 4 subjects (211, 888, and 521) for which this measure was recorded were bimodal, with one cluster of short IRTs (less than 1 s) and another cluster of longer IRTs (greater than 45 s). For Subject 999, the distributions of IRTs preceding point delivery were not bimodal, but the frequency of relatively long IRTs did increase under the NS conditions.

The bimodal distribution of reinforced IRTs is consistent with the suggestion made by Reed and Schachtman (1991) that response-rate differences on NS schedules and LVI schedules with similar reinforcement rates may be due to the differential reinforcement of relatively longer IRTs under the NS contingencies. In their experiment, interreinforcement intervals were yoked across NS and LVI components of a multiple schedule. Under some conditions, they also required that the IRT preceding reinforcement in the yoked component fall within 1 s of that obtained in the NS component. Overall response rates were more similar across the two components when the IRT requirement was present than when it was absent, suggesting that NS schedules differentially reinforce longer IRTs. In the present study, consistent with Reed and Schachtman's hypothesis, the distribution of IRTs preceding point delivery became bimodal and the range increased under the NS conditions (Figures 6 and 7) in 3 of these 4 subjects.

This bimodal response distribution appears to be characteristic of NS schedules, but much about the fine-grained structure of NS performance is still in question. Is the run–pause pattern itself an integrated unit, or is it the aggregate effect of reinforcement acting on different response tendencies? In the present procedure, overall reinforcement rate favors longer IRTs, but reinforcer immediacy favors shorter IRTs. If one assumes that the value of individual reinforcers diminishes with delay (Mazur, 1987), then it may not be too surprising that the rate-decreasing effects of overall reinforcement rate on NS schedules may be at least partially overridden by the rate-increasing effects of reinforcement immediacy (Silberberg, Warren-Boulton, & Asano, 1988). The run–pause response pattern may thus reflect trade-offs between these two competing response tendencies—bursting controlled by reinforcement immediacy and pausing controlled by overall reinforcement rate. Models that take into account both local and overall reinforcement variables appear to be necessary to predict response output on negative slope schedules.

## REFERENCES

Baum, W. M. (1973). The correlation-based law of effect. *Journal of the Experimental Analysis of Behavior, 20,* 137–153.

Baum, W. M. (1992). In search of the feedback function for variable-interval schedules. *Journal of the Experimental Analysis of Behavior, 57,* 365–375.

Bradshaw, C. M., Szabadi, E., & Bevan, P. (1976). Behavior of humans in variable-interval schedules of reinforcement. *Journal of the Experimental Analysis of Behavior, 26,* 135–141.

Bradshaw, C. M., Szabadi, E., & Bevan, P. (1977). Effect of punishment on human variable-interval performance. *Journal of the Experimental Analysis of Behavior, 27,* 275–279.

Bradshaw, C. M., Szabadi, E., & Bevan, P. (1979). The effect of punishment on free-operant choice in humans. *Journal of the Experimental Analysis of Behavior, 31,* 71–81.

Ettinger, R. H., Reid, A. K., & Staddon, J. E. R. (1987). Sensitivity to molar feedback functions: A test of molar optimality theory. *Journal of Experimental Psychology: Animal Behavior Processes, 13,* 366–375.

Fleshler, M., & Hoffman, H. S. (1962). A progression for generating variable-interval schedules. *Journal of the Experimental Analysis of Behavior, 5,* 529–530.

Herrnstein, R. J. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior, 13,* 243–266.

Herrnstein, R. J., Loewenstein, G. F., Prelec, D., & Vaughan, W., Jr. (1993). Utility maximization and melioration: Internalities in individual choice. *Journal of Behavioral Decision Making, 6,* 149–185.

Heyman, G. M., & Tanz, L. (1995). How to teach a pigeon to maximize overall reinforcement rate. *Journal of the Experimental Analysis of Behavior, 64,* 277–297.

Horne, P. J., & Lowe, C. F. (1993). Determinants of human performance on concurrent schedules. *Journal of the Experimental Analysis of Behavior, 59,* 29–60.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47,* 263–291.

Kudadjie-Gyamfi, E., & Rachlin, H. (1996). Temporal patterning in choice among delayed outcomes. *Organizational Behavior and Human Decision Processes, 65,* 61–67.

Madden, G. J., & Perone, M. (1999). Human sensitivity to concurrent schedules of reinforcement: Effects of observing schedule-correlated stimuli. *Journal of the Experimental Analysis of Behavior, 71,* 303–318.

Mazur, J. E. (1987). An adjusting procedure for studying delayed reinforcement. In M. L. Commons, J. E. Mazur, J. A. Nevin, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value* (pp. 55–73). Hillsdale, NJ: Erlbaum.

McDowell, J. J, & Wood, H. M. (1984). Confirmation of linear system theory prediction: Changes in Herrnstein's *k* as a function of changes in reinforcer magnitude. *Journal of the Experimental Analysis of Behavior, 41,* 183–192.

Rachlin, H., Raineri, A., & Cross, D. (1991). Subjective probability and delay. *Journal of the Experimental Analysis of Behavior, 55,* 233–244.

Reed, P., & Schachtman, T. R. (1991). Instrumental performance on negative schedules. *The Quarterly Journal of Experimental Psychology, 43,* 177–197.

Silberberg, A., Warren-Boulton, F. R., & Asano, T. (1988). Maximizing present value: A model to explain why moderate response rates obtain on variable-interval schedules. *Journal of the Experimental Analysis of Behavior, 49,* 331–338.

Vaughan, W., Jr. (1982). Choice and the Rescorla-Wagner model. In M. L. Commons, R. J. Herrnstein, & H. Rachlin (Eds.), *Quantitative analyses of behavior: Vol. 2. Matching and maximizing accounts* (pp. 263–279). Cambridge, MA: Ballinger.

Vaughan, W., Jr., & Miller, H. L., Jr. (1984). Optimization versus response-strength accounts of behavior. *Journal of the Experimental Analysis of Behavior, 42,* 337–348.

*Received February 2, 1998*
*Final acceptance March 23, 2000*

## APPENDIX A

To evaluate the accuracy of the feedback functions depicted here, we conducted simulations for LVI 30-s, NS FR 10, and NS FR 50 conditions. The software used to control the contingencies and record responses in the experiment proper was modified to simulate the occurrence of responses. Responses were programmed to occur with a constant probability per unit time. The probability was varied to produce overall response rates ranging from 1.5 per minute to 50 per minute. Twenty sessions were conducted at each schedule–response probability combination. The store value carried across sessions within conditions and reset to zero between conditions.

Programmed and obtained response and point rates for each combination of schedule and response probability are presented in Table 2. The values shown are means of data from the 20 sessions conducted at each condition. Across all conditions, the obtained response rates were within 0.24 responses per minute of the programmed overall response rate, indicating the simulated responses occurred appropriately.

The predicted point rates presented in Table 2 were based upon the obtained response rates. For LVI conditions, predicted point rate equaled the obtained response rate for all response rates less than two per minute and equaled two per minute for all response rates greater than two per minute. For NS conditions, the response rate ($R$) at which the apex of the function occurs was determined by

$$R = L/[1 + (1/F)], \qquad (1)$$

where $L$ equals the programmed rate of re-

Table 2

Mean obtained response rates, mean obtained point rates, and programmed point rates from 20 simulations conducted at each condition.

| Schedule | Programmed response rate (responses per minute) | Obtained response rate (responses per minute) | Programmed point rate (points per hour) | Obtained point rate (points per hour) |
|---|---|---|---|---|
| LVI 30 s | 1.5 | 1.47 | 88.2 | 88.2 |
| | 2 | 1.93 | 115.8 | 115.2 |
| | 3 | 2.95 | 120.0 | 119.4 |
| | 6 | 6.05 | 120.0 | 119.4 |
| | 10 | 9.97 | 120.0 | 119.4 |
| | 50 | 50.05 | 120.0 | 119.4 |
| NS FR 10 | 1.5 | 1.40 | 84.0 | 84.0 |
| | 2 | 1.95 | 108.0 | 108.0 |
| | 3 | 2.93 | 102.6 | 102.6 |
| | 6 | 5.97 | 84.0 | 84.6 |
| | 10 | 9.88 | 60.6 | 61.2 |
| NS FR 50 | 1.5 | 1.46 | 87.6 | 87.6 |
| | 2 | 1.94 | 116.4 | 115.8 |
| | 3 | 2.94 | 116.4 | 117.6 |
| | 6 | 5.89 | 112.8 | 113.4 |
| | 10 | 10.13 | 108.0 | 108.6 |
| | 25 | 24.81 | 90.0 | 91.2 |
| | 50 | 49.76 | 60.0 | 60.6 |

Table 3

The number and percentage of total responses unreinforced from sessions in which the obtained response rate is less than *R*. See text for details.

| Condition | Programmed response rate (responses per minute) | *R* | Sessions | Responses | Unreinforced responses | % of total responses unreinforced |
|---|---|---|---|---|---|---|
| LVI 30 s | 1.5 | 2.00 | 20 | 885 | 2 | 0.23 |
| | 2.0 | 2.00 | 12 | 636 | 5 | 0.79 |
| NS FR 10 | 1.5 | 1.82 | 20 | 841 | 1 | 0.12 |
| | 2.0 | 1.82 | 5 | 243 | 6 | 2.47 |
| NS FR 50 | 1.5 | 1.96 | 20 | 876 | 2 | 0.23 |
| | 2.0 | 1.96 | 9 | 474 | 2 | 0.42 |
| Overall | | | 86 | 3,955 | 18 | 0.46 |

inforcement (in points per minute) arranged by the LVI schedule and *F* equals the subtraction ratio. The predicted point rate equaled the obtained response rate for all response rates less than or equal to *R*. For response rates greater than *R*, predicted point rates (*P*, in points per minute) were determined by the following formula:

$$P = L - (B/F), \tag{2}$$

where *B* is the overall rate of responding in responses per minute and *L* and *F* are, again, the programmed rate of reinforcement arranged by the LVI schedule and the subtraction ratio value, respectively. Across all schedule and response probability combinations, the obtained mean point rates were within 1.2 points per hour of the predicted value.

Figure 9 shows the mean obtained point rates versus the mean obtained response rates for the LVI, NS FR 50, and NS FR 10 conditions. The molar feedback functions are included for reference. The means for each schedule–response probability combination fell along the molar feedback functions.

Although the condition means accorded well with the predicted values, molecular interactions between the obtained sequence of IRTs and the schedule may have produced deviations from the predicted functions as described above (see Figure 3). To assess the extent of these deviations in the simulation data, we examined the data from all sessions in which the predicted point rates fell along the FR 1 section of the molar feedback function. That is, we examined data from all sessions in which the obtained session-wide response rate was less than *R*. The occurrence of unreinforced responses during these sessions would cause the response rate to exceed the reinforcement rate, creating deviations from the predicted function.

The number of responses and percentage of total responses that did not produce a point are presented in Table 3 by condition. The number of sessions in which the obtained response rate was less than *R* and the total number of responses simulated in those sessions are also listed in Table 3 for each condition. Overall, responses that did not produce points were rare in these sessions. The obtained response rate was less than *R* in 86 sessions across conditions. In those 86 sessions, 3,955 responses were simulated, of which only 18 (0.46%) did not produce points. As expected, the frequency of responses not producing points tended to be higher in conditions in which the programmed response rate approached *R*.

## APPENDIX B

To ascertain whether the run–pause pattern of responding obtained in the study would produce significant departures from the depicted molar feedback functions, we conducted a series of simulations under NS FR 10 conditions. The software used to control events and record responses in the experiment proper was modified to simulate the occurrence of responses. Instead of arranging responses randomly, however, runs of responses were programmed to follow pauses. Across all simulations, the rate at which the apex of the NS FR 10 feedback function occurred ($R = 1.82$ responses per minute) was targeted. Pauses terminated (i.e., runs initiated) with a constant probability per unit

Table 4

Mean programmed and obtained pause lengths, response rate, point rate, and number and percentage of unreinforced responses from 20 sessions conducted at programmed pause lengths of 33, 66, 90, and 120 s.

| Programmed pause length (s) | Obtained pause length (s) | Response rate (responses per minute) | Point rate (points per hour) | Responses | Unreinforced responses | % of total responses unreinforced |
|---|---|---|---|---|---|---|
| 33 | 32.19 | 1.84 | 106.8 | 1,102 | 36 | 3.4 |
| 60 | 55.28 | 1.82 | 108.6 | 1,090 | 4 | 0.37 |
| 90 | 82.39 | 1.82 | 108.6 | 1,090 | 2 | 0.18 |
| 120 | 104.25 | 1.82 | 108.6 | 1,091 | 3 | 0.27 |
| Overall | | | | 4,373 | 45 | 1.03 |

time. Once a response run was initiated, responses occurred at a fixed tempo (IRT = 0.3 s) until the overall response rate equaled or exceeded *R*. Thus, overall response rate was controlled in a homeostatic fashion; pauses tended to reduce the overall response rate below *R*, but the ensuing run would persist until the overall response rate was equal to or slightly greater than *R*. The probability of pause termination per unit time was varied across simulations to target mean pause lengths of 33, 60, 90, and 120 s. Twenty sessions were simulated at each pause length.

The number of responses simulated, the number of unreinforced responses, and the percentage of unreinforced responses are shown in Table 4. The obtained mean pause lengths, response rates, and point rates are also shown. Obtained pause lengths tended to be less than programmed values because sessions typically ended during a pause, and this time was not included in the calculation. As before, obtained response rates closely approximated programmed values. Obtained overall response rate equaled the programmed rate in all but one of the conditions. In the 33-s pause condition, the overall rate was within 0.02 responses per minute of the targeted value. Of the 4,373 responses simulated across the four pause conditions, only 45 (1.03%) did not produce points. Thus, as in the earlier simulations, the vast majority of simulated responses produced points, and the resulting overall reinforcement rates were not significantly affected by the occurrence of the unreinforced responses. Across all pause conditions, overall reinforcement rates were within 2.4 points per hour of the predicted values.

The run–pause response pattern, therefore, does not appear to increase the likelihood or magnitude of deviations from the negative-slope molar feedback functions as depicted. In fact, the run–pause patterning may actually decrease such deviations. Of the 45 unreinforced responses, 36 (80%) occurred under the 33-s pause condition. Because the programmed pause length equaled the inverse of *R*, response patterning under this condition was similar to that obtained in the constant-probability simulations.

Figure 10 shows histograms of the IRTs from each of the four pause conditions. Across pause conditions, the frequency of very short (less than 1 s) and very long (greater than 100 s) IRTs increased, whereas the frequency of IRTs of intermediate length (between 1 and 100 s) decreased. A relative preponderance of intermediate IRTs, evident in the distribution for the 33-s pause condition, is likely to enhance the occurrence of unreinforced responses due to the greater likelihood of a response occurring before a schedule interval has timed out. As the run–pause pattern becomes better defined, points tend to be stored during pauses and collected during runs. Thus, relative to random responding, run–pause patterning decreases the likelihood of unreinforced responses, reducing deviations from the depicted molar feedback functions at response rates close to *R*.

In conclusion, data from both series of simulations suggest that although deviations from the proposed functions due to local variability in responding are possible, such deviations are negligible. The proposed molar feedback functions therefore appear to be adequate for evaluating sensitivity to overall reinforcement rate.